


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.


In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/authorsrights>



Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar


Optimal strategic reasoning with McNaughton functions

Tomáš Kroupa^a, Ondrej Majer^b^a Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Pod Vodárenskou věží 4, 182 08 Prague, Czech Republic^b Institute of Computer Science, Academy of Sciences of the Czech Republic, Pod Vodárenskou věží 2, 182 07 Prague, Czech Republic

ARTICLE INFO

Article history:

Received 31 December 2013

Received in revised form 7 April 2014

Accepted 7 April 2014

Available online 24 April 2014

Keywords:

Lukasiewicz logic

McNaughton function

Infinite game

Nash equilibria

ABSTRACT

The aim of the paper is to explore strategic reasoning in strategic games of two players with an uncountably infinite space of strategies the payoff of which is given by McNaughton functions—functions on the unit interval which are piecewise linear with integer coefficients. McNaughton functions are of a special interest for approximate reasoning as they correspond to formulas of infinitely valued Lukasiewicz logic. The paper is focused on existence and structure of Nash equilibria and algorithms for their computation. Although the existence of mixed strategy equilibria follows from a general theorem (Glicksberg, 1952) [5], nothing is known about their structure neither the theorem provides any method for computing them. The central problem of the article is to characterize the class of strategic games with McNaughton payoffs which have a finitely supported Nash equilibrium. We give a sufficient condition for finite equilibria and we propose an algorithm for recovering the corresponding equilibrium strategies. Our result easily generalizes to n -player strategic games which don't need to be strictly competitive with a payoff functions represented by piecewise linear functions with real coefficients. Our conjecture is that every game with McNaughton payoff allows for finitely supported equilibrium strategies, however we leave proving/disproving of this conjecture for future investigations.

© 2014 Elsevier Inc. All rights reserved.

0. Prologue

Strategic considerations of players in the theory of noncooperative games [20] represent one of the classical examples of reasoning under uncertainty. In this paper we focus on a particular class of games, where the goals of players can be expressed as a formula of a formal language. Typically we assume that the formula represents a claim and the goal of one player is to show that the claim is true, while the goal of the other one is to disprove it. Truth or falsity of the claim depend on parameters which players can control, such as the values of atomic formulas in the case the language in question is that of propositional logic. Each of the players can influence values of some parameters, while the remaining ones are under the control of the other player. Informally, we can imagine that a player makes a statement “there is a situation, which I can demonstrate by choosing appropriate values for the parameters under my control, in which the claim is true/false (independently of the remaining parameters)”.

This kind of game is well-known in literature. In [8] the assertion in question is expressed in classical propositional logic and parameters are Boolean variables divided into disjoint sets, each of which is under the control of one of the players.

From the logical point of view, we shall deal with a more general set-up in which the goals of players might generally be satisfied only up to some degree. Reasoning with claims evaluated on a non-dichotomic scale is the natural domain of many-valued logics [7], which offer versatile tools to model the statements which are not just true or false,

<http://dx.doi.org/10.1016/j.ijar.2014.04.009>

0888-613X/© 2014 Elsevier Inc. All rights reserved.

but which can take an intermediate truth value. Although we are considering games of two players with each of them controlling only a single propositional variable, the complexity of strategic considerations increases radically as the space of strategies is a continuum. This is in sharp contrast with the finite space considered in Boolean games. Marchioni and Wooldridge [11] present a generalization of Boolean games towards many-valued logics similar to our approach—the goals of players are represented by a formula in Łukasiewicz logic. In comparison to our approach they admit only finitely many truth values, but allow any finite number of players. Moreover, the game they introduce need not be strictly competitive.

The extension to uncountably-many truth-values has two important features from the viewpoint of game modeling and representation. First, the continuum of truth values enables us to capture a much larger space of payoff functions than with the Boolean games. Second, it also introduces a new kind of uncertainty into game-theoretical reasoning. Probability is standard equipment used in game theory—it appears naturally when the notion of a mixed strategy is introduced as randomizing over deterministic choices of (pure) strategies. Introducing many-valued framework, on the other hand, provides a broader scale for approximate reasoning about strategies: the players have to consider combining probability with graded truth values.

1. Introduction

Several interesting game-theoretic models arose from the investigations of reasoning under more than two degrees of truth, such as the bookmaking model over Łukasiewicz formulas [13] and the logical characterization of coherence [3]. In this paper we are going to develop a two-player strategic game in which the payoff functions are directly associated with formulas in infinite-valued Łukasiewicz logic. The free n -generated MV-algebra of Łukasiewicz logic [2,7] is the MV-algebra of McNaughton functions, which are many-valued analogues of Boolean functions. The expressive power of McNaughton functions is greater than that of Boolean functions so that many interesting strategic games can be expressed by formulas in Łukasiewicz logic. Since the associated utility functions are continuous and piecewise linear with integer coefficients, a Nash equilibrium exists by Glicksberg's Theorem [5], which is an infinite-dimensional generalization of the Nash Theorem [14]. However, the whole class of piecewise linear functions is not amenable to the known computational techniques for strategic games [17,18]. The main open questions are (i) whether finitely supported equilibria exist for each McNaughton function, (ii) provided that the answer to (i) is affirmative, are there algorithms for computing such equilibria? It is worth mentioning that “To find a computational technique of general applicability for zero-sum two-person games with infinite sets of strategies” was mentioned by Kuhn and Tucker as an interesting open problem in [10]. We will solve both problems (i) and (ii) for a special kind of McNaughton function whose domain allows for decomposition into an orthogonal grid subdividing the hypercube, where a suitable reduction to a matrix game can be used. Finiteness of the mixed equilibria means that the players can select their actions from the finite set with probability one.

This paper is structured as follows. Basic notions concerning Łukasiewicz logic and its functional representation are recalled in Section 2. The game model is investigated and examples are given in Section 3. In Section 4 we state the main result about the existence of finitely supported equilibria for a particular family of McNaughton functions (Theorem 4.2). Finally, the challenges for further research are mentioned in Section 5.

