A NEW KIND OF TURNPIKE THEOREM[1]

BY JOHN H. HOLLAND

Communicated by W. Givens, July 10, 1969

Briefly and informally an (economic or adaptive) *plan* can be described in terms of its action with respect to two sets:

$\mathcal{A}$, an effectively defined set of objects;

M, a set of functions $\mu_E: \mathcal{P} \to \mathcal{R}$, $E \in \mathcal{E}$, where $\mathcal{P}$ is the set of distributions over $\mathcal{A}$, and $\mathcal{R}$ is a ranking set of positive real numbers.

Under the intended interpretation, each $P \in \mathcal{P}$ corresponds to a mix of goods, chromosomes, strategies, or programs, $\mathcal{E}$ is a set of possible conditions or environments under which the plan is expected to operate, and the ranking of P, $\mu_E(P)$, specifies utility, expected offspring, payoff, or efficiency of the mix in condition or environment $E \in \mathcal{E}$. The plan, then, is a procedure (cf. sequential sampling procedure, dynamic programming policy) for searching $\mathcal{P}$ in an attempt to locate mixes of high rank in any given $E \in \mathcal{E}$; the object is to construct (if possible) a plan which is "robust with respect to M" in the sense that the search proceeds "efficiently" for any $\mu_E \in M$.

More formally, with any pair (τ, E) where τ is a plan and $E \in \mathcal{E}$, one can associate a trajectory through $\mathcal{P}$, $\langle \mathcal{P}(\tau, E) \rangle = \langle \mathcal{P}_1(\tau, E), \mathcal{P}_2(\tau, E), \cdots, \mathcal{P}_t(\tau, E), \cdots \rangle$; coordinated with the trajectory is the sequence of rankings $\langle \mu_E(\tau) \rangle = \langle \mu_{E,1}(\tau), \mu_{E,2}(\tau), \cdots, \mu_{E,t}(\tau), \cdots \rangle$ where $\mu_{E,t}(\tau) = \mu_E(\mathcal{P}_t(\tau, E))$. A plan τ_0 will be called *good* in E relative to a set of plans $\mathfrak{I}$ if

$$\sum_{t=1}^{\infty} [\overset{*}{\mu}_{E,t} - \mu_{E,t}(\tau_0)] = N_E(\tau_0) < \infty$$

where

$$\overset{*}{\mu}_{E,t} = \underset{\tau \in \mathfrak{I}}{\text{lub}} \{\mu_{E,t}(\tau)\}.$$

(The sense of good employed here is essentially that used, for example, by mathematical economists.) To assure that all the elements indexed by $\mathcal{E}$ represent nontrivial problems vis-à-vis the set of plans $\mathfrak{I}$ requires that, for all $t' > t$, $\overset{*}{\mu}_{E,t'} \geq \overset{*}{\mu}_{E,t} > 0$. (I.e., after a given time t,

[1] This report presents one phase of research on adaptation carried out at the Logic of Computers Group at The University of Michigan; the group is currently supported by ONR (N0014-67-A-0181-0011), ARO (DA-31-124-ARO-D-483) and NIH (GM-12236).

there will always be a plan which can give a performance of better than zero rank for a period of time.) A plan τ_0 will be called $(\mathcal{E}, \mathfrak{I})$-good if the plan is good in every $E \in \mathcal{E}$ relative to $\mathfrak{I}$. Under interpretation, an $(\mathcal{E}, \mathfrak{I})$-good plan is one which performs well in any $E \in \mathcal{E}$ relative to the set $\mathfrak{I}$ of (admissible or possible) plans under consideration. In these more rigorous terms, then, the objective is to construct (if possible) a plan which is $(\mathcal{E}, \mathfrak{I})$-good.

