NPS55Ma71091A

United States
Naval Postgraduate School

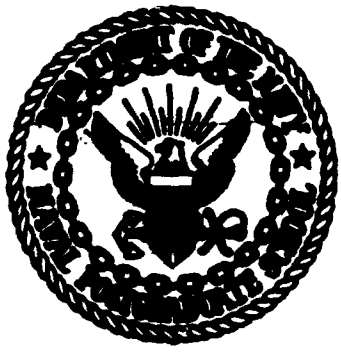

AN INTRODUCTION TO EVASION GAMES

by

Alan R. Washburn

September 1971

DISCLAIMER NOTICE

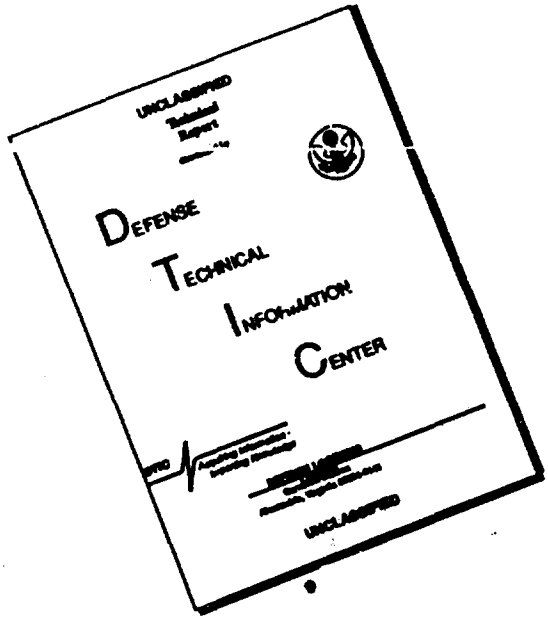

THIS DOCUMENT IS BEST QUALITY AVAILABLE. THE COPY FURNISHED TO DTIC CONTAINED A SIGNIFICANT NUMBER OF PAGES WHICH DO NOT REPRODUCE LEGIBLY.

NAVAL POSTGRADUATE SCHOOL
Monterey, California

Rear Admiral A. S. Goodfellow, USN
Superintendent

M. U. Clauser
Provost

ABSTRACT:

This paper is devoted mainly to summarizing the work that has been done to date on Evasion Games with a time lag. It also incorporates new results on the 4-step Discrete Evasion Game, and some asymptotic formulas for an Evader enforceable bound on the game in one dimension.

Prepared by:

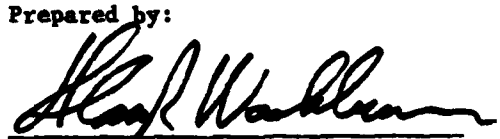

Alan R. Washburn
Department of Operations Research
and Administrative Sciences

Approved by:

J. R. Borsting, Chairman
Department of Operations Research
and Administrative Sciences

Released by:

C. E. Menneken
Dean of Research Administration

NPS55WS71091A

DOCUMENT CONTROL DATA - R & D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author)

Naval Postgraduate School
Monterey, California 93940

2a. REPORT SECURITY CLASSIFICATION

Unclassified

2b. GROUP

3. REPORT TITLE

An Introduction to Evasion Games

4. DESCRIPTIVE NOTES (Type of report and inclusive dates)

Technical Report

5. AUTHOR(S) (First name, middle initial, last name)

Alan R. Washburn

6. REPORT DATE	7a. TOTAL NO. OF PAGES	7b. NO. OF REFS
September 1971	44	11

8a. CONTRACT OR GRANT NO.	9a. ORIGINATOR'S REPORT NUMBER(S)
b. PROJECT NO.	
c.	9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)
d.	

10. DISTRIBUTION STATEMENT

Approved for public release; distribution unlimited.

11. SUPPLEMENTARY NOTES	12. SPONSORING MILITARY ACTIVITY

13. ABSTRACT

This paper is devoted mainly to summarizing the work that has been done to date on Evasion Games with a time lag. It also incorporates new results on the 4-step Discrete Evasion Game, and some asymptotic formulas for an Evader enforceable bound on the game in one dimension.

TABLE OF CONTENTS

14 KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Game						
Evasion						
Prediction						

1. Introduction.

The games that we are talking about are two person, zero-sum games that result in an Evader being either killed or not killed by an Attacker. Since the Attacker employs a weapon that requires a certain amount of time T for delivery, a major part of his task is the prediction of what the Evader will do during the next T, so that an equally good name for him would be "Predictor."

