Logics for Non-Cooperative Games with Expectations

Lluís Godo[1], Enrico Marchioni[2]

[1] IIIA, Artificial Intelligence Research Institute
CSIC, Spanish National Research Council
Campus de la Univ. Autònoma de Barcelona s/n
08193 Bellaterra, Spain
Email: godo@iiia.csic.es

[2] Institut de Recherche en Informatique de Toulouse
Université Paul Sabatier, 118 Route de Narbonne
31062 Toulouse, France
Email: enrico.marchioni@irit.fr

Abstract. We introduce the logics $\mathsf{E}(\mathfrak{G})$ for reasoning about probabilistic expectation over classes $\mathfrak{G}$ of games with discrete polynomial payoff functions represented by finite-valued Łukasiewicz formulas and provide completeness and complexity results. In addition, we introduce a new class of games where players' expected payoff functions are encoded by $\mathsf{E}(\mathfrak{G})$-formulas. In these games each player's aim is to randomise her strategic choices in order to affect the other players' expectations over an outcome as well as their own. We offer a logical and computational characterisation of this new class of games.

1 Introduction

In this work, we introduce the logics $\mathsf{E}(\mathfrak{G})$ for reasoning about probabilistic expectation over classes $\mathfrak{G}$ of games with discrete polynomial payoff functions represented by finite-valued Łukasiewicz formulas. This type of games, called Łukasiewicz games [14], is a generalisation of Boolean games [1, 12] and involves a finite set of players $\mathsf{P} = \{P_1, \ldots, P_n\}$, each controlling a finite set of propositional variables V_i, so that the sets V_i are mutually disjoint. Being in control of a set V_i of propositional variables means that P_i assigns to the variables in V_i values from the scale

$$L_k = \left\{0, \frac{1}{k}, \ldots, \frac{k-1}{k}, 1\right\}.$$

A strategy for a player P_i is a function $s : \mathsf{V}_i \to L_k$ that corresponds to a valuation of the variables controlled by P_i. Strategies can be interpreted as efforts or costs, and each player's strategic choice can be seen as an assignment to each controlled variable carrying an intrinsic cost.

Each player is assigned a finite-valued Łukasiewicz formula φ_i, with propositional variables from $\bigcup_i^n \mathsf{V}_i$, whose valuation is interpreted as the payoff function

for P_i and corresponds to the restriction over L_k of a continuous piecewise linear polynomial function with integer coefficients. Notice that not all variables in φ_i might be under P_i's control and, consequently, P_i's payoff from playing a certain strategy (i.e. choosing a certain variable assignment) also depends on the choices made by (some of) the other players.

In this work, we expand Łukasiewicz games by providing an explicit definition of a class $\mathfrak{G}$ of games, and also defining suitable concepts of a mixed strategy, best response and Nash Equilibrium, adapted to this framework. Then, we define a probabilistic logic $\mathsf{E}(\mathfrak{G})$ over many-valued formulas (of finite-valued Łukasiewicz logic) to reason about expectations in a class $\mathfrak{G}$ of Łukasiewicz games. This might be in principle a bit surprising since probabilities are not the same as expectations, but the reason is simple. The generalisation of a classical probability measure on Boolean algebras to the algebraic setting of MV-algebras [3] (the algebraic counterpart of Łukasiewicz logics) corresponds to the so-called *states* [16], which can be seen as averages of truth-values. Indeed, a state (or finitely-additive probability) σ over the set of Łukasiewicz logic formulas $\mathcal{L}$ (built from a given set of propositional variables V) is a mapping $\sigma : \mathcal{L} \to [0,1]$ such that:

- $\sigma(\overline{1}) = 1$,
- $\sigma(\varphi \oplus \psi) = \sigma(\varphi) + \sigma(\psi)$, if $\neg(\varphi \& \psi)$ is a Łukasiewicz tautology,
- $\sigma(\varphi) = \sigma(\psi)$, if $\varphi \leftrightarrow \psi$ is a Łukasiewicz tautology.

When we consider only finite-valued Łukasiewicz logics L_k, as proved in [18], a mapping $\sigma : \mathcal{L} \to [0,1]$ is a state on formulas iff there exists a probability distribution $\pi : \Omega_k \to [0,1]$ on the set of L_k-valuations Ω_k on $\mathcal{L}$ such that

$$\sigma(\varphi) = \sum_{w \in \Omega_k} p(w) \cdot w(\varphi).$$

The state $\sigma(\varphi)$ corresponds to a weighted average of the truth-values of φ under all possible valuations, and it can also be regarded as the expected value of the function $f_\varphi : \Omega_k \to [0,1]$, defined as $f_\varphi(w) = w(\varphi)$ for all $w \in \Omega_k$. If we look at Łukasiewicz formulas as a particular class of functions, to reason about the probability of these formulas amounts to reasoning about the expectation of the corresponding functions. This is the view we take in this paper.

Note that a logic, called $FP(\mathrm{L}_n, \mathrm{L})$, to reason about the probability of L_k-formulas was defined and axiomatised in [6]. The logic $\mathsf{E}(\mathfrak{G})$ introduced here is a (partially) syntactical and semantic expansion of $FP(\mathrm{L}_n, \mathrm{L})$, and its expressive power makes it possible to represent expectations in games where players' payoffs are given by discrete polynomial functions with integer coefficients over L_k (encoded by finite-valued Łukasiewicz formulas).

$\mathsf{E}(\mathfrak{G})$ is built from a set of non-modal formulas $\varphi, \psi, \dots$ that correspond to formulas of the finite-valued Łukasiewicz logic L_k^c (with additional truth constants for each element $c \in L_k$), and a set of modal formulas $\mathsf{E}\varphi, \mathsf{E}\psi, \dots$ to represent the expectation associated to each $\varphi, \psi, \dots$. Modal formulas are combined with the connectives of the $\mathrm{L}\Pi\frac{1}{2}$-logic [5], which makes it possible to express combinations of polynomial equalities and inequalities with rational coefficients.