The purpose of this note is to demonstrate that, under rather weak conditions on $\mathfrak{I}$, it is possible to constructively define an algorithm ρ which is $(\mathcal{E}, \mathfrak{I})$-good for any collection M of bounded continuous functions μ_E. The first of the conditions on $\mathfrak{I}$ requires that, for each $E \in \mathcal{E}$, there exists a subset of good plans, $\mathfrak{I}_E$, occurring with non-negligible density (initial probability) in $\mathfrak{I}$. Loosely, this condition will be violated, when $\mathfrak{I}$ is finite, only if $\mu^*_{E,t}$ is determined by different plans at different times; i.e., it will be violated only if the sequences $\langle \mu_E(\tau) \rangle$ "determining" $\mu^*_{E,t}$ are oscillatory. When $\mathfrak{I}$ is not finite, the condition can also be violated if the sequences determining $\mu^*_{E,t}$ occur with negligible density. Since goodness is defined relative to the set $\mathfrak{I}$, the requirement is easily met if, for given E, some plan or set of plans with positive density eventually "dominates" all others in $\mathfrak{I}$; i.e., they "determine" $\mu^*_{E,t}$ for all t large enough. The second of the conditions on $\mathfrak{I}$ requires that, for any given $E \in \mathcal{E}$ and any time t, the trajectories associated with any set of probability intervals for the $A \in \mathfrak{A}$ be a measurable subset of $\mathfrak{I}$ with respect to the initial distribution P_0. That is, for any $E \in \mathcal{E}$ and t, expectations for the distributions associated with various mixtures of trajectories are defined.

The theorem establishing the goodness of ρ can be looked upon as a kind of "turnpike" theorem (an analogue of von Neumann's prototype) which holds over a very broad range of conditions. It can also be given a Bayesian interpretation (first noticed by my colleague B. P. Zeigler) treating the $\mathcal{P}_t(\tau, E)$ as hypotheses and the $\mu_{E,t}(\tau)$ as evidence; the theorem then becomes a statement about the rate of convergence of Bayes's rule.

(A different version of the theorem can be obtained by changing the premise as follows: For each $\tau \in \mathfrak{I}$ and $E \in \mathcal{E}$ let

$$\mu_E(\tau) = \lim_{T\to\infty} \sum_{t=1}^{T} \mu_{E,t}(\tau)/T$$

and define

$$\mu^*_{E,t} = \operatorname*{lub}_{\mathfrak{I}_+ \in \operatorname{pos}_\epsilon(\mathfrak{I})} \ \operatorname*{lub}_{\tau \in \mathfrak{I}_+} \mu_E(\tau)$$

where $\text{pos}_\epsilon(\mathfrak{I})$ is the set of subsets of measure exceeding ϵ in $\mathfrak{I}$. Then there is a ρ such that

$$\sum_{t=1}^{\infty} [\mu^*_{E,\epsilon} - \mu_{E,t}(\rho)] = N'_{E,\epsilon}(\rho) < \infty.)$$

The algorithm ρ, which will be called a *reproductive plan*, can be defined as follows: Let the element $\mathcal{P}$ selected at t by ρ acting in $E\in\mathcal{E}$ be given in terms of the (observed) sequences $\langle\mu_{E,t'}(\tau)\rangle_{t'=1,\cdots,t}$, $\tau\in\mathfrak{I}$, by

$$\mathcal{P}_t(\rho, E) = \int_{\mathfrak{I}} \frac{\mathcal{P}_t(\tau, E)\cdot(\prod_{t'=1}^{t}\mu_{E,t'}(\tau))dP_0}{\left[\int_{\mathfrak{I}}(\prod_{t'=1}^{t}\mu_{E,t'}(\tau'))dP_0\right]}$$

where P_0 is a probability measure (the initial distribution) on $\mathfrak{I}$. Roughly, ρ produces a mixture of distributions at each time t in which distribution $\mathcal{P}_t(\tau, E)$ occurs with density

$$\frac{(\prod_{t'=1}^{t}\mu_{E,t'}(\tau))dP_0}{\left[\int_{\mathfrak{I}}(\prod_{t'=1}^{t}\mu_{E,t'}(\tau'))dP_0\right]}.$$