2. Basic notions

We repeat basic definitions and results concerning infinite-valued Łukasiewicz logic. The apparatus of polyhedral geometry, which paves the way for effective processing of McNaughton functions, is further developed. The reader is invited to consult [2] for details. Throughout the paper we use the standard notions and tools of convex geometry—see [22], for example.

2.1. Łukasiewicz logic

We consider only finitely many propositional variables A_1, \dots, A_n . Formulas φ, ψ, \dots are then constructed from these variables and the truth-constant $\bar{0}$ using the following basic connectives: negation \neg and strong disjunction \oplus . The set of all such formulas is denoted by FORM_n .

The standard semantics for connectives of Łukasiewicz logic is given by the corresponding operations of the *standard MV-algebra*, which is just the real unit interval $[0, 1]$ endowed with constant falsum $\bar{0}$ and the operations of negation \neg and strong disjunction \oplus defined as:

$$\bar{0} = 0, \quad \neg a = 1 - a, \quad a \oplus b = \min\{1, a + b\}.$$

Derived connectives are further introduced. They are listed together with their standard semantics and the same symbol is used to denote both the connective and the corresponding binary operation on $[0, 1]$:

strong conjunction	\odot	$\varphi \odot \psi = \neg(\neg\varphi \oplus \neg\psi)$	$\max\{0, a + b - 1\}$
lattice disjunction	\vee	$\varphi \vee \psi = (\varphi \rightarrow \psi) \rightarrow \psi$	$\max\{a, b\}$
lattice conjunction	\wedge	$\varphi \wedge \psi = \neg(\neg\varphi \vee \neg\psi)$	$\min\{a, b\}$
implication	\rightarrow	$\varphi \rightarrow \psi = \neg\varphi \oplus \psi$	$\min\{1, 1 - a + b\}$
equivalence	\leftrightarrow	$\varphi \leftrightarrow \psi = (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$	$1 - a - b $
verum	$\bar{1}$	$\bar{1} = -\bar{0}$	1

Operations \odot and \oplus are also known as Łukasiewicz *t-norm* and Łukasiewicz *t-conorm*, respectively [9]. A valuation is a mapping $v: \text{FORM}_n \rightarrow [0, 1]$ such that for each $\varphi, \psi \in \text{FORM}_n$:

$$v(\bar{0}) = 0, \quad v(\neg\varphi) = 1 - v(\varphi), \quad v(\varphi \oplus \psi) = \min\{1, v(\varphi) + v(\psi)\}.$$

Formulas $\varphi, \psi \in \text{FORM}_n$ are called *equivalent* when $v(\varphi) = v(\psi)$ for every valuation v . The *equivalence class* of φ is denoted by $[\varphi]$. Let us endow the set \mathcal{M}_n of all equivalence classes $[\varphi]$ with the operations

$$\bar{0} = [\bar{0}], \quad \neg[\varphi] = [\neg\varphi], \quad [\varphi] \oplus [\psi] = [\varphi \oplus \psi].$$

Due to the standard completeness theorem [2, Proposition 4.5.5], we may say that the MV-algebra \mathcal{M}_n is the *Lindenbaum algebra* of Łukasiewicz logic over n variables.

Since every valuation v is completely determined by its restriction to the propositional variables $v \mapsto (v(A_1), \dots, v(A_n)) \in [0, 1]^n$, it follows that v is mapped to a unique point $\mathbf{x}_v \in [0, 1]^n$. Conversely, for any $\mathbf{x} \in [0, 1]^n$, let $v_{\mathbf{x}}$ be the valuation uniquely defined by $(v_{\mathbf{x}}(A_1), \dots, v_{\mathbf{x}}(A_n)) = \mathbf{x}$. Therefore, for each formula φ , its equivalence class $[\varphi]$ can be viewed as a function $[\varphi]: [0, 1]^n \rightarrow [0, 1]$ defined through $[\varphi](\mathbf{x}) = v_{\mathbf{x}}(\varphi)$ for every $\mathbf{x} \in [0, 1]^n$. Accordingly, \mathcal{M}_n can be rendered as an algebra of functions with operations \neg and \oplus defined pointwise and with $\bar{0}$ being the constant zero function.

It is easy to prove by induction on the complexity of the formula φ that each function $[\varphi]$ is continuous and piecewise linear such that each linear piece has integer coefficients. The sufficiency of these conditions is the content of the McNaughton Theorem.

McNaughton theorem. (See [12].) Let $f: [0, 1]^n \rightarrow [0, 1]$. Then $f \in \mathcal{M}_n$ iff f is continuous and piecewise linear with each linear piece having integer coefficients.

2.2. McNaughton functions

Each function in \mathcal{M}_n is called an (*n-variable*) *McNaughton function*. Before we proceed further, we need to recall basic geometric facts about McNaughton functions. To this end, a few notions from convex geometry will be useful [22]. By a *convex polytope* in \mathbb{R}^n we mean the convex hull of finitely many points in \mathbb{R}^n . Throughout the paper, we omit the adjective “convex” so that we write just “polytope”. Every nonempty polytope A in \mathbb{R}^n must have at least one *vertex*, a point $\mathbf{v} \in A$ that cannot be expressed as a non-trivial convex combination of points in A . By $V(A)$ we denote the set of all vertices of A . We say that a polytope A is *rational* if $\mathbf{v} \in \mathbb{Q}^n$ for each $\mathbf{v} \in V(A)$.