Based on the logics $\mathsf{E}(\mathfrak{G})$, we also introduce a new class of games, called *Łukasiewicz games with expectations*, that generalise Łukasiewicz games. In the situations of strategic interactions modelled in Game Theory, the goal of each player is essentially the maximisation of her own expected payoff. Players, however, often care not only about maximising their own expectation, but also about influencing other players' expected outcomes. Łukasiewicz games with expectations offer a formalisation of these kinds of strategic interactions where each player's aim is to randomise her strategic choices in order to affect the other players' expectations over an outcome as well as their own expectation. Łukasiewicz games with expectations expand Łukasiewicz games by assigning to each player P_i a modal formula Φ_i of $\mathsf{E}(\mathfrak{G})$, whose interpretation corresponds to a piecewise rational polynomial function whose variables are interpreted as expected values. The modal formula Φ_i assigned to each player is then meant to represent a player's goal concerning the relation between her and other players' expectations.

This work is organised as follows.[1] The next section introduces the basic background notions about Łukasiewicz logics and the ŁΠ$\frac{1}{2}$-logic. In Section 3, we present the main definition of a Łukasiewicz game, define the concept of a class $\mathfrak{G}$ of games and build the logics $\mathsf{E}(\mathfrak{G})$ to represent expectations in each $\mathfrak{G}$. We provide both completeness and complexity results for $\mathsf{E}(\mathfrak{G})$. In Section 4, we introduce games with expectations based on $\mathsf{E}(\mathfrak{G})$ and offer some examples along with a logical and computational characterisation. We end with some final remarks.

2 Logical Background

The language of Łukasiewicz logic Ł (see [3]) is built from a countable set of propositional variables $\{p_1, p_2, \ldots\}$, the binary connective $\rightarrow$ and the truth constant $\bar{0}$ (for falsity). Further connectives are defined as follows:

$$\begin{array}{ll}
\neg\varphi \text{ is } \varphi \rightarrow \bar{0}, & \varphi \wedge \psi \text{ is } \varphi \& (\varphi \rightarrow \psi), \\
\varphi \& \psi \text{ is } \neg(\varphi \rightarrow \neg\psi), & \varphi \vee \psi \text{ is } ((\varphi \rightarrow \psi) \rightarrow \psi), \\
\varphi \oplus \psi \text{ is } \neg(\neg\varphi \& \neg\psi), & \varphi \leftrightarrow \psi \text{ is } (\varphi \rightarrow \psi) \& (\psi \rightarrow \varphi), \\
\varphi \ominus \psi \text{ is } \varphi \& \neg\psi, & d(\varphi, \psi) \text{ is } \neg(\varphi \leftrightarrow \psi).
\end{array}$$

Let *Form* denote the set of Łukasiewicz logic formulas. A valuation e from *Form* into $[0, 1]$ is a mapping $e : Form \rightarrow [0, 1]$ assigning to all propositional variables a value from the real unit interval (with $e(\bar{0}) = 0$) that can be extended to complex formulas as follows:

$$\begin{array}{ll}
e(\varphi \rightarrow \psi) = \min(1 - e(\varphi) + e(\psi), 1) & e(\neg\varphi) = 1 - e(\varphi) \\
e(\varphi \& \psi) = \max(0, e(\varphi) + e(\psi) - 1) & e(\varphi \oplus \psi) = \min(1, e(\varphi) + e(\psi)) \\
e(\varphi \ominus \psi) = \max(0, e(\varphi) - e(\psi)) & e(\varphi \wedge \psi) = \min(e(\varphi), e(\psi)) \\
e(\varphi \vee \psi) = \max(e(\varphi), e(\psi)) & e(d(\varphi, \psi)) = |e(\varphi) - e(\psi)| \\
e(\varphi \leftrightarrow \psi) = 1 - |e(\varphi) - e(\psi)| &
\end{array}$$

---

[1] Notice that the proofs of the main results are either simply sketched or omitted due to space constraints.

A valuation e *satisfies* a formula φ if $e(\varphi) = 1$. As usual, a set of formulas is called a theory. A valuation e satisfies a theory T, if $e(\psi) = 1$, for every $\psi \in T$.

Infinite-valued Łukasiewicz logic has the following axiomatisation:

(Ł1) $\varphi \to (\psi \to \varphi)$,   (Ł2) $(\varphi \to \psi) \to ((\psi \to \chi) \to (\varphi \to \chi))$,

(Ł3) $(\neg\varphi \to \neg\psi) \to (\psi \to \varphi)$, (Ł4) $((\varphi \to \psi) \to \psi) \to ((\psi \to \varphi) \to \varphi)$.

The only inference rule is *modus ponens*, i.e.: from $\varphi \to \psi$ and φ derive ψ.

A *proof* in Ł is a sequence $\varphi_1, \ldots, \varphi_n$ of formulas such that each φ_i either is an axiom of Ł or follows from some preceding φ_j, φ_k ($j, k < i$) by modus ponens. We say that a formula φ can be derived from a theory T, denoted as $T \vdash \varphi$, if there is a proof of φ from a set $T' \subseteq T$. A theory T is said to be consistent if $T \nvdash \bar{0}$.

Łukasiewicz logic is complete with respect to deductions from finite theories for the given semantics, i.e.: for every finite theory T and every formula φ, $T \vdash \varphi$ iff every valuation e that satisfies T also satisfies φ.

For each $k \in \mathbb{N}$, the finite-valued Łukasiewicz logic $Ł_k$ is the schematic extension of Ł with the axiom schemas:

(Ł5) $(n-1)\varphi \leftrightarrow n\varphi$,   (Ł6) $(k\varphi^{k-1})^n \leftrightarrow n\varphi^k$,

for each integer $k = 2, \ldots, n-2$ that does not divide $n-1$, and where $n\varphi$ is an abbreviation for $\varphi \oplus \cdots \oplus \varphi$ (n times) and φ^k is an abbreviation for $\varphi \& \ldots \& \varphi$, ($k$ times). The notions of valuation and satisfiability for $Ł_k$ are defined as above just replacing $[0, 1]$ by

$$L_k = \left\{ 0, \frac{1}{k}, \ldots, \frac{k-1}{k}, 1 \right\}$$

as set of truth values. Every $Ł_k$ is complete with respect to deductions from finite theories for the given semantics.

It is sometimes useful to introduce constants in addition to $\bar{0}$ that will denote values in the domain L_k. Specifically, we will denote by $Ł_k^c$ the Łukasiewicz logic obtained by adding constants $\bar{c}$ for every value $c \in L_k$. We assume that valuation functions e interpret such constants in the natural way: $e(\bar{c}) = c$.