The name "reproductive" comes from the observation that the given densities result if each trajectory τ is thought of as increasing its proportion in a population of trajectories by reproducing at each time t according to its "payoff" $\mu_{E,t}(\tau)$. (When $\mathfrak{I}$ is discrete, the plan ρ reduces to assigning the probability

$$\sum_{\mathfrak{I}} \mathcal{P}_t(A, \tau, E)\cdot\frac{\prod_{t'=1}^{t}\mu_{E,t'}(\tau)}{[\sum_{\mathfrak{I}}\prod_{t'=1}^{t}\mu_{E,t'}(\tau')]}$$

to each $A\in\mathcal{A}$. The reproductive plan can be defined for any case where $\mathcal{A}$ and $\mathfrak{I}$ are both effectively defined and infinite by using the appropriate random functions with $\mathcal{A}$ as the index set.) That the definition is meaningful is established by the

LEMMA. *Define $f_{E,t}\colon \mathfrak{I}\to\mathcal{P}$ by $f_{E,t}(\tau) = \mathcal{P}_t(\tau, E)$ and require that, for all $E\in\mathcal{E}$ and $t=1, 2, \cdots$, $f_{E,t}$ is P_0-measurable with respect to $\mathfrak{I}$. (I.e., let the trajectories be so defined that the sets $\{\tau\colon \mathcal{P}_t(A_\alpha, \tau, E)\leq P_\alpha$, all $A_\alpha\in\mathcal{A}\}$ are measurable subsets of $\mathfrak{I}$ for all $E\in\mathcal{E}$, $t=1, 2, \cdots$.) Then the expectation $\mathcal{P}_t(\rho, E)$ is defined and is a distribution for all $E\in\mathcal{E}$, $t=1, 2, \cdots$.*

PROOF. $\prod_{t'=1}^{t}\mu_{E,t'}(\tau)$ is integrable with respect to P_0 so that

$$\int_{\mathfrak{J}}\left(\prod_{t'=1}^{t}\mu_{E,t'}(\tau)\right)dP_0$$

is defined and hence the function $g_{E,t}\colon \mathfrak{J}\rightarrow[0, 1]$ defined by

$$g_{E,t}(\tau) = \frac{f_{E,t}(\tau)\prod_{t'=1}^{t}\mu_{E,t'}(\tau)}{\int_{\mathfrak{J}}(\prod_{t'=1}^{t}\mu_{E,t'}(\tau))dP_0}$$

is measurable and bounded (it is in fact a distribution over $\mathfrak{A}$). But then (the expectation)

$$\mathcal{P}_t(\rho, E) = \int_{\mathfrak{J}} g_{E,t}(\tau)dP_0$$

is defined. ▲

To show that ρ is $(\mathcal{E}, \mathfrak{J})$-good it must be established that the corresponding expected accumulation of losses $(\mu_{E,t}^{*}-\mu_{E,t}(\rho))$ is finite.

THEOREM. *Let M be any set of bounded continuous ranking functions μ_E indexed by $\mathcal{E}$ and such that for any $E\in\mathcal{E}$:*

(i) *there exists t for which $\mu_{E,t'}^{*}\geqq\mu_{E,t}^{*}>0$, for all $t'>t$. Let $\mathfrak{J}$ be any set of plans such that for each $E\in\mathcal{E}$:*

(i) *there exists a subset of good plans $\mathfrak{J}_E$ with $P_0(\mathfrak{J}_E)>0$,*

(ii) *the associated functions $f_{E,t}$ are P_0-measurable for $t=1, 2, \cdots$.*
Then the reproductive plan ρ is $(\mathcal{E}, \mathfrak{J})$-good.

PROOF. $\mu_{E,t}(\tau)=\mu_{E,t}(f_{E,t}(\tau))$ is a continuous function of a P_0-measurable function; hence the expected rank $\mu_{E,t}(\rho)$ of the mixture of distributions produced by ρ at time t is defined and is given by

$$\frac{\int_{\mathfrak{J}}\mu_{E,t}(\tau)(\prod_{t'=1}^{t}\mu_{E,t'}(\tau))dP_0}{\int_{\mathfrak{J}}(\prod_{t'=1}^{t}\mu_{E,t'}(\tau'))dP_0}.$$

Thus to show ρ is $(\mathcal{E}, \mathfrak{J})$-good it must be shown that

$$\lim_{T\rightarrow\infty}\sum_{t=1}^{T}\left[\frac{\int_{\mathfrak{J}}(\mu_{E,t}^{*}-\mu_{E,t}(\tau))(\prod_{t'=1}^{t}\mu_{E,t'}(\tau))dP_0}{\int_{\mathfrak{J}}(\prod_{t'=1}^{t}\mu_{E,t'}(\tau'))dP_0}\right] = N_E(\rho) < \infty.$$