The classic example of such a game is "bomber vs. battleship." The bomber is assumed to have one bomb, infinite endurance, and perfect aim, so that the battleship's only chance of survival lies in the bomber's inability to predict the motion of the battleship during the next T. The game is trivial if the battleship knows the time at which the bomb is dropped; both the bomb and the battleship should be placed randomly within whatever area A the battleship can reach in T, and the payoff is the ratio of the lethal area to A. If the battleship does not know when the bomb is dropped, then the problem is not trivial. We shall discuss this problem further in the sequel.

A more modern counterpart of the above game is the ICBM vs. trailed ship game, where perhaps several nuclear weapons are launched in a surprise attack against some sort of naval target. Nuclear weapons are very powerful, but a modern ship can also travel some distance in half an hour, so that the outcome is not obvious. The mathematical assumptions that have to be made in this game are

probably closer to the real world than they were in the days of bombers and battleships, so that the game is of more than academic interest. It is, strictly speaking, unsolved, as is the ICBM vs. railroad train game, which is the one dimensional equivalent.

Another example of a Time Tagged Evasion Game is the anti-aircraft gun vs. airplane game, where the lethal radius and time lag are both scaled down considerably from the bomber vs. battleship. This game differs fundamentally from the latter in that the AA gun fires at many distinct points of time, with the times of firing bearing little relationship to the maneuvers of the target. This distinction results in the RMS error being the appropriate payoff and leads to the introduction of filtering techniques.

The examples above should make it clear that the time lag is what makes these games unique and interesting. It also makes them non-intuitive, at least in the sense that nature provides very little instruction in how to play them. One can presumably learn something about pursuit by watching natural predators and their prey, and there are many other combat situations where nature provides guidelines. But the ability to kill at long range is unique to modern man, and we must therefore expect intuition to be a fallible guide in what follows.

2. Models Derived from the Bomber vs. Battleship Game.

We will refer to the Bomber as A and the Battleship as E, and we make the following assumptions:

1) Both sides have infinite endurance.

2) E's motion is unrestrained except that his speed must never exceed V. In particular, E can make sharp turns if he likes.

3) A can warp the shape of the lethal region in any way he likes, so long as the area of the region is S.

It is not intuitively obvious that the above game is difficult. In fact, it is clear that A will always have E located to within a circle of radius VT, so that he can achieve a kill probability of $S/\pi(VT)^2$ by simply choosing the lethal region randomly within the circle (a wedge of random orientation will do). On the other hand, if E maneuvers in such a way that the probability density of his position is uniform within the circle of uncertainty at the moment of impact, then the kill probability will be $S/\pi(VT)^2$ no matter what lethal region A chooses. Therefore, $S/\pi(VT)^2$ is the value of the game, as long as E can behave as described. The trouble with this analysis is that E cannot make the probability density of his position increment uniform at the moment of impact. It would be easy if E knew when the bomb was dropped because his assumed ability to make sharp turns would permit him to simply pick a point at random within the circle of uncertainty and go there. But E doesn't have that critical piece of information, and therefore doesn't know when to start the maneuver. Still, intuition may

argue that E can achieve the same effect by "moving around completely at random all the time." Unfortunately, this instruction is not specific, and attempts to make it specific lead to non-uniform distributions at certain times.

To illustrate how this happens, we consider the following strategy for E in the one dimensional analog (Bomber vs. Railroad Train), where S is now a length, rather than an area: At time 0(t=0) pick a velocity uniformly from [-V,V], and stick to it for T. This will certainly make the position at T uniform throughout the interval of uncertainty. At T, repeat the procedure with an independent choice of velocity, and continue ad infinitum. Let us explore the consequences of this policy. At time T/2, the increment to E's position over the next T is $\frac{T}{2} V_1 + \frac{T}{2} V_2$, where V_1 and V_2 are the first two random velocities. Since the velocities are independent, the probability density of the sum can be found by convolving the uniform density with itself. The result is triangular, and is shown as one of the densities in Figure 1A.