A *polyhedral complex* in \mathbb{R}^n is a nonempty finite family \mathcal{P} of polytopes in \mathbb{R}^n such that the face of every polytope $A \in \mathcal{P}$ belongs to \mathcal{P} , and the intersection of any two polytopes $A, B \in \mathcal{P}$ is a common face of both A and B . We say that a polyhedral complex \mathcal{P} is *rational* whenever each polytope $A \in \mathcal{P}$ is rational. A *support* of \mathcal{P} is the set $|\mathcal{P}| = \bigcup_{A \in \mathcal{P}} A$. A polyhedral complex \mathcal{P}' is a *subdivision* of \mathcal{P} when $|\mathcal{P}| = |\mathcal{P}'|$ and each polytope of \mathcal{P}' is contained in some polytope of \mathcal{P} . Let $V(\mathcal{P}) = \bigcup_{A \in \mathcal{P}} V(A)$ denote the finite set of *vertices* of polyhedral complex \mathcal{P} .

Every McNaughton function $f \in \mathcal{M}_n$ induces a polyhedral complex whose support is $[0, 1]^n$. Specifically, since f is piecewise linear, there exist distinct linear polynomials f_1, \dots, f_k such that each represents one linear piece of f . For every permutation π of $\{1, \dots, k\}$, let

$$A_{\pi}(f) = \{\mathbf{x} \in [0, 1]^n \mid f_{\pi(1)}(\mathbf{x}) \geq f_{\pi(2)}(\mathbf{x}) \geq \dots \geq f_{\pi(k)}(\mathbf{x})\}.$$

Each $A_{\pi}(f)$ is a rational polytope in $[0, 1]^n$ since it is a set containing solutions of the finite system of linear inequalities with integer coefficients. Given $f \in \mathcal{M}_n$, put

$$\mathcal{P}'(f) = \{A_{\pi}(f) \mid \pi \text{ permutation of } \{1, \dots, k\}\}$$

and consider the collection $\mathcal{P}(f)$ of all faces of the polytopes in $\mathcal{P}'(f)$:

$$\mathcal{P}(f) = \bigcup_{A \in \mathcal{P}'(f)} \{B \mid B \text{ is a face of } A\}.$$

As mentioned in [2, p. 65], the family $\mathcal{P}(f)$ is a rational polyhedral complex such that for each polytope $A \in \mathcal{P}(f)$ there is an index $i \in \{1, \dots, k\}$ for which the functions f and f_i coincide over A . We can refer to polytopes in $\mathcal{P}(f)$ as *linear regions* of f . We write $V(f)$ in place of $V(\mathcal{P}(f))$ without the risk of confusion and the points $\mathbf{x} \in V(f)$ are called the *vertices* of f .




Fig. 1. McNaughton function $f(x, y)$ and the polyhedral complex $\mathcal{P}(f)$.

Example 2.1. Let $f(x, y) = 0 \vee (1 \wedge (2x + 4y - 1))$. Then f is a two-variable McNaughton function—see Fig. 1. This function has exactly three linear pieces: $f_1(x, y) = 0$, $f_2(x, y) = 2x + 4y - 1$, and $f_3(x, y) = 1$, which give rise to the polyhedral complex $\mathcal{P}(f)$ also depicted in Fig. 1.

3. Strategic games with McNaughton payoffs

Game theory admits a presentation of games in an extensive or a strategic form. Extensive games consist of a series of players' moves, which is in fact a series of decision problems with (possibly incomplete) information about the previous moves of the other players. Strategic games, on the other hand, represent games in a concise form. It is assumed that each player makes one decision at a time and chooses a global strategy for the whole game independently of the other players and without information about their decisions.

Let $N = \{1, \dots, n\}$ be a finite set whose elements are called *players*. Each player $i \in N$ can adopt a *strategy* x_i , an element of a possibly infinite set \mathcal{A}_i , in order to achieve his/her goal. A *strategic game* is a triple $\Gamma = (N, (\mathcal{A}_i)_{i \in N}, (f_i)_{i \in N})$, where $f_i : \prod_{j \in N} \mathcal{A}_j \rightarrow \mathbb{R}$ is a payoff function of player $i \in N$, which determines the payoff $f_i(x_1, \dots, x_n)$ provided that player i chooses strategy $x_i \in \mathcal{A}_i$ and the other players select strategies $(x_j)_{j \in N \setminus \{i\}} \in \prod_{j \in N \setminus \{i\}} \mathcal{A}_j$. A *finite game* is a strategic game Γ in which each strategy set \mathcal{A}_i has finite cardinality. Such games have been the subject of study since the pioneering work in game theory by von Neumann and Morgenstern and there are many solution concepts and effective computational methods for finite games. In contrast with finite games requiring only a finite amount of information to describe the players' strategies, we are going to analyze a particular class of infinite strategic games where each player makes strategic choices from the real unit interval $[0, 1]$. We will detail our assumptions as follows.

Definition 3.1. Let $f \in \mathcal{M}_2$ be a two-variable McNaughton function. A *strategic game with McNaughton payoff function* f is the strategic game Γ_f with the player set $N = \{1, 2\}$, the strategy spaces $\mathcal{A}_1 = \mathcal{A}_2 = [0, 1]$, and the payoff functions $f_1 = f$ and $f_2 = \neg f = 1 - f$.

The conditions above mean that every strategic game Γ_f is a game of two players, each of whom selects a real number from $[0, 1]$ as his/her strategy. Moreover, the game is a constant-sum one, $f_1 + f_2 = f + (\neg f) = \bar{1}$, which implies that the interests of the two players are strictly opposing and there is no incentive for cooperation.

We can introduce the natural logical interpretation of the game Γ_f associated with f . Since $f = [\varphi]$ for some Łukasiewicz formula $\varphi \in \text{FORM}_2$, Player 1 seeks a truth valuation ν such that $\nu(\varphi) = f(\mathbf{x}_\nu) = f((x, y)_\nu)$ is maximal, while controlling only the values of the propositional variable x . Hence Player 1 tries to verify the formula. On the contrary, Player 2 seeks a truth valuation ν' which falsifies the formula φ , while controlling only the values of the second propositional variable y . Besides their logical contents, the games with McNaughton functions capture strategic conflicts in which continuous piecewise linear payoff functions are natural. We will illustrate such situations with examples.