A McNaughton function [3] is a continuous piecewise linear polynomial functions with integer coefficients over the nth-cube $[0, 1]^n$. To each Łukasiewicz formula $\varphi(p_1, \ldots, p_n)$ we can associate a McNaughton function f_φ so that, for every valuation e

$$f_\varphi(e(p_1), \ldots, e(p_n)) = e(\varphi(p_1, \ldots, p_n)).$$

Every Ł-formula is then said to define a McNaughton function. The converse is also true, i.e. every continuous piecewise linear polynomial function with integer coefficients over $[0, 1]^n$ is definable by a formula in Łukasiewicz logic. In the case of finite-valued Łukasiewicz logics, the functions defined by formulas are just the restrictions of McNaughton functions over $(L_k)^n$. In this sense, we can associate

to every formula $\varphi(p_1, \ldots, p_n)$ from Ł$_k$ a function $f_\varphi : (L_k)^n \to L_k$. As for each Ł$_k^c$, the functions defined by a formula are combinations of restrictions of McNaughton functions and, in addition, the constant functions for each $c \in L_k$. The class of functions definable by Ł$_k^c$-formulas exactly coincides with the class of all functions $f : (L_k)^n \to L_k$, for every $n \geq 0$.

The expressive power of infinite-valued Łukasiewicz logic lies in, and is limited to, the definability of piecewise linear polynomial functions. Expanding Ł with the connectives $\odot, \to_\Pi$ of Product logic [10], interpreted as the product of reals and as the truncated division, respectively, significantly augments the expressive power of the logic. The ŁΠ$\frac{1}{2}$ logic [5] is the result of this expansion, obtained by adding the connectives $\odot, \to_\Pi, \overline{\frac{1}{2}}$, whose valuations e extend the valuations for Ł as follows:

$$e(\varphi \odot \psi) = e(\varphi) \cdot e(\psi) \qquad e(\varphi \to_\Pi \psi) = \begin{cases} 1 & e(\varphi) \leq e(\psi) \\ \frac{e(\psi)}{e\varphi} & \text{otherwise} \end{cases}$$

$$e\left(\overline{\frac{1}{2}}\right) = \frac{1}{2} \qquad\qquad e(\Delta\varphi) = \begin{cases} 1 & e(\varphi) = 1 \\ 0 & \text{otherwise} \end{cases}$$

Notice that the presence of the constant $\overline{\frac{1}{2}}$ makes it possible to define constants for all rationals in $[0,1]$ (see [5]). ŁΠ$\frac{1}{2}$'s axioms include the axioms of Łukasiewicz and Product logics (see [10]) as well as the following additional axioms, where $\Delta\varphi$ is $\neg\varphi \to_\Pi \overline{0}$:

(ŁΠ1) $(\varphi \odot \psi) \ominus (\varphi \odot \chi) \leftrightarrow \varphi \odot (\psi \ominus \chi)$,
(ŁΠ2) $\Delta(\varphi \to \psi) \to (\varphi \to_\Pi \psi)$,
(ŁΠ3) $\Delta(\varphi \to_\Pi \psi) \to (\varphi \to \psi)$,
(ŁΠ4) $\overline{\frac{1}{2}} \leftrightarrow \neg\overline{\frac{1}{2}}$.

The deduction rules are modus ponens for & and $\to$, and the necessitation rule for Δ, i.e.: from φ derive $\Delta\varphi$. ŁΠ$\frac{1}{2}$ is complete with respect to deductions from finite theories for the given semantics [5].

While Ł is the logic of McNaughton functions, ŁΠ$\frac{1}{2}$ is the logic of piecewise rational functions with integer coefficients over $[0,1]^n$ (see [15]). In fact, the function defined by each ŁΠ$\frac{1}{2}$-formula corresponds to a supremum of rational fractions

$$\frac{P(x_1, \ldots, x_n)}{Q(x_1, \ldots, x_n)}$$

over $[0,1]^n$, where $P(x_1, \ldots, x_n), Q(x_1, \ldots, x_n)$ are polynomials with rational coefficients. Conversely, every piecewise rational function with rational coefficients over the unit cube $[0,1]^n$ can be defined by an ŁΠ$\frac{1}{2}$-formula.

3 Logics for Representing Expectation

In this section we first introduce the concept of a Łukasiewicz game on Ł$_k^c$ along with the notion of a class $\mathfrak{G}$ of Łukasiewicz games. Then, we define the logic $\mathsf{E}(\mathfrak{G})$ to represent expected payoffs for games in $\mathfrak{G}$, and provide completeness and complexity results.

3.1 Łukasiewicz Games

A Łukasiewicz game $\mathcal{G}$ on $Ł_k^c$ [14] is a tuple

$$\mathcal{G} = \langle \mathsf{P}, \mathsf{V}, \{\mathsf{V}_i\}, \{\mathsf{S}_i\}, \{\varphi_i\}\rangle$$

where:

1. $\mathsf{P} = \{P_1, \ldots, P_n\}$ is a set of *players*;
2. $\mathsf{V} = \{p_1, \ldots, p_m\}$ is a finite set of propositional variables;
3. For each $i \in \{1, \ldots, n\}$, $\mathsf{V}_i \subseteq \mathsf{V}$ is the set of propositional variables under control of player P_i, so that the sets V_i form a partition of V, $|\mathsf{V}_i| = m_i$, and $\sum_{i=1}^{n} m_i = m$.
4. For each $i \in \{1, \ldots, n\}$, S_i is the strategy set for player P_i that includes all valuations $s : \mathsf{V}_i \to L_k$ of the propositional variables in V_i, i.e.

$$\mathsf{S}_i = \{s \mid s : \mathsf{V}_i \to L_k\}.$$

5. For each $i \in \{1, \ldots, n\}$, $\varphi_i(p_1, \ldots, p_t)$ is an $Ł_k^c$-formula, built from variables in V, whose associated function

$$f_{\varphi_i} : (L_k)^t \to L_k$$

corresponds to the *payoff function* of P_i, and whose value is determined by the valuations in $\{\mathsf{S}_1, \ldots, \mathsf{S}_n\}$.