If this policy for E were optimal, he could announce it to A without hurting his chances (A can figure it out for himself anyway, since he has lots of time available). In response, A might use the firing policy "straddle the present position." If S = VT, it can be seen that this policy would cut out 3/4 of the probability density function at t = T/2, so the highest kill probability is at least 3/4. Figure 1B is a plot of kill probability (P_K) vs. time. It shows that

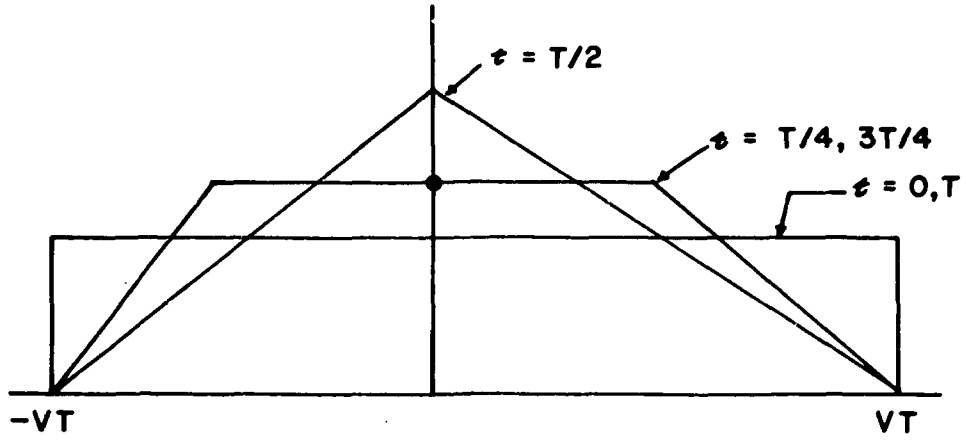

FIGURE 1A

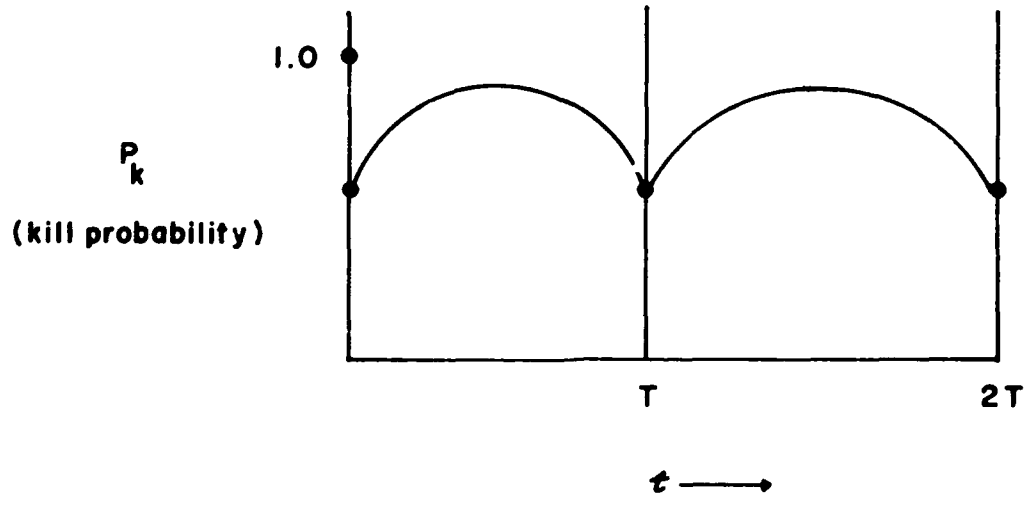

FIGURE 1B

1) The highest P_K available is 3/4.

2) $P_K = 1/2 (= S/2VT)$ is available only at multiples of T.

3) The average kill probability is closer to 3/4 than to 1/2, since the scallops are concave.

In other words, the stated strategy for E does not make the T-increment of his position "uniform all the time."

A has a better strategy than straddling that involves extrapolation. Specifically, at time $t = \delta$, will have observed the first velocity V_1, and can consequently predict exactly where E will be at T. By killing an interval of length $2V\delta$ around that point, he can guarantee a kill, and since δ is arbitrary, A can guarantee a kill as long as $S > 0$ (assuming a noiseless tracking system that is able to measure a velocity in an arbitrarily small amount of time). Thus, by paying careful attention to E's track, A can actually guarantee to kill E if E uses the stated policy.

Obviously, the trouble with E's strategy is that his motion is predictable over long periods of time; once he picks a velocity, he sticks to it for T. A natural way to improve E's strategy is to make many independent velocity decisions within each T, rather than just one. Intuition may lead to the conclusion that, if E makes enough of these incremental decisions within each interval, then the T-increment of his position (I_T) will be uniform for all t. This is false; what actually happens is that E can reliably be expected to not go anywhere. If there are N independent, identically distributed velocity choices V_i in T, then

$$\mathrm{Var}(I_T) = \sum \mathrm{Var}\left(\frac{T}{N} V_i\right) = N \mathrm{Var}\left(\frac{T}{N} V_1\right) = \frac{T^2}{N} \mathrm{Var}(V_1),$$

so that $\lim_{N\to\infty} \mathrm{Var}(I_T) = 0$. Consequently, this sort of strategy will also lead to E's being killed with probability 1 even when S is very small.

