SOME THEORETICAL RESULTS CONCERNING AUTOMATED GAUE-PLAYING

T. B. Boffey
University of Liverpool,
Liverpool, England.

Abstract

Though evaluation functions have been used in many game-playing and problem-solving programs the aim has generally been to construct suitable functions for particular situations. A formal framework is outlined in which general properties of evaluation functions can be studied. An application of these concepts is made to the study of when decision rules (as exhibited by a players actual play) can be represented by evaluation functions

1. Introduction.

When computers have been used to play games such as chess, checkers etc., explicit use of evaluation functions (most frequently linear in form) has usually been made. They have also featured in the solution of non-game problems[1]. In other cases a decision rule for playing games (and solving problems) has been used, but it is shown in 3 how decision rules are in many cases, equivalent to the use of suitable EF's (from now on the abbreviation EF will be used for evaluation function). EF s have been used in many circumstances and in varying ways and lie at the heart of programming machines to play games and solve problems, and their importance was stressed as early as 1950 by Shannon[2]. However, on the whole they have been introduced for specific applications and their meaning and properties neglected other than that a position or state with a higher value is likely to be a better one.

It is the purpose of this paper to introduce a formal framework in which a general study of EF's can be made. The full length version of this paper shows how this framework can be used when studying existing techniques and how it is potentially useful for introducing new techniques. In this paper the basic concepts of EF's are presented and then a single application is choBen for further consideration, namely the study of when a players actual play can be represented by an EF. The main results are given in Theorems 3.1.3, 3.2.4 and 3.2.2.

2. Basic Properties of Evaluation Functions.

A game usually consists of a set of definite states together with transitions between them. Consequently it is natural to represent a game by means of a directed grapht[1] (X, P) whose vertices correspond to the states or position! and whose arcs correspond to the possible transitions or moves. However, it may be of importance to consider the development of a game (or the nature of the solution path in problem-solving) in which case the more usual formulation in terms of a (rooted) tree (V, T) is likely to be advantageous. (V, T) can be obtained from (X,P) by starting at the root (the start position) and 'growing' the tree according to the successor mapping except that a vertex is never identified with a previously considered one. (X, P) will be used when the discussion concerns graphs (and also trees as a special case) and (V, T) will be used if trees only are involved.

The games considered in this paper will be 2-person games with a fixed start position between players whom we name MAX and MIN. Moreover, it will be assumed that the games are finite in the sense that their game trees have a finite number of vertices, that MAX always has the first move, and thereafter every other move. Most discussions will be in terms of the game as seen from MAX's point of view. Certain vertices of (X, Γ) will correspond to a state in which the game is over. These vertices, called terminals, will have no successors, and a final set F is defined by

$$F = \{x \mid \Gamma x = \phi\}.$$

In a game between two players there will, in general, be several different pay-offs (frequently 'win', 'draw' and 'loss'), that is there is defined a function

$$\lambda : F \rightarrow \mathcal{P}$$

where $\mathcal{P}$ is the pay-off set. However, once a terminal position has been reached the situation is clear, and what is really required is some means of assessing non-terminal positions with respect to the likely outcomes from those positions.

A value can (in principle) be assigned to any vertex by backing-up from the terminal vertices using some suitable procedure. However, the most commonly used procedure (the minimax procedure cf. § 2.1) requires maxima and minima to be found over subsets of $\mathcal{P}$, and so $\mathcal{P}$ must possess an order structure. The minimum requirement, to avoid 'inconsistencies', is that $\mathcal{P}$ should be weakly ordered. Other backing-up rules may involve the use of probability combinations $\langle u|p|v\rangle$ (representing the gamble in which u is chosen with probability p and v with probability 1 - p). For such cases it is sufficient for the von-Neumann-Morgenstern axioms[5] to be satisfied over the set of gambles $\mathcal{P}^*$.