We denote by $\mathsf{S} = \mathsf{S}_1 \times \cdots \times \mathsf{S}_n$ the product of the strategy spaces. A tuple $\boldsymbol{s} = (s_1, \ldots, s_n) \in \mathsf{S}$ of strategies is called a *strategy combination*. s_{-i} denotes the tuple of strategies $(s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ not including s_i. Aside from $\boldsymbol{s}$, we use the form (s_i, s_{-i}) do denote a strategy combination. With an abuse of notation, we denote by $f_{\varphi_i}(s_i, s_{-i})$ (or, equivalently, $f_{\varphi_i}(\boldsymbol{s})$) the value of the payoff function f_{φ_i} under the valuation corresponding to the strategy combination (s_i, s_{-i}) [or, equivalently, $\boldsymbol{s}$].[2]

Given a game $\mathcal{G}$, let

$$\delta : \mathsf{P} \to \{1, \ldots, m\}$$

be a function assigning to each player P_i an integer from $\{1, \ldots, m\}$ that corresponds to the number of variables in V_i: i.e.:

$$\delta(P_i) = m_i.$$

δ is called a *variable distribution function*. Given a game $\mathcal{G}$, the *type* of $\mathcal{G}$ is the triple $\langle n, m, \delta \rangle$, where n is the number of players, m is the number of variables in V, and δ is the variable distribution function for $\mathcal{G}$.

---

[2] $f_{\varphi_i}(s_i, s_{-i})$ can be also seen as the value of $e(\varphi_i)$, when e coincides with the valuation $\boldsymbol{s}$.

Definition 1 (Class). *Let $\mathcal{G}$ and $\mathcal{G}'$ be two Łukasiewicz games $\mathcal{G}$ and $\mathcal{G}'$ on L_k^c of type $\langle n, m, \delta \rangle$ and $\langle n, m, \delta' \rangle$, respectively. We say that $\mathcal{G}$ and $\mathcal{G}'$ belong to the same class $\mathfrak{G}$ if there exists a permutation $\jmath$ of the indices $\{1, \ldots, n\}$ such that, for all P_i,*

$$\delta(P_{\jmath(i)}) = \delta'(P_i).$$

Notice that what matters in the definition of a type is not which players are assigned certain variables, but rather their distribution. For instance, take two games $\mathcal{G}$ and $\mathcal{G}'$ all having three players P_1, P_2, P_3 and the same variables $p_1, \ldots, p_6$. Suppose that, in $\mathcal{G}$, P_1 controls p_1, P_2 controls p_2, p_3, and P_3 controls p_4, p_5, p_6, while in $\mathcal{G}'$, P_2 controls p_3, P_3 controls p_4, p_5, and P_1 controls p_1, p_2, p_6. $\mathcal{G}$ and $\mathcal{G}'$ have the same type, since they have the same number of players, the same number of variables, and the permutation $\jmath$, where

$$1 \overset{\jmath}{\longmapsto} 2 \qquad 2 \overset{\jmath}{\longmapsto} 3 \qquad 3 \overset{\jmath}{\longmapsto} 1,$$

is such that $\delta\left(P_{\jmath(i)}\right) = \delta'\left(P_i\right)$ for all P_i.

Given the above definitions, we can adapt some well-known game-theoretic concepts to this settings. We introduce the notion of mixed strategy in order to define the concept of expected payoff along with the notions of best response and equilibrium.

Let $\mathcal{G} = \langle \mathsf{P}, \mathsf{V}, \{\mathsf{V}_i\}, \{\mathsf{S}_i\}, \{\varphi_i\} \rangle$ be a Łukasiewicz game on L_k^c. A *mixed strategy* π_i for player P_i is a probability distribution on the strategy space S_i. By π_{-i}, we denote the tuple of mixed strategies $(\pi_1, \ldots, \pi_{i-1}, \pi_{i+1}, \ldots, \pi_n)$.

Given a tuple $(\pi_1, \ldots, \pi_n)$ of mixed strategies for $P_1, \ldots, P_n$, respectively, the *expected payoff* for P_i of playing π_i, when P_{-i} play π_{-i}, is given by

$$exp_{\varphi_i}(\pi_i, \pi_{-i}) = \sum_{\boldsymbol{s} = (s_1, \ldots, s_n) \in \mathsf{S}} \left(\left(\prod_{j=1}^{n} \pi_j(s_j) \right) \cdot f_{\varphi_i}\left(\boldsymbol{s}\right) \right)$$

Let Σ_i denote the set of mixed strategies of P_i. P_i's *best response* to P_{-i}'s mixed strategy combination π_{-i} is a mixed strategy π_i such that, for all strategies $\pi_i' \in \Sigma_i$:

$$exp_{\varphi_i}(\pi_i, \pi_{-i}) \geq exp_{\varphi_i}(\pi_i', \pi_{-i}).$$

Definition 2 (Nash Equilibrium). *Let $\mathcal{G}$ be a Łukasiewicz game on L_k^c. We call a tuple of mixed strategies $(\pi_1^*, \ldots, \pi_n^*)$ a* Nash Equilibrium *for $\mathcal{G}$ if each player's mixed strategy π_i^* is a best response to the other players' mixed strategy combination π_{-i}^*.*

Example: Matching Pennies. The following game is a generalisation of Matching Pennies[3], a classic example of a zero-sum game without a pure strategy equilibrium. In the original game, two players P_1 and P_2 both have a penny and must secretly choose to turn it to head or tails and reveal their choice at the same

---

[3] The idea behind this generalisation comes from [13].

time. If their choices agree, P_1 takes both pennies, but if they do not match, P_2 is the one winning both.

Now, imagine that both players have a dice with $n + 1$ faces, they both choose one and reveal it at the same time. P_1's overall strategy is to be as close as possible to P_2's choice, who, instead, wants to keep the greater possible distance between the outcomes. Clearly, we can represent each player's strategy space with the set L_k. P_1's payoff function is given by the formula $\neg d(p_1, p_2)$, whose associated function is $1 - |e(p_1) - e(p_2)|$, while P_2's payoff is defined by the formula $d(p_1, p_2)$, which corresponds to the function $|e(p_1) - e(p_2)|$. The game is formally defined on L_k^c as follows:

$$\mathcal{G} = \Big\langle \{P_1, P_2\}, \{p_1, p_2\}, \{p_1\}_1, \{p_2\}_2, \{s_1 : \{p_1\} \to L_k\}, \{s_2 : \{p_2\} \to L_k\}, \{\neg d(p_1, p_2), d(p_1, p_2)\} \Big\rangle.$$

Table 1 shows the payoff matrix for this generalised version of Matching Pennies with $k = 5$ (the values in each cell correspond to the first and second player payoffs under the valuation to the variables p_1, p_2).