Example 3.2 (Model of electoral competition). First, we will define the values of function f_E at vertices of the unit square $[0, 1]^2$:

$$f_E(x, y) = \begin{cases} 0 & x, y \in \{0, 1\}, \\ \frac{1}{2} & x = y = \frac{1}{2}, \\ 1 & x = \frac{1}{2}, y \in \{0, 1\}. \end{cases}$$




Fig. 2. $f_E(x, y)$.




Fig. 3. $f_R(x, y)$.

Second, the function f_E is determined by linear interpolation elsewhere on the unit square. Then $f_E \in \mathcal{M}_2$. The game given by f_E captures the following simplified version of Hotelling's electoral competition model [19]. Assume that Players 1 and 2 are candidates choosing policies $x, y \in [0, 1]$ in order to win the elections. Each citizen has preferences over policies and votes for either Player 1 or 2. In the latest poll, the preferences show that Player 1

1. cannot attract extreme right or extreme left voters at all,
2. is preferred by centrist voters whenever Player 2 chooses any of the two extreme policies,
3. ties with Player 2 if both adopt the centrist policy.

The values $f_E(x, y)$ express the ratio of votes for Player 1—see Fig. 2.

Example 3.3 (Matching reals). This game is a continuous version of the game of Matching pennies [20, Example 17.1]. Player 1 chooses $x \in [0, 1]$. Player 2, not knowing Player 1's choice, also chooses a value $y \in [0, 1]$. The more dissimilar these two choices are, the greater the first player's payoff is: $f_R(x, y) = |x - y|$. On the contrary, the second player tries to minimize their distance since his payoff is given by $\neg f_R(x, y) = 1 - |x - y|$. Since $|x - y| = (x \odot \neg y) \oplus (\neg x \odot y)$, the function f_R is a McNaughton function (Fig. 3), which is known as the *Chang distance* [2]. The corresponding formula is a many-valued counterpart of the Boolean symmetric difference since

$$f_R = [(\varphi \odot \neg \psi) \oplus (\neg \varphi \odot \psi)] = [\neg(\varphi \leftrightarrow \psi)].$$

Dresher discusses a slightly more general version of the game with the McNaughton function f_R in [4, Chapter 8.8].

Which strategies should players choose in order to maximize their payoffs? The standard answer to this fundamental problem requires introducing the main solution concept of game theory.

Definition 3.4 (Nash equilibrium). Let $f \in \mathcal{M}_2$. A *Nash equilibrium* of a game with McNaughton function Γ_f is a pair of strategies $(x^*, y^*) \in [0, 1]^2$ such that

$$f(x, y^*) \leq f(x^*, y^*) \leq f(x^*, y), \quad \text{for every } (x, y) \in [0, 1]^2. \tag{1}$$

When both players adopt the Nash equilibrium strategies x^* and y^* , respectively, then each of them avoids the situation with incurring a smaller payoff than that guaranteed by the value $f(x^*, y^*)$. If a player's strategy is a part of a certain Nash equilibrium, then we say that the strategy is *optimal* for the player. For every game Γ_f , we denote

$$NE(f) = \{(x^*, y^*) \in [0, 1]^2 \mid (x^*, y^*) \text{ satisfies (1)}\}.$$

For any $(x^*, y^*) \in NE(f)$, the number $f(x^*, y^*)$ is called the *value* of game Γ_f .

Example 3.5 (*Example 3.2 contd.*). It can easily be verified that the only pair of optimal strategies in the game Γ_{f_E} is $(\frac{1}{2}, \frac{1}{2})$, thus $NE(f_E) = \{(\frac{1}{2}, \frac{1}{2})\}$. Indeed, the center of the square is the only saddle point of f_E .

We will see later that existence of a Nash equilibrium in a strategic game Γ_f is, in general, not guaranteed. However, if $NE(f) \neq \emptyset$, then the assumption that f is a McNaughton function implies a special form of this set.

Proposition 3.6 (*Properties of Nash equilibria*). Let $f \in \mathcal{M}_2$. If the strategic game Γ_f admits at least one Nash equilibrium, then:

1. $NE(f) \cap V(f) \neq \emptyset$. At least one equilibrium pair lies in a vertex of the polyhedral complex corresponding to f .
2. The value of Γ_f is a rational number.
3. $NE(f)$ is a finite union of rational rectangles.
4. If $\dim(NE(f)) = 2$, then the value of game Γ_f is either 0 or 1.

Proof. 1. Define

$$\begin{aligned} X_f &= \{x \in [0, 1] \mid (x, y) \in V(f) \text{ for some } y \in [0, 1]\}, \\ Y_f &= \{y \in [0, 1] \mid (x, y) \in V(f) \text{ for some } x \in [0, 1]\}. \end{aligned} \tag{2}$$

Let $(x^*, y^*) \in NE(f)$ and $v = f(x^*, y^*)$. The function $g(x) = f(x, y^*)$ is continuous and piecewise linear. Hence there exists some $x_i \in X_f$ such that

$$f(x, y^*) \leq f(x_i, y^*), \quad \text{for every } x \in [0, 1],$$

and thus $f(x_i, y^*) = v$. Analogously, the function $h(y) = f(x_i, y)$ must attain its minimum at a point $y_j \in Y_f$, so that $f(x_i, y_j) = v$. Since $x_i \in X_f$, the pair (x_i, y_j) is a vertex of $\mathcal{P}(f)$.

2. The value v is uniquely determined by Γ_f . In particular, $v = f(x^*, y^*)$ for any $(x^*, y^*) \in NE(f) \cap V(f)$. Since the vertices of f are rational points and the linear pieces of function f have integer coefficients, the value v is necessarily rational.