It is possible to find an order-preserving function $\varphi : \mathcal{P} \rightarrow R$ (or $\mathcal{P}^* \rightarrow R$) where R is the set of real numbers, and nothing is lost by replacing values $a \in \mathcal{P}$ (or $\mathcal{P}^*$) by the values $\varphi(a)$. In what follows it will be assumed that this transformation has been made, and that all pay-offs are real numbers.

2.1. Evaluation functions.

Backing up procedures are used to obtain values at non-terminal vertices, that is to extend the mapping $\lambda : F \rightarrow R$ to a mapping $\chi^* : X \rightarrow R$ with $\chi^*|_F = \lambda$. It is convenient at this point to define two subsets of X as follows.

Definition 2.1.1. The maxset X^M is that subset of X at which it is MAX's turn to make a move. The minset X^m is that subset at which it is MIN's turn to make a move.

The following result is immediate.

Theorem 2.1.1. (i) $X^M \cup X^m = X$

(ii) $X^M \cap X^m = \phi$ if (X, Γ) has no cycles with an odd number of arcs.

For graph-theoretic concepts see Berge[4].

For game graphs (X, Γ^+) satisfying $X^m \cap X^{m} = \phi$, λ^* can be extended as follows:

1. $\lambda^*(x) = \lambda(x)$ if $x \in F$;
2. Suppose $\lambda^*(y)$ is known for all $y \in \Gamma x$ then
   (i) $\lambda^*(x) = \max_{y \in \Gamma x} \lambda^*(y)$ if $x \in X^M$
   (ii) $\lambda^*(x) = \min_{y \in \Gamma x} \lambda^*(y)$ if $x \in X^m$.

λ^* formed in this will be called the minimax extension of λ and will be denoted by λ^*_{mm}. In the case of game trees the above procedure is just the familiar minimax rule. However, with the above formulation, it is easily seen that the minimax extension can be formed for any game graph satisfying $X^M \cap X^m = \phi$. (Boffey[4] gives a fuller discussion of the class of graphs over which the minimax rule is applicable).

From now on it will be assumed that all graphs satisfy the condition $X^M \cap X^m = \phi$ unless otherwise stated.

Although, in principle, such backing-up can always be performed over a game graph this is not, in general, practically feasible. For example Good[7] estimates that there are about 10^{44} possible chess positions and so the associated game graph will have 10^{44} vertices. In such cases estimates of the values of positions encountered may be considered. These lead, in principle, to a function $f: X \to R$ (where R is the set of real numbers) such that f in some sense approximates λ^*_{mm}.

Definition 2.1.2. An evaluation function (abbreviated EF) over a game graph (X, Γ^+) is a single valued function $f : X \to R$.

If $f(x) > f(y)$ then x is said to be better than y (under f).

The merits of a particular EF will depend upon the way it is to be used. If it is suspected that the EF is a good approximation to λ^*_{mm} (in a sense to be made more precise below) and the opponent is a 'minimax like' player then it is natural to use the strategy Δ_0 where

Definition 2.1.3. Δ_0 is that strategy which leads to a move xy^* being made from position x where

$$f(y^*) = \max_{y \in \Gamma x} f(y) \quad \text{if } x \in X^M,$$

$$\text{and} \quad f(y^*) = \min_{y \in \Gamma x} f(y) \quad \text{if } x \in X^m.$$

Ties for y^* are broken by some subsidiary rule.

Definition 2.1.4. EF's f and g are locally equivalent if

$$f(x) > f(y) \quad \text{whenever} \quad g(x) > g(y)$$

where $x, y \in \Gamma w$ for some $w \in X$. f is locally perfect if it is locally equivalent to λ^*_{mm}.

Theorem 2.1.2. Using strategy Δ_0 locally equivalent EF's lead to the same play, provided the same subsidiary rule is used to split ties. A locally perfect EF leads to optimal (minimax) play.