Not only have we failed to discover a strategy for E that will guarantee $p_K \leq S/(2VT)$, but we have failed to produce a strategy for E that will permit him to survive with any positive probability, even against a weak opponent. E seems to be caught in a dilemma between turning infrequently, in which case he is vulnerable to simple extrapolation, and turning frequently, in which case he has to fight the Law of Large Numbers. In addition, he is handicapped by the fact that A can wait for a maximum of p_K before he fires, since A has infinite endurance and can drop the bomb whenever he likes.

The previous paragraphs will have served their purpose if the reader is now convinced that the game is not simple, and that "moving around at random" will not guarantee that $p_K \leq S/(2VT)$ (or $p_K \leq S/\pi(VT)^2$ in two dimensions).

2.1 A Reasonable Evader Strategy.

Since A can be expected to pick the maximum of the p_K vs. t curve, a reasonable strategy for E is to behave in such a manner that p_K does not depend on t; i.e., the curve is flat. E can accomplish this by making independent direction change

choices at times corresponding to the jumps in a Poisson Process. This follows from the fact that in a Poisson Process the time until the next jump is always an exponential random variable with mean $1/\lambda$, regardless of when jumps have occurred in the past. The T-increment of E's position follows a probability law that does not depend on t, and the constant kill probability can be found by maximizing the amount of probability that A can "cut out" with S. In one dimension, E is left with the single parameter λ with which to minimize that maximum (the "turn probability" at each decision point can be taken to be 1.0, since a smaller number is equivalent to changing λ). In two dimensions, E controls both λ and the common D.F. $F(\theta)$ of the successive direction changes. In one dimension, the probability density of random variable $X = I_T/(VT)$ is [1]

$$f_X(x) = e^{-\alpha}\delta(x-1) + \alpha \frac{e^{-\alpha}}{2} \left(I_0(\alpha\sqrt{1-x^2}) + \sqrt{\frac{1+x}{1-x}}\, I_1(\alpha\sqrt{1-x^2})\right), \qquad (1)$$

where $\alpha = \lambda T$, I_K is the modified Bessel function of order K, and the positive direction for x is the last direction of travel ($-1 \le x \le 1$). The δ-function term corresponds to the probability $e^{-\alpha}$ that E will not turn at all in T. Figure 2 shows plots of the continuous part of $f_X(x)$ for 3 values of α, the optimal area for A to cut out in each case when $S = VT$, and the associated p_K (A also covers the single point $x = 1$, getting

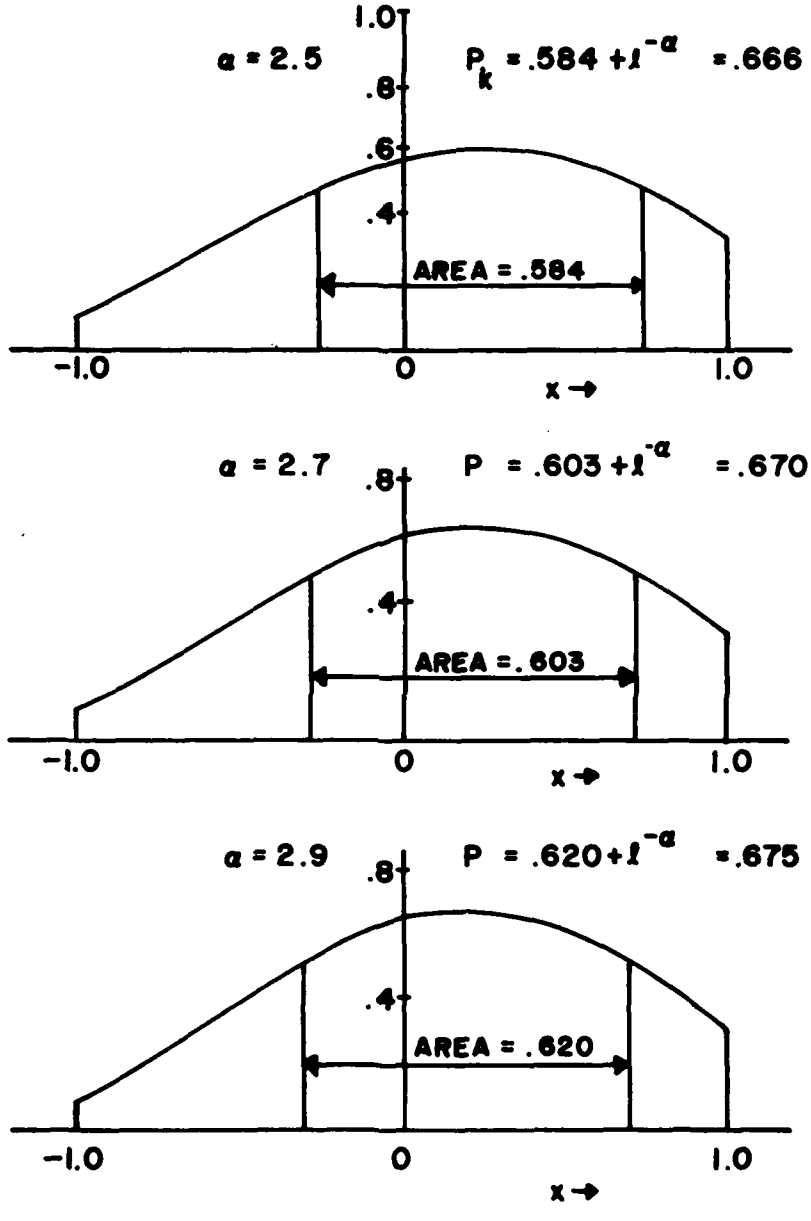
α = 2.5
$P_k = .584 + 1^{-\alpha} = .666$
1.0
.8
.6
.4
AREA = .584
-1.0
0
1.0
x →
α = 2.7
$P = .603 + 1^{-\alpha} = .670$
.8
.4
AREA = .603
-1.0
0
1.0
x →
α = 2.9
$P = .620 + 1^{-\alpha} = .675$
.8
.4
AREA = .620
-1.0
0
1.0
x →