By Nash's Theorem [17], we know that an equilibrium with mixed strategies always exists for the generalised version of Matching Pennies.

				P_2			
		0	1/5	2/5	3/5	4/5	1
	0	1, 0	4/5, 1/5	3/5, 2/5	2/5, 3/5	1/5, 4/5	0, 1
	1/5	4/5, 1/5	1, 0	4/5, 1/5	3/5, 2/5	2/5, 3/5	1/5, 4/5
P_1	2/5	3/5, 2/5	4/5, 1/5	1, 0	4/5, 1/5	3/5, 2/5	2/5, 3/5
	3/5	2/5, 3/5	3/5, 2/5	4/5, 1/5	1, 0	4/5, 1/5	3/5, 2/5
	4/5	1/5, 4/5	2/5, 3/5	3/5, 2/5	4/5, 1/5	1, 0	4/5, 1/5
	1	0, 1	1/5, 4/5	2/5, 3/5	3/5, 2/5	4/5, 1/5	1, 0

Table 1. Generalised Matching Pennies.

3.2 The Logic $\mathsf{E}(\mathfrak{G})$

The aim of this section is to introduce the logic $\mathsf{E}(\mathfrak{G})$ for reasoning about expected utility in games with Łukasiewicz strategies. Notice that, while sometimes we refer to $\mathsf{E}(\mathfrak{G})$ as "a logic", we are actually defining a whole family of logics, one for each class $\mathfrak{G}$.

Syntax The construction of $\mathsf{E}(\mathfrak{G})$ mimics the one provided for logics for reasoning about uncertainty in [7]. The syntax of $\mathsf{E}(\mathfrak{G})$ is built by taking the m-variable fragment $m\mathrm{L}_k^{c\,4}$ of L_k^c as inner logic and $\mathrm{L}\Pi\frac{1}{2}$ as outer logic. Its language is defined as follows:

1. The set NModF of non-modal formulas corresponds to the set of $m\mathrm{L}_k^c$-formulas built from the propositional variables $p_1, \ldots, p_m$.

2. The set ModF of modal formulas is built from the atomic modal formulas $\mathsf{E}\varphi$, with $\varphi \in \mathrm{NModF}$, using the $\mathrm{L}\Pi\frac{1}{2}$ connectives. $\mathsf{E}\varphi$ is meant to encode a player's expected payoff from playing a mixed strategy, given the payoff function associated to φ. Nested modalities are not allowed.

Semantics Given a class of games $\mathfrak{G}$ on $m\mathrm{L}_k^c$, a model $\mathbf{M}$ for $\mathsf{E}(\mathfrak{G})$ is a tuple $\langle \mathsf{S}, e, \{\pi_i\} \rangle$, such that:

1. S is the set of all strategy combinations $\{ \boldsymbol{s} = (s_1, \ldots, s_n) \mid (s_1, \ldots, s_n) \in \mathsf{S}_1 \times \cdots \times \mathsf{S}_n \}$.

2. $e : (\mathrm{NModF} \times \mathsf{S}) \to L_k$ is a valuation of non-modal formulas, such that, for each $\varphi \in \mathrm{NModF}$

$$e(\varphi, \boldsymbol{s}) = f_\varphi(\boldsymbol{s}),$$

   where f_φ is the function associated to φ and $\boldsymbol{s} = (s_1, \ldots, s_n)$.

3. $\pi_i : \mathsf{S}_i \to [0,1]$ is a probability distribution, for each P_i.

Given a formula Φ, the truth value of Φ in $\mathbf{M}$ at the combination $\boldsymbol{s}$, denoted $\|\Phi\|_{\mathbf{M},\boldsymbol{s}}$, is inductively defined as follows:

1. If Φ is a non-modal formula $\varphi \in \mathrm{NModF}$, then

$$\|\varphi\|_{\mathbf{M},\boldsymbol{s}} = e(\varphi, \boldsymbol{s}),$$

2. If Φ is an atomic modal formula $\mathsf{E}\varphi$, then

$$\|\mathsf{E}\varphi\|_{\mathbf{M},\boldsymbol{s}} = exp_\varphi(\pi_1, \ldots, \pi_n) = \sum_{\boldsymbol{s}=(s_1,\ldots,s_n)\in\mathsf{S}} \left(\left(\prod_{j=1}^{n} \pi_j(s_j) \right) \cdot e(\varphi, \boldsymbol{s}) \right). ^5$$

3. If Φ is a non-atomic modal formula, its truth value is computed by evaluating its atomic modal subformulas and then by using the truth functions associated to the $\mathrm{L}\Pi\frac{1}{2}$-connectives occurring in Φ.

---

[4] Notice that the functions definable in $m\mathrm{L}_k^c$ are exactly all the functions $f : (L_k)^t \to L_k$, where $t \leq m$.

[5] Notice that the notation is slightly different from the one used above for the definition of expected payoff but the meaning is the same.

Since the valuation of a modal formula Φ does not depend on a specific strategy combination but only on the model $\mathbf{M}$, we will often simply write $\|\Phi\|_{\mathbf{M}}$ to denote the valuation of Φ in $\mathbf{M}$. As usual, we say that a formula Φ is satisfiable if there exists a model $\mathbf{M}$ such that $\|\Phi\|_{\mathbf{M}} = 1$. Similarly, a modal theory Γ, i.e. a set of modal formulas, is satisfiable if there exists a model that satisfies each and every one of the formulas contained in Γ. Validity clearly means satisfiability in all models.

Axiomatisation The axioms of $\mathsf{E}(\mathfrak{G})$ are the following:

1. All the L_k^c-tautologies for $m\mathrm{L}_k^c$, i.e.: all the L_k^c-tautologies in the variables $p_1, \ldots, p_m$, for non-modal formulas.