2.2. Probabilistic evaluation functions.

The above discussion is relevant to the situation in which a player always chooses the same move from a given position. However, in practice this will generally not be the case, a player's move depending partly on his assessment of his opponents ability and temperament, partly on his frame of mind at the time of the game and partly on other factors such as a draw in the last round of a tournament being sufficient to secure first place. A player will also vary his play simply for the sake of variety. Although the definition of a game could be extended to incorporate 'environmental' factors this will not be attempted here. Instead a model is chosen which will allow players to use variable strategies with the actual choices of moves being governed by probability distributions. Such a model provides a generalization of the usual concept of automated game-playing and will be discussed further in §3.2.

It is now appropriate to relax the constraint that each position should have a single value associated with it, and to replace this by the condition that each position has a probability distribution attached to it.

Definition 2.2.1. A probabilistic evaluation function (PEF) over a game graph (X, Γ^+) is a bounded function $p: X \times R \to R$, with $p(x,r) \geq 0$,

$$\int p(x,r)\, dr = 1$$

and $p(x, r) = 0$ outside some finite range $r_1(x) \leq r \leq r_2(x)$.

$p(x, r)dr$ can be interpreted as the probability that the value of x, now denoted by $v(x)_p$, lies in the range $r \leq v(x)_p \leq r + dr$. When required $v(x)_p$ is obtained by selecting, at random, a number from the distribution $p(x, \ )$. Now just as for EF's, a PEF is meaningless unless a strategy for its use is specified. We shall only consider strategy Δ_0^p, the analogue of Δ_0.

Definition 2.2.2. Δ_0^p is that strategy which leads to a move xy^* being made from position x where

$$v(y^*)_p = \max_{y \in \Gamma x} v(y)_p \quad \text{if } x \in X^M$$

$$v(y^*)_p = \min_{y \in \Gamma x} v(y)_p \quad \text{if } x \in X^m.$$

Notice how the successor position y^* selected by Δ_0^p can vary from game to game due to the fact that the $v(y)_p$ are stochastic variables.

§3.2 is devoted to a study of how PEF's can be used to represent actual play.

3. Decision Rules.

Game graphs are usually very large and game trees tend to be even larger. It is therefore (as noted in §2) not in general feasible to consider the whole of (X, Γ^+) in order to derive values of λ^*_{mm}. Consider for example the case where (X, Γ^+) is a tree with a constant branching factor b and constant depth D (ie. every path from the root of the tree to a leaf contains D arcs). Then the final set F will contain b^D elements and more than b^D vertices will have to be considered in order to obtain $\lambda^*_{mm}(S)$. On the other hand a decision rule, which can be applied at any vertex, need only be applied D times during the course of a single game. Thus for a single game the ratio of the amounts of computation required by the two approaches is less than $k \times D : b^D$ where k is a measure of the effort involved in a single application of the decision rule. It follows that the decision rule requires much less computational effort to apply if k is not too large.

Sometimes a decision rule may be implicit, as is the case for example when one is observing the play of a master (whether directly or from written records) with a view to improving ones own play. In this situation it is relevant to pose the following question:- how can the underlying decision rule as exemplified by the masters observed choices of moves be characterized, and under what circumstances can such a decision rule be represented by an EF (with strategy Δ_0). Although the masters play is likely to be observed over only a small subset of all possible positions, the more constrained problem of establishing the existence of an EF to represent a decision rule over the whole of (X, Γ) will be tackled.

3.1. Representing the decision rule of a constant player.

A master will in general not always play the same move from a given position (for example a chess player who likes to open P-K4 is likely to vary his play by sometimes opening P-Q4 or P-QB4 etc.). Such play will be termed variable play and is treated in § 3.2. In this section it will be assumed that at a given position a player will always make the same move.

Definition 3.1.1. A (constant) decision rule or preference function for MAX is a single valued function $\sigma_M : X \to X$ with $\sigma_M(x) \in \Gamma x$. $x \to \sigma_M(x)$ is MAX's preferred move from x.

A preference function σ_m for MIN can be defined in an analogous way.

Definition 3.1.2. If σ_M and σ_m are preference functions for MAX and MIN respectively then $\sigma = (\sigma_M, \sigma_m)$ is termed a (constant) preference pair on (X, Γ).