Divide the numerator and denominator of the general term of the sum by $\prod_{t'=1}^{t} \overset{*}{\mu}_{E,t'}$ and apply the definition $(\overset{*}{\mu}_{E,t} - \mu_{E,t}(\tau)) = \overset{*}{\mu}_{E,t}\epsilon_{E,t}(\tau)$, obtaining the new general term:

$$L_{E,t}(\tau) = \frac{\overset{*}{\mu}_{E,t}\epsilon_{E,t}(\tau) \prod_{t'=1}^{t} (1 - \epsilon_{E,t'}(\tau))dP_0}{\int_{\mathfrak{Z}} \prod_{t'=1}^{t}(1 - \epsilon_{E,t'}(\tau'))dP_0}.$$

(When $\mathfrak{Z}$ is finite, the theorem reduces to showing that the above sum is bounded above by

$$\frac{\overset{*}{\mu}_E}{\sum_{\mathfrak{Z}} \prod_{t'=1}^{\infty} (1 - \epsilon_{E,t'}(\tau'))} \cdot \sum_{t=1}^{\infty} \left(\epsilon_{E,t}(\tau) \prod_{t'=1}^{t} (1 - \epsilon_{E,t'}(\tau))\right),$$

where $\overset{*}{\mu}_E$ is the highest possible rank in E. In this expression the first term is finite if there is a good plan for E and the second term is a convergent series (cf. Chapter IX of Knopp's *Infinite series*).) Now

$$(1 - \epsilon_{E,t'}(\tau)) = \exp[\ln(1 - \epsilon_{E,t'}(\tau))] \leqq \exp[-\epsilon_{E,t'}(\tau)].$$

Hence

$$L_{E,t}(\tau) \leqq \left[\frac{\overset{*}{\mu}_{E,t}}{\int_{\mathfrak{Z}} \prod_{t'=1}^{t}(1-\epsilon_{E,t'}(\tau'))dP_0}\right] \cdot \left[\frac{\epsilon_{E,t}(\tau)dP_0}{\exp(\sum_{t'=1}^{t} \epsilon_{E,t'}(\tau))}\right].$$

Note that $\prod_{t'=1}^{t}(1-\epsilon_{E,t'}(\tau')) \geqq \prod_{t'=1}^{\infty}(1-\epsilon_{E,t'}(\tau'))$, all t. Thus

$$\lim_{T\to\infty} \sum_{t=1}^{T} \int_{\mathfrak{Z}} L_{E,t}(\tau)dP_0 \leqq \left[\frac{\overset{*}{\mu}_E}{\int_{\mathfrak{Z}} \prod_{t'=1}^{\infty} (1 - \epsilon_{E,t'}(\tau))dP_0}\right]$$

$$\cdot\left[\lim_{T\to\infty} \sum_{t=1}^{T} \int_{\mathfrak{Z}} \frac{\epsilon_{E,t}(\tau)dP_0}{\exp(\sum_{t'=1}^{t} \epsilon_{E,t'}(\tau))}\right]$$

$$\leqq \left[\frac{\overset{*}{\mu}_E}{\int_{\mathfrak{Z}} \prod_{t'=1}^{\infty} (1 - \epsilon_{E,t'}(\tau))dP_0}\right]$$

$$\cdot\left[\lim_{T\to\infty} \int_{\mathfrak{Z}} \left(\sum_{t=1}^{T} \frac{\epsilon_{E,t}(\tau)}{\exp(\sum_{t'=1}^{t} \epsilon_{E,t'}(\tau))}\right) dP_0\right].$$