FIGURE 2.

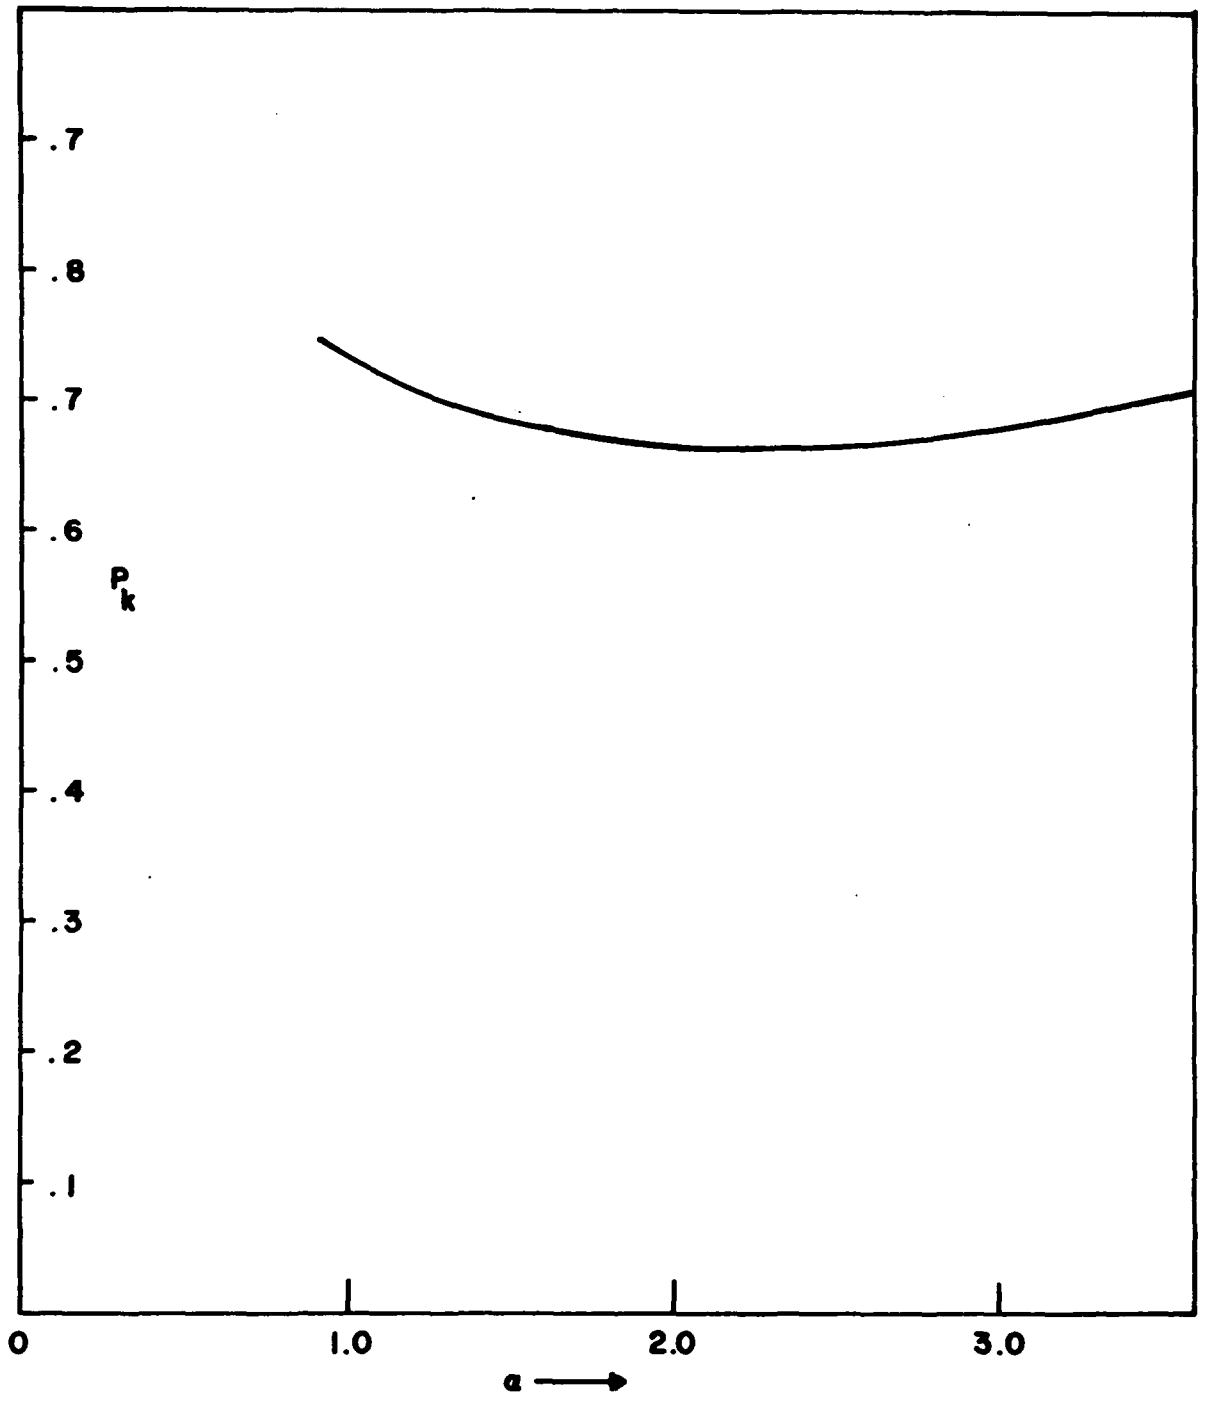

FIGURE 3.

$e^{-\alpha}$ at no cost). Figure 3 shows p_K vs. α when $S = VT$; the best α is 2.3, and the minimum p_K is about 2/3. Note that α is the expected number of turns within any period of length T; too few turns leads to vulnerability to extrapolation ($e^{-\alpha}$ is large), and too many turns leads to trouble with the Law of Large Numbers ($f_X(x)$ becomes concentrated around $x = 0$). The optimum α decreases with S; it is a curious fact that E exhibits the most frantic behavior against the weakest opponent.

2.1.1 Approximation for Weak Attacker.

The minmax p_K is a function of the ratio $S/(2VT) \equiv s$. There is no analytic representation in general, but it is possible to obtain a limiting form for small s. To do so, we first obtain the symmetrical form of (1) that results from randomizing E's first move; the resulting advantage for E will be negligible when α is large. The resulting density function is $1/2(f_X(x) + f_X(-x))$, the continuous part of which is

$$f(x) = \alpha \frac{e^{-\alpha}}{2} \left(I_0(\alpha\sqrt{1-x^2}) + \frac{1}{\sqrt{1-x^2}} I_1(\alpha\sqrt{1-x^2})\right) \qquad (2)$$

Since the power series for I_1 has no constant term, $f(x)$ is decreasing for $x \geq 0$, and the best strategy for A is therefore to straddle the origin in addition to hitting the points -1 and $+1$. An asymptotic expansion for $I_K(z)$ is

$$\frac{e^z}{\sqrt{2\pi z}}\,(1 + a_K/z + \ldots),$$

so that $f(0) \approx \sqrt{\frac{\alpha}{2\pi}}$ for large α. For the moment, we assume that $f(x)$ is well approximated by $f(0)$ within the straddle, so that the kill probability for given α, s is approximately

$$e^{-\alpha} + 2s\sqrt{\frac{\alpha}{2\pi}}\,, \tag{3}$$

where $e^{-\alpha}$ is the probability that no turn will be made.