3. It is well-known that the set of all Nash equilibria of a constant-sum game is a Cartesian product, that is, $NE(f) = X \times Y$, for some $X, Y \subseteq [0, 1]$. Let v be the value of game Γ_f and consider the set $M_{y^*} = \{x \in X \mid f(x, y^*) = v\}$ for a fixed optimal strategy y^* of Player 2. Since f is continuous piecewise linear with finitely many linear pieces, we have

$$M_{y^*} = \bigcup_{i=1}^m [a_i, b_i]$$

for certain values $a_i, b_i \in [0, 1]$. The endpoints of each interval must be rational since the value is rational and the function f has integer coefficients. The conclusion follows by observing that $X = M_{y^*}$ for every choice of the optimal y^* and by repeating the same construction for Y .

4. Suppose that $\dim(NE(f)) = 2$. Let $v = f(x^*, y^*)$, where $(x^*, y^*) \in NE(f)$. Since the value of the game Γ_f is uniquely determined and $\dim(NE(f)) = 2$, the function f must be constant over $NE(f)$, and the number v must satisfy $v \in \{0, 1\}$. Indeed, otherwise the function f would have a linear piece given by a real constant $c \in (0, 1)$, and a contradiction would follow from the definition of McNaughton function f (all pieces have integer coefficients). \square

Observe that the previous theorem gives a very easy but inefficient method for finding a Nash equilibrium in Γ_f : the exhaustive search over the vertices in $V(f)$. However, this technique may fail badly when Γ_f has no equilibria—examples of McNaughton functions f with this property are not difficult to find.

Example 3.7 (*Example 3.3 contd.*). The game associated with Chang distance $f_R(x, y) = |x - y|$ does not have any Nash equilibrium. Indeed, for every $x, y \in [0, 1]$ there exists x' such that $f_R(x', y) > f_R(x, y)$ or there exists y' such that $f_R(x, y') > f_R(x, y)$. We can distinguish three cases. Let $x, y \in \{0, 1\}$ and suppose that $x = y = 0$. Then it suffices to put $x' = 1$ and similarly for other cases when x and y are 0 or 1. Let $0 < x \leq y < 1$. Then Player 1 is better off by selecting $x' = 1$. If $0 < y \leq x < 1$, then Player 1 can strictly improve his payoff under $x' = 0$.

Strategies x in the $[0, 1]$ are also called *pure strategies*. When a game allows for no equilibrium in pure strategies, players may still randomize over them. This leads to the introduction of the Nash equilibrium in mixed strategies.

Since the strategy space $[0, 1]$ is uncountable, it is theoretically possible that any measurable set of strategies can be adopted with nonzero probability. Thus it is natural to define a *mixed strategy* as a Borel probability measure μ on $[0, 1]$. By Δ we denote the set of all mixed strategies, which is a weak*-compact and convex subspace of the Banach space of all Borel measures over $[0, 1]$. If Player 1 and Player 2 use mixed strategies μ and ν , respectively, then the *expected payoff* of Player 1 in the game Γ_f is

$$E_f(\mu, \nu) = \int_{[0,1]^2} f \, d(\mu \times \nu), \tag{3}$$

where $\mu \times \nu$ denotes the product measure on $[0, 1]^2$. If one of the players, say Player 1, uses a pure strategy $x \in [0, 1]$, then we may identify x with the Dirac measure δ_x and (3) becomes

$$E_f(\delta_x, \nu) = \int_{[0,1]} f(x, y) \, d\nu(y).$$

Throughout the rest of this article the notation $E_f(x, \nu)$ is used instead of $E_f(\delta_x, \nu)$ above. We say that a pair of Borel probability measures $(\mu^*, \nu^*) \in \Delta^2$ is a *mixed-strategy equilibrium* if

$$E_f(\mu, \nu^*) \leq E_f(\mu^*, \nu^*) \leq E_f(\mu^*, \nu), \quad \text{for every } (\mu, \nu) \in \Delta^2. \tag{4}$$

We can test the property (4) with respect to the pure strategies only.

Proposition 3.8. *Let $f \in \mathcal{M}_2$. Then the following assertions are equivalent:*

1. $(\mu^*, \nu^*) \in \Delta^2$ is a mixed-strategy equilibrium in Γ_f .
2. For every pair of pure strategies $(x, y) \in [0, 1]^2$,

$$E_f(x, \nu^*) \leq E_f(\mu^*, \nu^*) \leq E_f(\mu^*, y).$$

Proof. The only non-trivial direction is a direct consequence of weak*-compactness of the space Δ and the fact that all convex combinations of Dirac measures are weak*-dense in Δ [1, Theorem 15.10]. \square

Does every game Γ_f corresponding to a McNaughton function $f \in \mathcal{M}_2$ possess a mixed strategy equilibrium? The existence of the equilibrium pair is a direct consequence of Glicksberg's Theorem [5], which requires only compactness of the strategy spaces and continuity of the payoff functions. Therefore we obtain the following existence result.

Theorem 3.9. *Let $f \in \mathcal{M}_2$. Then the game Γ_f has a mixed strategy equilibrium.*

While the concept of mixed strategies is general enough to provide a solution to any game Γ_f , the equilibrium pair (μ^*, ν^*) is not amenable to a direct calculation as the proof of Glicksberg's Theorem is not constructive. As far as the authors' knowledge goes, in general, there is no algorithm for computing the mixed equilibria in games with continuous payoff functions over the unit interval. The existing methods are well-tailored to very specific classes of functions, such as polynomial and separable games [15], which do not include the class of games with McNaughton functions discussed in this paper. Therefore, in order to recover mixed equilibria in the examples below, we will first rely on a combination of a guess and good luck.