The terms $\epsilon_{E,t'}(\tau)$, $t'=1, \cdots, t$, in the sum occurring in the denominator of the second half of the product can be rearranged so that they are monotone decreasing. To show that the second half of the product is bounded above choose a monotone nonincreasing continuous $\beta_{\tau,E}(t)$ defined on $0 \leqq t \leqq T$ so that $\beta_{\tau,E}(t) = \epsilon_{E,t}(\tau)$ for $t=1, 2, \cdots, T$. Then

$$\int_{t'=1}^{t+1} \beta_{\tau,E}(t')dt' \leqq \sum_{t'=1}^{t} \epsilon_{E,t'}(\tau) \leqq \int_{t'=0}^{t} \beta_{\tau,E}(t')dt'.$$

Or, for any t and any $t-1 \leqq b \leqq t$,

$$\frac{\epsilon_{E,t}(\tau)}{\exp(\sum_{t'=1}^{t} \epsilon_{E,t'}(\tau))} \leqq \frac{\epsilon_{E,t}(\tau)}{\exp\left(\int_{t'=1}^{t+1} \beta_{\tau,E}(t')dt'\right)} \leqq \frac{\beta_{\tau,E}(b)}{\exp\left(\int_{t'=1}^{b} \beta_{\tau,E}(t')dt'\right)}.$$

Hence, noting that

$$\frac{\epsilon_{E,1}(\tau)}{\exp(\epsilon_{E,1}(\tau))} \leqq 1,$$

$$\sum_{t=1}^{T} \frac{\epsilon_{E,t}(\tau)}{\exp(\sum_{t'=1}^{t} \epsilon_{E,t'}(\tau))} \leqq 1 + \sum_{t=2}^{T} \frac{\epsilon_{E,t}(\tau)}{\exp(\sum_{t'=1}^{t} \epsilon_{E,t}(\tau))} \leqq 1 + \int_{t=1}^{T} \frac{\beta_{\tau,E}(t)dt}{\exp\left(\int_{t'=1}^{t} \beta_{\tau,E}(t')dt'\right)}.$$

Substituting $X(t)$ for $\int_{t'=1}^{t} \beta_{\tau,E}(t')dt'$ yields for the integral on the right

$$\int_{X(1)}^{X(T)} \frac{dX}{e^X} = -\frac{1}{e^X}\Bigg|_{X(1)}^{X(T)} \leqq e^{-X(1)} = 1,$$

since $X(1) = \int_{t'=1}^{1} \beta_{\tau,E}(t')dt' = 0$. But then

$$\lim_{T\to\infty} \int_3 \left(\sum_{t=1}^{T} \frac{\epsilon_{E,t}(\tau)}{\exp(\sum_{t'=1}^{t} \epsilon_{E,t'}(\tau))} \right) dP_0 \leqq \lim_{T\to\infty} \int_3 (1+1)dP_0 = 2.$$

Therefore

$$\lim_{T\to\infty} \sum_{t=1}^{T} \int_{\mathfrak{J}} L_{E,t}(\tau)dP_0 \leqq \frac{2\mu_E^*}{\int_{\mathfrak{J}} \left(\prod_{t'=1}^{\infty}(1-\epsilon_{E,t'}(\tau'))\right)dP_0}.$$

By hypothesis there exists a measurable subset $\mathfrak{J}_E$ of good plans for E (i.e., $\sum_{t'=1}^{\infty}\mu_{E,t}^*\epsilon_{E,t'}(\tau')$ converges on $\mathfrak{J}$) which implies that $\sum_{t'=1}^{\infty}\epsilon_{E,t'}(\tau)$ converges on $\mathfrak{J}_E$ because $\mu_{E,t'}^*$ is bounded below by $\mu_{E,t}^*>0$ for t' sufficiently large. Therefore $\Pi_{t'=1}^{\infty}(1-\epsilon_{E,t'}(\tau'))$ converges (to a value greater than zero) on $\mathfrak{J}_E$ and the expectation

$$\int_{\mathfrak{J}} \left(\prod_{t'=1}^{\infty} (1-\epsilon_{E,t'}(\tau')) \right) dP_0$$

is greater than zero, which proves the theorem. ▲

The theorem can be generalized in several obvious ways (e.g., the μ_E need only be Baire functions and $\mathfrak{A}$, if given an appropriate structure, can be made uncountable) but these generalizations seem to have little to do with the intended interpretations.

UNIVERSITY OF MICHIGAN, ANN ARBOR, MICHIGAN 48104