For convenience, let $x = \frac{s}{\sqrt{2\pi}}$ and $v^2 = \alpha$. Then the minimum kill probability is $p_K = \min_{\alpha \geq 0} (e^{-\alpha} + 2x\sqrt{\alpha}) = \min_{v \geq 0} (e^{-v^2} + 2vx)$. Equating the v-derivative to 0, we find that $x = ve^{-v^2}$ and $p_K = e^{-v^2}(1 + 2v^2)$. There will be two solutions to $x = ve^{-v^2}$ if $x \leq \frac{1}{\sqrt{2}}\, e^{-1/2}$, and none otherwise. Only the larger of the two solutions can be a minimum, since the initial slope is positive. In other words, $p(x)$ could be generated parametrically for small x by using large values for v. But it is possible to obtain an explicit solution. Let $y = p_K/x$. Then $y = \frac{1 + 2v^2}{v}$, which number is never less than $\sqrt{8}$. Solving for v, we find that $v = \frac{y + \sqrt{y^2-8}}{4}$, taking the larger of the two roots. Since $x = ve^{-v^2}$, we therefore have

$$\left(\frac{y + \sqrt{y^2-8}}{4}\right)^2 - \log\left(\frac{y + \sqrt{y^2-8}}{4}\right) = -\log x. \tag{4}$$

When $s = .5$, we earlier found that the correct kill probability is .665 at $\alpha = 2.3$. For $s = .5$, (4) gives $x = .2$, $v^2 = 1.8$, $y = 3.42$, and $p_K = .685$, thus advising E to turn somewhat less frequently than he ought to, and slightly exaggerating the minimum kill probability.

When x is very small, y is very large. If we approximate v with $y/2$ and ignore the log v term in (4), the result is $y \approx 2\sqrt{-\log x}$. Recalling that $y = p_K/x$ and $x = s/\sqrt{2\pi}$, this is

$$p_K = 2\left(\frac{s}{\sqrt{2\pi}}\right)\sqrt{\log\sqrt{2\pi} - \log s}$$

$$= .796\, s\sqrt{.92 - \log s} \qquad (5)$$

E might have hoped for better, since the kill probability would be only s if he could achieve the ideal of making his position increment a uniform random variable.

When E uses the specified α, $f(0)$ is a good approximation to $f(x)$ within the straddle. To prove this, note that $f''(0)/f'(0) = c\alpha$, so that $\lim_{\alpha\to\infty} s^2 f''(0)/f'(0) = \lim_{\alpha\to\infty} c'\alpha(\sqrt{\alpha}\, e^{-\alpha})^2 = 0$.

2.2 The Problem in Two Dimensions.

In two dimensions, E controls a directional D.F. $F(\theta)$, as well as the rate of turning λ. The problem of finding the probability law for the T-increment of the Evader's position is unsolved, except for the case $F(\theta) = \theta/2\pi$ for $0 \le \theta \le 2\pi$ (the uniform D.F.).

In this case, the joint density function of $(X,Y) = (I_T^{(x)}/VT, I_T^{(y)}/VT)$ is [2]

$$f_{X,Y}(x,y) = e^{-\alpha}\delta(y)\delta(x-1) + \frac{\alpha}{2\pi}\frac{1}{1-x}\, e^{-\alpha(1-\sqrt{1-x^2-y^2})}$$

for $x^2 + y^2 \leq 1$, where the positive x direction is the last direction of travel, and $\alpha = \lambda T$ as before. By numerically integrating this function, and minimizing the integral with respect to α, Figure 4 can be computed. That figure represents the current "state of the art" as far as the Bomber vs. Battleship game is concerned; value of the game (assuming one exists) is a survival probability $(1-p_K)$ somewhere between the upper and lower bounds that are shown. It is not known whether or not the optimal strategy for E is a Poisson Strategy of the type just considered, or even whether the uniform distribution on angles is optimal within the class.