Example 3.10 (Example 3.3 contd.). The game given by f_R does not have a unique pair of equilibrium in mixed strategies, unlike the classical game of Matching pennies. It is easy to see that the mixed strategy $\mu^* = \frac{1}{2}(\delta_0 + \delta_1)$ is optimal for both players so that the value of the game is

$$E_f(\mu^*, \mu^*) = \sum_{(x,y) \in \{0,1\}^2} \frac{f_R(x, y)}{4} = \frac{1}{2}.$$

Another equilibrium forms the pair $(\mu^*, \frac{1}{2})$. Note that the optimal strategies are not symmetric since $(\frac{1}{2}, \mu^*)$ is not a pair of equilibrium strategies. We can show that the uniform distribution—the Lebesgue measure λ on $[0, 1]$ —is never optimal for Player 1: indeed, we have

$$E_{f_R}(\lambda, y) = y^2 - y + \frac{1}{2} \leq \frac{1}{2},$$




Fig. 4. $g_R(x, y)$.

which gives $E_{f_R}(\lambda, y) < \frac{1}{2}$ for every y such that $0 < y < 1$. On the contrary, the pair (μ^*, λ) is in equilibrium since $E_{f_R}(\mu^*, y) = \frac{1}{2}$, for every $y \in [0, 1]$, and

$$E_{f_R}(x, \lambda) = x^2 - x + \frac{1}{2} \leq \frac{1}{2}, \quad \text{for every } x \in [0, 1].$$

While the game Γ_{f_R} given by the Chang distance is fair since its value is equal to $\frac{1}{2}$, the two players' sets of optimal strategies are not the same. Below we will introduce a more convincing continuous generalization to the game of Matching pennies: the payoff must measure the degree of similarity and dissimilarity at the same time.

Example 3.11. Put $g_R = 2(f_R \wedge \neg f_R)$. The payoff function of the first player in the game given by g_R is thus

$$g_R(x, y) = 2 \min\{|x - y|, 1 - |x - y|\},$$

see Fig. 4. This McNaughton function corresponds to the formula

$$((\varphi \leftrightarrow \psi) \wedge (\varphi \leftrightarrow \psi)) \oplus ((\varphi \leftrightarrow \psi) \wedge (\varphi \leftrightarrow \psi))$$

in which the double appearance of $(\varphi \leftrightarrow \psi) \wedge (\varphi \leftrightarrow \psi)$ is for normalization purposes only. In game Γ_{g_R} it is optimal for each player to select a pure strategy in $[0, 1]$ according to the uniform distribution over the unit interval. Thus one Nash equilibrium pair is (λ, λ) and the value of Γ_{g_R} is

$$E_{g_R}(\lambda, \lambda) = \frac{1}{2}.$$

However, other equilibria exist too. For every $a, b \in [0, 1]$ with $a - b = \frac{1}{2}$, put $\mu_{a,b} = \frac{1}{2}(\delta_a + \delta_b)$. Then we can routinely check that the pair $(\mu_{a,b}, \mu_{a,b})$ is equilibrium.

All the examples of games shown so far reveal the existence of a special pair of optimal strategies in which the players randomize among finitely many pure strategies. Specifically, we say that a mixed strategy $\mu \in \Delta$ is *finitely supported* if the support of μ is finite, that is,

$$\text{spt } \mu = \bigcap \{A \subseteq [0, 1] \mid \mu(A) = 1, A \text{ closed}\} \text{ is finite.} \tag{5}$$

The games in which optimal strategies μ exist such that (5) holds true have convenient algorithmic properties, since processing Borel probability measures is in general a very difficult task. It is known that not every game with continuous payoff functions admits at least one pair of finitely supported strategies [6]. On the other hand, the optimal finitely supported strategies are present in the class of separable games for which even computational algorithms exist [17,18]. Deciding whether a given game over the unit square has a pair of optimal finitely supported strategies thus remains a difficult task.




Fig. 5. Grid $\mathcal{G}(f)$ and complex $\mathcal{P}(f)$ (in gray) of the function from Fig. 1.




Fig. 6. Grid $\mathcal{G}(g)$ and complex $\mathcal{P}(g)$ (dashed) from Example 4.1.

We are going to address this question in the next section, where our goal is to single out a class of games with McNaughton functions for which finitely supported optimal strategies exist.

4. Existence and computation of finitely supported mixed strategy equilibria

As a point of departure, notice that the optimal strategies in the games discussed above can be chosen so that the players select the pure strategies only from those points that are projections of the vertices $V(f)$. Recalling the notation from (2), this means that the supports of such optimal strategies are contained in the sets X_f and Y_f , respectively. It turns out that there is a class of McNaughton functions with a special geometric structure that makes it possible to directly recover finite mixed strategies.

Let $f \in \mathcal{M}_2$ and let $\mathcal{P}(f)$ be the complex of linear regions of f . Without loss of generality, we assume that the elements of vertex projections X_f and Y_f are

$$0 = x_1 \leq \dots \leq x_p = 1 \quad \text{and} \quad 0 = y_1 \leq \dots \leq y_q = 1, \tag{6}$$

respectively, where the natural numbers $p, q \geq 2$. The *orthogonal grid* of f is the polyhedral complex $\mathcal{G}(f)$ such that

1. $\{x_i\} \times [y_j, y_{j+1}] \in \mathcal{G}(f)$ for every $i = 1, \dots, p$ and $j = 1, \dots, q - 1$,
2. $[x_i, x_{i+1}] \times \{y_j\} \in \mathcal{G}(f)$ for every $i = 1, \dots, p - 1$ and $j = 1, \dots, q$,
3. $(x_i, y_j) \in \mathcal{G}(f)$ for every $i = 1, \dots, p$ and $j = 1, \dots, q$.

The motivation for the introduction of the grid $\mathcal{G}(f)$ is to investigate whether function f is linear over the line segments connecting the vertices of f —see Fig. 5. The next example shows that a McNaughton function need not be linear over the line segments of the grid. Let co denote the convex hull operator.