Definition 3.1.3. If $u, v \in X$ and σ is a preference pair on (X, Γ), then $u \succ_\sigma v$ if and only if there exists an element $t \in X$ such that

(i) $u, v \in \Gamma t$;

(ii) $u = \sigma_M(t)$ or $v = \sigma_m(t)$.

Definition 3.1.4. An EF f over (X, Γ) reproduces a preference pair σ (with respect to strategy $\mathcal{S}_0$) if

$$f(x) > f(y) \quad \text{whenever} \quad x \succ_\sigma y.$$

If $u \succ_\sigma v$ then u may be said to be directly preferred to v. However for constant play and representation of a preference pair by an EF it is necessary to consider indirect preferences and this motivates the following definition.

Definition 3.1.5. If $u, v \in X$ and σ is a preference pair on (X, Γ) then $u \succ_\sigma^* v$ if and only if there is a sequence $u = w_0, w_1, \ldots, w_s = v$ such that

$$w_i \succ_\sigma w_{i-1}, \quad i = 1, \ldots, s.$$

A preference pair induces an ordering on the set X. However, since not all pairs of vertices will be compared it is found that it is sufficient for the ordering to be partial as is shown by the following two theorems.

Theorem 3.1.1. Given a partial order relation Q over a finite set Y then there exists a mapping $\psi: Y \to R$ such that

(i) $\psi(x) \geq \psi(y)$ if xQy;

(ii) $\psi(x) \neq \psi(y)$ if $x \neq y$.

Theorem 3.1.2. There exists an EF f^σ over (X, Γ) which reproduces the preference pair σ if and only if $\succeq_\sigma^*$ is a partial order relation on X. ($u \succeq_\sigma^* v$ is equivalent to ($u \succ_\sigma^* v$ or $u = v$)).

Proof. The necessity is clear.
The sufficiency follows from Theorem 3.1.1. by taking Q to be the relation $\succeq_\sigma^*$ and f^σ to be the function ψ.

Corollary. If σ is a preference pair over a tree (V, T) then there always exists an EF f^σ which reproduces σ.

Proof. For a tree the relation $\succeq_\sigma^*$ reduces to $\succeq_\sigma$ and is always a partial order relation.

The above result is appropriate to the case in which the strategy $\mathcal{S}_0$ is being used. However, strategies involving lookahead are frequently used and so we turn to the problem of finding an EF f^σ so that f_d^σ † reproduces the preference pair σ.

Theorem 3.1.3. If σ is a preference pair over a tree (V, T) then there exists an EF $\hat{f}^\sigma$ over (V, T) such that $\hat{f}_d^\sigma$ reproduces σ for each $d \geq 0$.

Proof. Two functions

$$\hat{\sigma}_M : V^M \times V^m \to \{0, 1\}$$
$$\hat{\sigma}_m : V^M \times V^m \to \{0, 1\}.$$

associated with σ are defined by

$\hat{\sigma}_M(x, u) = 0$, if $u \neq \sigma_M(x)$ and $\hat{\sigma}_M(x, \sigma_M(x)) = 1$

$\hat{\sigma}_m(v, y) = 0$, if $v \neq \sigma_m(y)$ and $\hat{\sigma}_m(\sigma_m(y), y) = 1$.

$\hat{f}^\sigma$ is now defined by

$$\left.\begin{array}{lll} \text{(i)} & \hat{f}^\sigma(S) = 0; & \\ \text{(ii)} & \hat{f}^\sigma(u) = \hat{f}^\sigma(x) + \hat{\sigma}_M(x, u) & \text{if } u \in \Gamma x \subset V^m; \\ \text{(iii)} & \hat{f}^\sigma(v) = \hat{f}^\sigma(y) - \hat{\sigma}_m(v, y) & \text{if } v \in \Gamma y \subset V^M. \end{array}\right\} (F)$$

and clearly reproduces σ.
In order to prove that $\hat{f}^\sigma$ has the desired properties we first show that for all $x \in V$,