2. All the $\mathrm{L}\Pi\frac{1}{2}$-axioms and rules for modal formulas.

3. Probabilistic axioms for E, with $\varphi, \psi, \bar{r} \in \mathrm{NModF}$:

   (a) $\mathsf{E}(\neg\varphi) \leftrightarrow \neg\mathsf{E}\varphi$
   (b) $\mathsf{E}(\varphi \oplus \psi) \leftrightarrow [(\mathsf{E}\varphi \to \mathsf{E}(\varphi \& \psi)) \to \mathsf{E}\psi]$
   (c) $\mathsf{E}\bar{r} \leftrightarrow \bar{r}$

4. Independence axioms for E, where $p_{1_i}, \ldots, p_{m_i}$ is the tuple of variables assigned to P_i, for all tuples $r_{1_1}, \ldots, r_{m_1}, \ldots, r_{1_n}, \ldots, r_{m_n} \in (L_k)^m$:

   (a) $\mathsf{E}\left(\bigwedge_{i=1}^n \left(\bigwedge_{j_i=1_i}^{m_i} (\Delta\,(p_{j_i} \leftrightarrow \bar{r}_{j_i}))\right)\right) \leftrightarrow \bigodot_{i=1}^n \left(\mathsf{E}\left(\bigwedge_{j_i=1_i}^{m_i} \Delta\,(p_{j_i} \leftrightarrow \bar{r}_{j_i})\right)\right)$

5. The following inference rules for E, with $\varphi, \psi \in \mathrm{NModF}$:

   (a) Necessitation: from φ derive $\mathsf{E}\varphi$
   (b) Monotonicity: from $\varphi \to \psi$ derive $\mathsf{E}\varphi \to \mathsf{E}\psi$

   Notice that each

   $$(\dagger) \qquad \bigwedge_{j_i=1_i}^{m_i} (\Delta\,(p_{j_i} \leftrightarrow \bar{r}_{j_i}))$$

encodes a particular strategy for player P_i. In fact, each $\Delta\,(p_{j_i} \leftrightarrow \bar{r}_{j_i})$ is satisfiable if and only if $e(p_{j_i}) = r_{j_i}$. So, given a tuple $r_{1_i}, \ldots, r_{m_i}$, $(\dagger)$ is satisfiable by a strategy s_i if and only if $s_i(p_{j_i}) \longmapsto r_{j_i}$, for all p_{j_i}.

The notion of *proof* in $\mathsf{E}(\mathfrak{G})$ is defined as usual. For any modal theory Γ and formula φ, we write

$$\Gamma \vdash_{\mathsf{E}(\mathfrak{G})} \varphi$$

to denote that φ is a consequence of Γ in $\mathsf{E}(\mathfrak{G})$.

Completeness Before we prove completeness for $\mathsf{E}(\mathfrak{G})$ we provide an axiomatic characterisation for the expectation of mixed strategies over $m\mathrm{L}_k^c$-formulas.

Theorem 1. *Let $\mathfrak{G}$ be a class of Łukasiewicz games on L_k^c and let $m\mathrm{L}_k^c$ be the m-variable fragment of L_k^c. The following statements are equivalent:*

1. *There exists a state*

$$\sigma : m\mathrm{L}_k^c \to [0,1].$$

such that for all tuples $r_{1_1}, \ldots, r_{m_1}, \ldots, r_{1_n}, \ldots, r_{m_n} \in (L_k)^m$

$$\sigma\left(\bigwedge_{i=1}^{n}\left(\bigwedge_{j_i=1_i}^{m_i}\left(\Delta\left(p_{j_i} \leftrightarrow \overline{r}_{j_i}\right)\right)\right)\right) = \prod_{i=1}^{n}\left(\sigma\left(\bigwedge_{j_i=1_i}^{m_i}\Delta\left(p_{j_i} \leftrightarrow \overline{r}_{j_i}\right)\right)\right),$$

where $p_{1_i}, \ldots, p_{m_i}$ *is the tuple of variables assigned to* P_i.

2. *There exists a probability distribution*

$$\pi_i : \mathsf{S}_i \to [0,1]$$

for each P_i, *such that, for all* $\varphi \in m\mathrm{L}_k^c$, $\sigma(\varphi) = exp_\varphi(\pi_1, \ldots, \pi_n)$, *i.e.:*

$$\sigma(\varphi) = \sum_{\boldsymbol{s}=(s_1,\ldots,s_n)\in\mathsf{S}}\left(\left(\prod_{j=1}^{n}\pi_j(s_j)\right)\cdot f_\varphi(\boldsymbol{s})\right),$$

where f_φ *is the function associated to* φ.

Proof. It is easy to check that (2) implies (1).

To prove the converse, suppose that there exists a state

$$\sigma : m\mathrm{L}_k^c \to [0,1].$$

As shown by Paris in [18, Appendix 2] there exists a probability distribution $\pi : \mathsf{S}_1 \times \cdots \times \mathsf{S}_n \to [0,1]$ such that

$$\sigma(\varphi) = \sum_{\boldsymbol{s}\in\mathsf{S}}\pi(\boldsymbol{s})\cdot f_\varphi(\boldsymbol{s}),$$

for all $\varphi \in m\mathrm{L}_k^c$.

Now, let $m_i\mathrm{L}_k^c[\boldsymbol{p}_i]$ be the m_i-variable fragment of L_k^c in the variables $\boldsymbol{p}_i = \{p_{1_i}, \ldots, p_{m_i}\}$, i.e. the variables assigned to P_i. Let $\sigma_{\downarrow i}$ be the restrictions of σ to $m_i\mathrm{L}_k^c[\boldsymbol{p}_i]$. It is clear that each $\sigma_{\downarrow i}$ is still a probability measure. It follows, again, from [18, Appendix 2] that, for each i, there exists a probability distribution

$$\pi_i : \mathsf{S}_i \to [0,1]$$

such that, for all $\psi \in m_i\mathrm{L}_k^c[\boldsymbol{p}_i]$

$$\sigma(\psi) = \sigma_{\downarrow 1}(\psi) = \sum_{s_i\in\mathsf{S}_i}\pi_i(s_i)\cdot f_\psi(s_i).$$

By assumption, σ satisfies

$$\sigma\left(\bigwedge_{i=1}^{n}\left(\bigwedge_{j_i=1_i}^{m_i}\left(\Delta\left(p_{j_i} \leftrightarrow \overline{r}_{j_i}\right)\right)\right)\right) = \prod_{i=1}^{n}\left(\sigma\left(\bigwedge_{j_i=1_i}^{m_i}\Delta\left(p_{j_i} \leftrightarrow \overline{r}_{j_i}\right)\right)\right),$$

for all tuples $r_{1_1}, \ldots, r_{m_1}, \ldots, r_{1_n}, \ldots, r_{m_n} \in (L_k)^m$.