Example 4.1. Let

$$\begin{aligned} T_1 &= \text{co}\{(0, 0), (0, 1), (1, 1)\}, \\ T_2 &= \text{co}\{(0, 0), (\frac{2}{3}, \frac{1}{3}), (1, 1)\}, \\ T_3 &= \text{co}\{(0, 0), (\frac{2}{3}, \frac{1}{3}), (1, 0)\}, \\ T_4 &= \text{co}\{(1, 1), (\frac{2}{3}, \frac{1}{3}), (1, 0)\}, \end{aligned} \quad g(x, y) = \begin{cases} y - x & (x, y) \in T_1, \\ 3x - 3y & (x, y) \in T_2, \\ x + y & (x, y) \in T_3, \\ 2 - x - y & (x, y) \in T_4. \end{cases}$$

See Fig. 6.

Theorem 4.2. Let $f \in \mathcal{M}_2$ be such that every line segment in $\mathcal{G}(f)$ is included in some polytope of $\mathcal{P}(f)$. Then the game Γ_f has Nash equilibrium consisting of finitely supported mixed strategies.

Proof. We define the following two-player constant-sum game: the strategy spaces are X_f and Y_f , the payoff function of Player 1 is the restriction f_0 of f to $X_f \times Y_f$, and the payoff function of the second player is $-f_0$. Since each strategy set is finite, this game is solvable in mixed strategies by the Nash Theorem [14]. Assume that an optimal pair of strategies of Players 1 and 2 is given by the probability vectors $(\alpha_1, \dots, \alpha_p)$ and $(\beta_1, \dots, \beta_q)$, respectively. Put

$$\mu^* = \sum_{i=1}^p \alpha_i \delta_{x_i} \quad \text{and} \quad \nu^* = \sum_{j=1}^q \beta_j \delta_{y_j},$$

where x_i and y_j are as in (6). We will show that $(\mu^*, \nu^*) \in \Delta^2$ is a finitely supported Nash equilibrium in the game Γ_f . First, we have to show that the inequality

$$E_f(\mu^*, \nu^*) \geq E_f(x, \nu^*) \tag{7}$$

is true for every choice of pure strategy $x \in [0, 1]$. Clearly, the inequality

$$E_f(\mu^*, \nu^*) \geq E_f(x_i, \nu^*) \tag{8}$$

holds true for every $i = 1, \dots, p$ since (μ^*, ν^*) is an equilibrium of the game associated with f_0 and the value of this game is $E_f(\mu^*, \nu^*)$. For every $x \in [0, 1]$, there exist some $\gamma \in [0, 1]$ and some $i = 1, \dots, p - 1$ such that $x = \gamma x_i + (1 - \gamma)x_{i+1}$. Hence

$$E_f(x, \nu^*) = \sum_{j=1}^q \beta_j f(x, y_j) = \sum_{j=1}^q \beta_j f(\gamma x_i + (1 - \gamma)x_{i+1}, y_j). \tag{9}$$

For each $j = 1, \dots, q$, the line segment with endpoints (x_i, y_j) and (x_{i+1}, y_j) is included in some polytope $A_j \in \mathcal{P}(f)$. As the function f is linear over A_j , the last sum in (9) becomes

$$\begin{aligned} \sum_{j=1}^q \beta_j (\gamma f(x_i, y_j) + (1 - \gamma)f(x_{i+1}, y_j)) &= \gamma \sum_{j=1}^q \beta_j f(x_i, y_j) + (1 - \gamma) \sum_{j=1}^q \beta_j f(x_{i+1}, y_j) \\ &= \gamma E_f(x_i, \nu^*) + (1 - \gamma)E_f(x_{i+1}, \nu^*) \\ &\leq \gamma E_f(\mu^*, \nu^*) + (1 - \gamma)E_f(\mu^*, \nu^*) = E_f(\mu^*, \nu^*), \end{aligned}$$

where the inequality follows from (8). This proves (7). The proof that ν^* is an optimal strategy for the second player is completely analogous and we skip it. \square

Observe that in the proof we have not used the assumption about the integer values of linear coefficients of McNaughton function f . Hence the theorem holds true even for piecewise linear continuous functions with real coefficients. Moreover, for the sake of simplicity, we have always assumed that there are only two players in the game and that the game is strictly competitive, i.e., the sum of players' payoffs remains constant for all strategy choices. These conditions can be relaxed so that the results obtained above remain unchanged even for finitely many players $1, \dots, m$ and piecewise linear continuous payoff functions f_1, \dots, f_m . We will formulate this precisely—the proof is, however, only an inessential modification of the proof of Theorem 4.2.

Corollary 4.3. Let Γ be a strategic game with n players, each having the strategy space $[0, 1]$ and the payoff function $f_i: [0, 1]^n \rightarrow [0, 1]$ that is continuous and piecewise linear. If every line segment in each $\mathcal{G}(f_i)$ is included in some polytope of $\mathcal{P}(f_i)$, then the game Γ_f has Nash equilibrium consisting of finitely supported mixed strategies.

The proof of Theorem 4.2 provides an algorithm for finding Nash equilibria in Γ_f . Indeed, once we check that function f satisfies the assumption about the associated grid, it suffices to recover the equilibria in a finite matrix constant-sum game, which is a standard task solved by linear programming [21, Chapter III].

Not every game Γ_f meets the condition of Theorem 4.2—see Example 4.1. Even in such cases, however, finite optimal strategies may exist.