$$\left.\begin{array}{ll} \hat{f}_j^\sigma(x) = \hat{f}^\sigma(x) & \text{if } j = 2k \\ \quad = \hat{f}^\sigma(x) + (-1)^s & \text{if } j = 2k + 1 \end{array}\right\} (A)$$

where $s = 0$ if $x \in V^M$ and 1 if $x \in V^m$. An inductive proof will be given, in which H_r denotes the hypothesis that the equations (A) above are true when $k = r$.
Firstly, H_0 is true since $\hat{f}_0^\sigma = \hat{f}^\sigma$ and

$$\hat{f}_1^\sigma(x) = \max_{u \in \Gamma x} \hat{f}^\sigma(u) = \hat{f}^\sigma(x) + 1 \quad \text{if } x \in V^M$$
$$\hat{f}_1^\sigma(x) = \min_{u \in \Gamma x} \hat{f}^\sigma(u) = \hat{f}^\sigma(x) - 1 \quad \text{if } x \in V^m$$

It will now be shown that H_r implies H_{r+1}. If $x \in V^M$, then

$$\hat{f}_{2r+2}^\sigma(x) = \max_{u \in \Gamma x} \min_{w \in \Gamma u} \hat{f}_{2r}^\sigma(w) = \max_{u \in \Gamma x} \min_{w \in \Gamma u} \hat{f}^\sigma(w) \quad (\text{by } H_r)$$
$$= \max_{u \in \Gamma x} (\hat{f}^\sigma(u) - 1) = \hat{f}^\sigma(x)$$

and similar reasoning shows that $\hat{f}_{2r+2}^\sigma(x) = \hat{f}^\sigma(x)$ if $x \in V^m$. By using similar arguments it is also readily shown that $\hat{f}_{2r+3}^\sigma(x) = \hat{f}^\sigma(x) + (-1)^s$ and hence H_{r+1} follows if H_r is true. It is concluded that H_r is true for all $r \geq 0$.
$\hat{f}^\sigma$ reproduces σ and hence $\hat{f}_d^\sigma$ must also reproduce σ if d is even. If d is odd $\hat{f}_d^\sigma$ and $\hat{f}^\sigma$ differ by a constant over V^M and by a different constant over V^m and so $\hat{f}_d^\sigma$ reproduces σ since $\hat{f}^\sigma$ does. The fact that $\hat{f}_d^\sigma$ and $\hat{f}^\sigma$ differ by different constants over V^M and V^m is immaterial since an element of V^M is never compared with an element of V_m.
It has thus been established that there exists an EF $\hat{f}^\sigma$ over (V, T) such that $\hat{f}_d^\sigma$ reproduces σ for all $d \geq 0$.

For a game graph (X, Γ) the situation is more complex since the condition that $\succeq^*$ be a partial order relation is no longer sufficient and 'higher order consistency relations' are required. However, if

† For any EF g, g_d will denote the d-th derived EF. (that is the EF obtained by backing-up, using the minimax rule, over a lookahead tree of depth d).

the definition (F) given in the proof of Theorem 3.1.3 results in $\hat{p}^{\sigma}$ being defined unambiguously over X then $\hat{p}^{\sigma}_{d}$ reproduces σ over (X, Γ).

Suppose σ is a preference pair on a graph (X, Γ) and let π be a path from vertex x to vertex y and C_{π} be (the number of arcs of π corresponding to preferred moves for MAX minus the number of arcs of π corresponding to preferred moves by MIN) then a definition of consistency can be given as follows.

Definition 3.1.4. σ is fully consistent if $C_{\pi_1} = C_{\pi_2}$ for any pair of paths π_1, π_2 with the same endpoints.

Theorem 3.1.4. If σ is a fully consistent preference pair on (X, Γ) then there exists an EF $\hat{p}^{\sigma}$ over (X, Γ) such that $\hat{p}^{\sigma}_{d}$ reproduces σ for all $d \geq 0$.