It is possible to check that the above condition guarantees the fact that all probability distributions π_i are independent, and so for all $\boldsymbol{s} \in \mathsf{S}$

$$\pi(\boldsymbol{s}) = \prod_{i=1}^{n} \pi_i(s_i).$$

Therefore

$$\sigma(\varphi) = \sum_{\boldsymbol{s} \in \mathsf{S}} \pi(\boldsymbol{s}) \cdot f_\varphi(\boldsymbol{s}) = \sum_{\boldsymbol{s}=(s_1,\ldots,s_n) \in \mathsf{S}} \left(\left(\prod_{j=1}^{n} \pi_j(s_j) \right) \cdot f_\varphi(\boldsymbol{s}) \right).$$

We can now proceed to proving the Completeness Theorem.

Theorem 2 (Completeness). *Let Γ and Φ be a finite modal theory and a modal formula in $\mathsf{E}(\mathfrak{G})$. Then, $\Gamma \vdash_{\mathsf{E}(\mathfrak{G})} \Phi$ if and only if for every model $\mathbf{M}$ such that, for each $\Psi \in \Gamma$, $\|\Psi\|_{\mathbf{M}} = 1$, also $\|\Phi\|_{\mathbf{M}} = 1$.*

Proof. The proof follows from an adaptation of the strategy laid out in [10] and generalised in [7].

We now study the computational complexity of certain kinds of satisfiability problems for $\mathsf{E}(\mathfrak{G})$. Let $r \in \mathbb{Q} \cap [0,1]$ and let $\flat \in \{<,>,\leq,\geq,=\}$. We call an $\mathsf{E}(\mathfrak{G})$-modal formula Φ $\flat r$-*satisfiable* if there is a model $\mathbf{M}$ such that

$$\|\Phi\|_{\mathbf{M}} \flat r.$$

Following a strategy similar to the one laid out by Hájek in [11] for probability logics, satisfiability of an $\mathsf{E}(\mathfrak{G})$-formula Φ can be translated into an existential formula in the theory of the reals whose size is exponential in the number of non-modal propositional variables in Φ. Decidability for the existential theory of the reals $\mathsf{Th}_{\exists}(\mathbb{R})$ was shown by Canny to be in PSPACE [2]. Therefore, we obtain the following result:

Theorem 3. *Checking $\flat r$-satisfiability for $\mathsf{E}(\mathfrak{G})$ is in EXPSPACE.*

4 Games with Expectations

In this section we introduce a class of games with polynomial constraints over expectations. These games expand Lukasiewicz games by assigning to each player a formula of $\mathsf{E}(\mathfrak{G})$, which is a piecewise rational polynomial function whose variables correspond to expected values. The idea is that in a situation of strategic interaction players might be interested not only in maximizing their own expectation, but also in influencing others'. The modal formula assigned to each player is then meant to represent a player's goal concerning the relation between her and other players' expectations.

A game with expectations $\mathcal{E}_{\mathcal{G}}$ on $\mathsf{E}(\mathfrak{G})$ is a tuple

$$\mathcal{E}_{\mathcal{G}} = \langle \mathsf{P}, \mathsf{V}, \{\mathsf{V}_i\}, \{\mathsf{S}_i\}, \{\varphi_i\}, \{\mathsf{M}_i\}, \{\Phi_i\}\rangle,$$

where:

1. $\mathcal{G} = \langle \mathsf{P}, \mathsf{V}, \{\mathsf{V}_i\}, \{\mathsf{S}_i\}, \{\varphi_i\}\rangle$ is a Łukasiewicz game on L_k^c, with $\mathcal{G} \in \mathfrak{G}$,
2. for each $i \in \{1, \ldots, n\}$, M_i is the set of all mixed strategies on S_i of player P_i,
3. for each $i \in \{1, \ldots, n\}$, Φ_i is an $\mathsf{E}(\mathfrak{G})$-formula such that every atomic modal formula occurring in Φ_i has the form $\mathsf{E}\psi$, with $\psi \in \{\varphi_1, \ldots, \varphi_n\}$.

Let $\mathcal{E}_{\mathcal{G}}$ be a game with expectations on $\mathsf{E}(\mathfrak{G})$. A model $\mathbf{M} = \langle \mathsf{S}, e, \{\pi_i\}\rangle$ for $\mathsf{E}(\mathfrak{G})$ is called a *best response model* for a player P_i whenever, for all models $\mathbf{M}' = \langle \mathsf{S}, e, \{\pi_i'\}\rangle$ with $\pi_{-i}' = \pi_{-i}$,

$$\|\Phi_i\|_{\mathbf{M}'} \leq \|\Phi_i\|_{\mathbf{M}}.$$

Definition 3 (Equilibrium). *A game with expectations $\mathcal{E}_{\mathcal{G}}$ on $\mathsf{E}(\mathfrak{G})$ is said to have a* Nash Equilibrium, *whenever there exists a model $\mathbf{M}^*$ that is a best response model for each player P_i.*

Example 1. Let $\mathcal{E}_{\mathcal{G}}$ be any game with expectations where each P_i is simply assigned the formula $\Phi_i := \mathsf{E}\varphi_i$. This game corresponds to the the situation where each player cares only about her own expectation and whose goal is its maximisation. Clearly, by Nash's Theorem [17], every $\mathcal{E}_{\mathcal{G}}$ of this form admits an Equilibrium, since it offers a formalisation of the classical case where equilibria are given by tuples of mixed strategies over valuations in a Łukasiewicz game.

Example 2. Not every game with expectations has an equilibrium. In fact, consider the following game

$$\mathcal{E}_{\mathcal{G}} = \langle \mathsf{P}, \mathsf{V}, \{\mathsf{V}_i\}, \{\mathsf{S}_i\}, \{\varphi_i\}, \{\mathsf{M}_i\}, \{\Phi_i\}\rangle,$$

with $i \in \{1, 2\}$, where:

1. $\varphi_1 := p_1$ and $\varphi_2 := p_2$, and
2. $\Phi_1 := \neg d(\mathsf{E}(p_1), \mathsf{E}(p_2))$ and $\Phi_2 := d(\mathsf{E}(p_1), \mathsf{E}(p_2))$.

The above game can be regarded as a particular version of Matching Pennies with expectations. In fact, while P_1 aims at matching P_2's expectation, P_2's goal is quite the opposite, since she wants their expectations to be as far as possible.

It is easy to see that there is no model $\mathbf{M}$ that gives an equilibrium for $\mathcal{E}_{\mathcal{G}}$.