Example 4.4. Put

$$\begin{aligned} T_1 &= \text{co}\{(0, 0), (0, 1), (1, 1)\}, \\ T_2 &= \text{co}\{(0, 0), (1, \frac{1}{2}), (1, 1)\}, \\ T_3 &= \text{co}\{(0, 0), (1, \frac{1}{2}), (1, 0)\}, \end{aligned} \quad f(x, y) = \begin{cases} y - x & (x, y) \in T_1, \\ 0 & (x, y) \in T_2, \\ x - 2y & (x, y) \in T_3. \end{cases}$$

The game Γ_f meets the assumptions of [16, Theorem 1], where the payoff function of Player 1 is supposed to be convex, the game is constant-sum and hence the payoff function of Player 2, which is $\neg f = 1 - f$, is concave. This implies that there exists a pair of optimal strategies (μ^*, y) such that the support of measure μ^* has at most two elements and $y \in [0, 1]$ is a pure strategy of Player 2.

In the next example we show that there are games with McNaughton functions for which both the condition of Theorem 4.2 and [16, Theorem 1] fail.

Example 4.5. Consider the function g from Example 4.1. It is easy to see that at least one line segment in $\mathcal{G}(g)$ is not fully contained in one of the polytopes of $\mathcal{P}(g)$. Moreover, the one-variable functions $g(x, \cdot)$ and $g(\cdot, y)$ are neither convex nor concave and thus [16, Theorem 1] does not apply to this case. We may, however, routinely verify that the optimal mixed strategy for each player is $\frac{1}{2}(\delta_0 + \delta_1)$.

5. Conclusions

Our aim was to characterize the class of strategic games with McNaughton payoffs which have a finitely supported Nash equilibrium. We have shown that existence of an orthogonal grid is a sufficient condition for finite equilibria and we have proposed an algorithm for recovering the corresponding equilibrium strategies. Our result can easily be generalized to n -player strategic games which do not need to be strictly competitive with payoff functions represented by piecewise linear functions with real coefficients. Proving/disproving the conjecture about existence of finitely supported strategies for every McNaughton function remains the main challenge and we leave it for future investigations. In our future research we would also like to explore the relationships of the strategic games we proposed to the background logic. It is, for example, easy to see that games with monotone McNaughton payoff functions correspond to the formulas without negation and implication, and always have an equilibrium payoff of 0 or 1. The question is whether there are any other classes of formulas such that the corresponding games have a specific form of equilibria.

Acknowledgements

The work of both authors was supported by Grant No. P402/12/1309 of the Czech Science Foundation.

References

- [1] C.D. Aliprantis, K.C. Border, *Infinite Dimensional Analysis: A Hitchhiker's Guide*, third edition, Springer, Berlin, 2006.
- [2] R.L.O. Cignoli, I.M.L. D'Ottaviano, D. Mundici, *Algebraic Foundations of Many-Valued Reasoning*, Trends in Logic—Studia Logica Library, vol. 7, Kluwer Academic Publishers, Dordrecht, 2000.
- [3] M. Fedel, H. Hosni, F. Montagna, A logical characterization of coherence for imprecise probabilities, *Int. J. Approx. Reason.* 52 (8) (2011) 1147–1170.
- [4] M. Drescher, *Games of Strategy: Theory and Applications*, Prentice-Hall Applied Mathematics Series, Prentice Hall Inc., Englewood Cliffs, NJ, 1961.
- [5] I.L. Glicksberg, A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points, *Proc. Am. Math. Soc.* 3 (1952) 170–174.
- [6] I. Glicksberg, O. Gross, Notes on games over the square, in: H. Kuhn, A. Tucker (Eds.), *Contributions to the Theory of Games*, vol. II, Princeton University Press, 1950, pp. 173–183.
- [7] P. Hájek, *Metamathematics of fuzzy logic*, Trends in Logic—Studia Logica Library, vol. 4, Kluwer Academic Publishers, Dordrecht, 1998.
- [8] P. Harrenstein, W. van der Hoek, J.-J. Meyer, C. Witteveen, Boolean games, in: *Proceedings of the 8th Conference on Theoretical Aspects of Rationality and Knowledge, TARK '01*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2001, pp. 287–298.
- [9] E.P. Klement, R. Mesiar, E. Pap, *Triangular Norms*, Trends in Logic—Studia Logica Library, vol. 8, Kluwer Academic Publishers, Dordrecht, 2000.
- [10] H.W. Kuhn, A.W. Tucker (Eds.), *Contributions to the Theory of Games*, Princeton University Press, Princeton, NJ, 1950.
- [11] E. Marchioni, M. Wooldridge, Łukasiewicz games, in: *Proceedings of the Thirteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2014)*, 2014, in press.
- [12] R. McNaughton, A theorem about infinite-valued sentential logic, *J. Symb. Log.* 16 (1951) 1–13.
- [13] D. Mundici, Bookmaking over infinite-valued events, *Int. J. Approx. Reason.* 43 (3) (2006) 223–240.
- [14] J. Nash, Non-cooperative games, *Ann. Math. (2)* 54 (1951) 286–295.
- [15] P. Parrilo, Polynomial games and sum of squares optimization, in: *45th IEEE Conference on Decision and Control*, 2006, pp. 2855–2860.
- [16] T. Parthasarathy, T.E.S. Raghavan, Equilibria of continuous two-person games, *Pac. J. Math.* 57 (1) (1975) 265–270.
- [17] N. Stein, *Characterization and computation of equilibria in infinite games*, PhD thesis, Cornell University, 2005.
- [18] N.D. Stein, A. Ozdaglar, P.A. Parrilo, Separable and low-rank continuous games, *Int. J. Game Theory* 37 (4) (2008) 475–504.
- [19] M.J. Osborne, *An Introduction to Game Theory*, Oxford University Press, 2004.
- [20] M.J. Osborne, A. Rubinstein, *A Course in Game Theory*, MIT Press, Cambridge, MA, 1994.
- [21] G. Owen, *Game Theory*, third edition, Academic Press Inc., San Diego, CA, 1995.
- [22] G. Ziegler, *Lectures on Polytopes*, Graduate Texts in Mathematics, vol. 152, Springer-Verlag, New York, 1995.