Proof. Since σ is fully consistent it is easily seen that definition (F) of Theorem 3.1.3 with X replacing V leads to an unambiguous definition for $\hat{p}^{\sigma}$. The rest of the proof closely follows that of Theorem 3.1.3.

It may be noted that the condition that $X^{M} \cap X^{m} = \phi$ is essential for this proof.

3.2 Representing the decision rule of a variable player.

As discussed in §2.2 a player tends in practice to vary his play, and so attention is now concentrated on studying how variable play can be simulated by the use of PEF's. First, variable play needs to be formalized and the following definition is presented.

Definition 3.2.1. A (variable) preference function for MAX is a single-valued function $\tau_M: X^M \times \Gamma X^M \to I$ where I is the unit interval

$$I = \{x \mid 0 \leq x \leq 1\}$$

and

(i) $\tau_M(x, u) = 0$ if $u \notin \Gamma x$

(ii) $\tau_M(x, u) \geq 0$ if $u \in \Gamma x$

(iii) $\sum_{u \in \Gamma x} \tau_M(x, u) = 1$.

$\tau_M(x, u)$ can be thought of as the probability that, when in position x, MAX will choose to move to position u. A variable preference function τ_m for MIN can be defined in an analogous way.

Definition 3.2.2. If τ_M and τ_m are variable preference functions for MAX and MIN respectively then $\tau = (\tau_M, \tau_m)$ is termed a (variable) preference pair on (X, Γ).

Definition 3.2.3. A PEF p^{τ} reproduces a variable preference pair τ if

$\pi(x, u) = \tau_M(x, u)$ all $x \in X^M$ and all $u \in \Gamma x$;

$\pi(x, u) = \tau_m(x, u)$ all $x \in X^m$ and all $u \in \Gamma x$,

where $\pi(x, u)$ denotes the probability that, using p^{τ}, the move $x \to u$ will be selected.

Theorem 3.2.1. For each preference pair τ on a tree (V, T) there exists a PEF p^{τ} which reproduces τ.

Proof. Let $x \in V^M$ and $Tx = \{u_1, \ldots\ldots, u_s\}$. Then assuming the form of $p^{\tau}(u_i, \ )$, the distribution attached to u_i, has been decided upon it remains only to specify the means

$$\mu_i = \int r\, p^{\tau}(u_i, r)\, dr$$

such that

$$\pi(x, u_i) = \tau_M(x, u_i) \quad i = 1, \ldots\ldots, s. \qquad (A)$$

The μ_i will be found by using an inductive process. Let H_k be the hypothesis that μ_i, $i = 1, \ldots, s$ can be found such that

$$\pi(x, u) = \xi\, \tau_M(x, u_i) \quad i = 1, \ldots, k \qquad (B_1)$$
$$= 0 \quad i = k+1, \ldots, s \qquad (B_2)$$

B_2 is easily satisfied by setting μ_j, $j > k$ to be sufficiently small, and by using continuity arguments it is clear that H_2 is true. It remains to be shown that H_k implies H_{k+1}.

Let μ_1 be chosen so that B_2 is satisfied and let $\mu_2 = \gamma$. The combination of the distributions $p^{*}(u_1, \ )$ and $p^{*}(u_2, \ )$ is a new distribution, p^{*} say, which satisfies the conditions of Definition 2.2.1, and so by H_k, $\mu_j = \mu_j(\gamma)$, $i = 3, \ldots, k+1$ can be chosen such that

$$\left.\begin{aligned} \pi(x, u_i) &= \xi' \tau_M(x, u_i) && i = 3, \ldots, k+1 \\ &= 0 && i > k+1. \end{aligned}\right\} (C)$$

Defining

$$\varphi(\gamma) = (\pi(x, u_1) - \pi(x, u_2)) - (\tau_M(x, u_1) - \tau_M(x, u_2))$$

then, $\varphi(\gamma) > 0$ for sufficiently small γ and $\varphi(\gamma) < 0$ for sufficiently large γ. Since φ is clearly a continuous function of γ it follows that there exists a value γ^* of γ such that $\varphi(\gamma^*) = 0$, that is

$$\pi(x, u_1) - \pi(x, u_2) = \tau_M(x, u_1) - \tau_M(x, u_2) \qquad (D)$$

and since the relations (C) are maintained throughout

$$\pi(x, u_1) + \pi(x, u_2) = \tau_M(x, u_1) + \tau_M(x, u_2) \qquad (E)$$

From (D) and (E) it follows that

$$\pi(x, u_i) = \tau_M(x, u_i) \qquad i = 1, 2$$

and this together with (C) establishes the truth of H_{k+1}. It follows that H_k is true, all $k \geq 0$.