Proposition 1. *There exist games with expectations on $\mathsf{E}(\mathfrak{G})$ that do not admit a Nash Equilibrum.*

As mentioned above, the satisfiability of every $\mathsf{E}(\mathfrak{G})$-formula can be translated into the validity of an existential formula of the theory $\mathsf{Th}(\mathbb{R})$ of real closed fields. In a similar way, we can express the existence of an equilibrium in a game

with expectations $\mathcal{E}_\mathcal{G}$ through a first-order sentence ξ of $\mathsf{Th}(\mathbb{R})$ having exponentially many variables and a fixed alternation of quantifiers. By using the fact that the general decidability problem for $\mathsf{Th}(\mathbb{R})$ is singly exponential in the number of variables when the alternation of quantifiers is fixed [8], we obtain the following result.

Theorem 4 (Complexity). *Checking the existence of a Nash Equilibrium in a game with expectations $\mathcal{E}_\mathcal{G}$ on $\mathsf{E}(\mathfrak{G})$ is in 2-EXPTIME.*

By exploiting the connection between $\mathrm{L\Pi}\frac{1}{2}$ and real closed fields it is also possible to express the existence of an equilibrium through an $\mathrm{L\Pi}\frac{1}{2}$-formula (see [4]). Therefore, we obtain the following logical characterisation.

Theorem 5. *For every game with expectations $\mathcal{E}_\mathcal{G}$ on $\mathsf{E}(\mathfrak{G})$ there exists an $\mathrm{L\Pi}\frac{1}{2}$-formula ϕ such that $\mathcal{E}_\mathcal{G}$ admits a Nash Equilibrium if and only if ϕ is satisfiable.*

5   Final Remarks

In this work, we presented a logic $\mathsf{E}(\mathfrak{G})$ for reasoning about expectations in a class $\mathfrak{G}$ of Łukasiewicz games. We have also introduced a new class of games based on $\mathsf{E}(\mathfrak{G})$ that expand Łukasiewicz games. These games capture strategic interactions in which players randomise their choices in order to influence the expectations of the other players as well as their own.

While our approach to representing expectation in games in a logical framework is certainly novel, other works have dealt with similar topics. In [9], Halpern and Pucella introduced a propositional logic for reasoning about probabilistic expectation. Their work, though, is mainly concerned with modelling expectation in general and not in a game-theoretic setting. Certainly closer to our paper is the work by Sack and van der Hoek [19], where the authors study a modal logic to reason about mixed strategies in games. Although (each instance of) their logic is based on a fixed game and a fixed set of mixed strategies, it makes it possible to represent the concept of a Nash Equilibrium through logical formulas: a feature that is not possible in our logic.

In our future work, we plan to extend the logical study of expectation for Łukasiewicz games to those situations where payoff formulas are functions defined over the whole unit cube $[0, 1]^n$, and, consequently, players have an infinite strategy space.

Acknowledgments

References

1. E. Bonzon, M.C. Lagasquie-Schiex, J. Lang, B. Zanuttini. Boolean games revisited. In *Proceedings of the Seventeenth European Conference on Artificial Intelligence* (ECAI-2006), Riva del Garda, Italy, 265–269, 2006.
2. J.F. Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of the 20^{th} ACM Symposium on Theory of Computing*, 460–467, 1988.
3. R. Cignoli, I. M. L. D'Ottaviano, D. Mundici. *Algebraic Foundations of Many-alued Reasoning*, Volume 7 of *Trends in Logic*, Kluwer Academic Publishers, Dordrecht, 2000.
4. F. Esteva, L. Godo, E. Marchioni. Fuzzy logics with enriched language. In *Handbook of Mathematical Fuzzy Logic*, Cintula P., Hájek P., and Noguera C. (Editors), College Publications, 627–712, 2011.
5. F. Esteva, L. Godo, F. Montagna. The ŁΠ and ŁΠ$\frac{1}{2}$ logics: two complete fuzzy logics joining Łukasiewicz and product logic. *Archive for Mathematical Logic*, 40: 39–67, 2001.
6. T. Flaminio, L. Godo. A logic for reasoning about the probability of fuzzy events. Fuzzy Sets and Systems, 158: 625–638, 2007.
7. T. Flaminio, L. Godo, E. Marchioni. Reasoning about uncertainty of fuzzy events: an overview. In *Reasoning under Vagueness - Logical, Philosophical, and Linguistic Perspectives*, Cintula P., Fermüller C., Godo L., Hájek P. (Editors), Studies in Logic, College Publications, 367–400, 2011.
8. D. Y. Grigor'ev. Complexity of deciding Tarski algebra. *Journal of Symbolic Computation*, 5: 65–108, 1988.
9. J.Y. Halpern, R. Pucella. Characterizing and reasoning about probabilistic and non-probabilistic expectation. *Journal of the ACM*, Article No. 15, 54(3), 2007.
10. P. Hájek *Metamathematics of Fuzzy Logic*. Volume 4 of *Trends in Logic*, Kluwer Academic Publishers, Dordrecht, 1998.
11. P. Hájek. Complexity of fuzzy probability logics II. *Fuzzy Sets and Systems*, 158(23): 2605–2611, 2007.
12. P. Harrenstein, W. van der Hoek, J.J.Ch. Meyer, C. Witteveen. Boolean games. In *Proceeding of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge* (TARK VIII), J. van Benthem (Ed.), Siena, Italy, 287–298, 2001.
13. T. Kroupa, O. Majer. Nash Equilibria in a class of constant-sum games represented by McNaughton functions. In *Logic, Algebra and Truth Degrees 2012, Book of Abstracts*, K. Terui and N. Preining (Eds.), 2012.
14. E. Marchioni, M. Wooldridge. Łukasiewicz Games. Submitted.
15. F. Montagna, G. Panti. Adding structures to MV-algebras. *Journal of Pure and Applied Algebra*, 164: 365–387, 2001.
16. D. Mundici, Averaging the truth-value in Łukasiewicz logic. *Studia Logica* 55(1): 113–127, 1995.
17. J. Nash. Non-cooperative games. *The Annals of Mathematics*, Second Series, 54(2): 286–295, 1951.
18. J. B. Paris. A note on the Dutch Book method (revised version). In *Second International Symposium on Imprecise Probability and their Application*, Cornell University, Ithaca, NY (USA), 2001.
19. J. Sack, W. van der Hoek. A Modal Logic for Mixed Strategies. *Studia Logica*, accepted.