A similar result holds for positions $y \in V^m$ and so there exists a PEF p^{τ} which reproduces τ.

Theorem 3.2.2. For each variable preference pair τ over a tree (V, T) there exists a PEF $\hat{p}^{\tau}$ such that $\hat{p}^{\tau}_{d}$ † reproduces τ for all $d \geq 0$.

Proof. The proof will proceed inductively. H_m is the hypothesis that there exists a PEF $\hat{p}^{\tau}$ such that $\hat{p}^{\tau}_{d}$ reproduces τ for all $0 \leq d \leq m$, then by Theorem 3.2.1. H_0 is true and it remains to proof that H_m implies H_{m+1}.

Let $\Gamma x = \{u_1, \ldots\ldots, u_s\}$, $U_i = \Gamma^{m+1} u_i$ and p be a PEF satisfying H_m. Then, if the transformation

$$p'(y, r) = p\,(y, r - \nu_i) \qquad y \in U_i$$
$$= p\,(y, r) \qquad \text{otherwise},$$

is applied, p' also satisfies H_m. $p'_{m+1}(z, r)$ differs from $p_{m+1}(z, r)$ for each u_i, but for no other positions. The ν_i can now be chosen (in a way similar to that used in the proof of Theorem 3.2.1) so that

$$\pi'(x, u_i) = \tau_M(x, u_i) \qquad i = 1, \ldots\ldots, s \qquad (A)$$

where $\pi'(x, u_i)$ is the probability of move $x \to u_i$ being selected when using p'. A similar statement for positions $x \in V^m$ follows by analogy, and hence the result follows.

In the above two theorems the restrictions placed upon the distributions attached to the vertices of the tree were weak, and in practice it would probably be reasonable to restrict oneself to the rectangular

† $\hat{p}^{\tau}_{d}$ denotes the d-th derived PEF of $\hat{p}^{\tau}$.

distributions

$$p(x, r) = 1 \quad \text{if} \quad \mu_x - k_x \leq r \leq \mu_x + k_x$$
$$= 0 \quad \text{otherwise.}$$

For game graphs (X, Γ) which are not trees the situation is much more complex. It is not even clear as to what constitutes a sufficient condition for the existence of a PEF which will reproduce a preference pair though It seem that it is heavily dependent on the geometry of the graph.

4. Summary.

Evaluation functions and the more general probabilistic evaluation functions were introduced in $ 2. A single anpiication of these concepts, to the representation of decision rules, was then considered in detail. It was found that, to a large extent, the use of an EF or PEF is equivalent to the use of a decision rule.

Itefarenoes.

(1) Sandewall, K. J. (1969) A planning problem Solver based on Lookahead in Stochastic Game Trees, J. ACM 16, 364-K2.
(2) Shannon, C. (1 95O) prograraming a Computer for Playing Chess, ph11. Hag. 7-th Ser. 41, 256-75.
(3) Boffey, T. B. Evaluation functions - to be submitted shortly.
(4) Berge, C. (1962) The Theory of Graphs, Methuen.
(5) von Neunann, *J.* 4 Morganstern, 0. (1947) Theory of Games and Economic Behaviour, Princeton Univ. press.
(6) Boffey, T. B. (1973) Applying the Minimax Rule over Graphs which are not Trees -unpublished.
(7) Good, I. J. (1967) A Five-Year plan for Automatic Chess, Machine Intelligence 1, Edinburgh University Press.