2.3 The Discrete Evasion Game.

The Bomber vs. Battleship game discussed above was the subject of some effort at the RAND Corporation in the 1950's. Finding the game to be too difficult for exact solution, Isaacs and others decided to formulate approximate games that could be solved exactly, rather than to try to approximate the value of the exact game. In Isaacs' words [3],

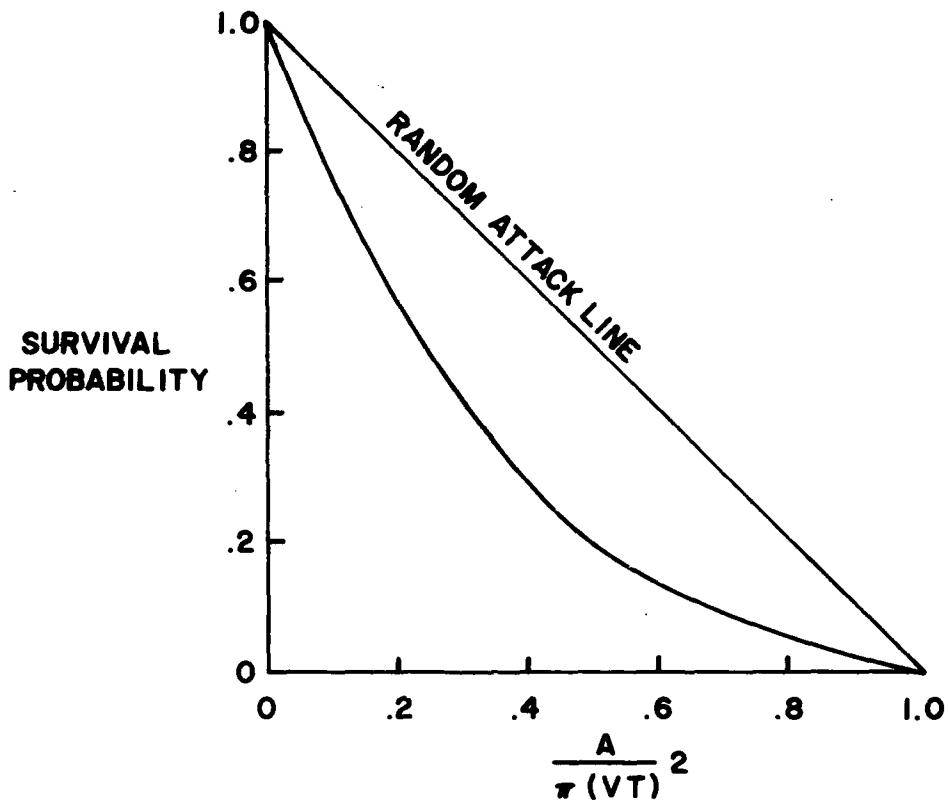

FIGURE 4.

SURVIVAL PROBABILITY FOR OPTIMUM α T.

"To gain a foothold, we simplified it further. We made the ocean one-dimensional and discrete. That is, we supposed the battleship to be located on one of a long row of points and at each unit of time he hops to one adjoining one, enjoying the sole choice of a right or left jump. The time lag was to be an integral number n of time units, or--the same thing--of jumps. This is tantamount to saying that the bomber knows all positions of the battleship which precede his present one by n jumps or more. If n = 1, the bomber knows all but the most recent of the ship's positions and there are but two possibilities for that: one space to right or left of the last observed one.

This case--n = 1--is trivial. The ship makes each decision--left or right--by the toss of a coin. The bomber can bomb at any time and when he does he also decides between the two possibilities with a coin.* Then the value of the game (hit probability) is 1/2.

Our intention was now to take up n = 2,3,4,... and, from the knowledge gained, proceed to the continuous case. Thence we hoped to restore planarity to the ocean and approach practicality by more realistic assumptions about the ship's kinematics, accuracy of the bomber, number of bombs, etc.

But the case of n = 2 proved to be an incubus. A considerable amount of effort by several people was expended before its shell began to crack. This paper will be the third one devoted to it; see [1,2]. We can expect the general class of aiming--and--evasion problems to be more difficult than anticipated, but by no means hopeless."

---

*(For the game theory tyro only.) If at some time, the ship elected, say, the probabilities: Left: .6; Right: .4, the bomber need only wait for this time and bomb on the left; then hit probability = .6. Similar considerations hold vice versa. Thus the unique optimal strategies require 50-50 decisions on the parts of both players.

The reader is referred to [3] for a complete discussion of the case $n = 2$. Briefly, the solution was found by beginning with the intuitively plausible Markov Hypothesis: In the n step game, the probability that E will go right or left on any given move may depend on the previous $n - 1$ steps, but will not depend on steps further in the past than the $n - 1^{st}$. The assumption is plausible because, "Why should E let his behavior depend on information that A already knows?" Note that the assumption holds when $n = 1$, since E flips a coin each time regardless of what he did the previous time. When $n = 2$, E will presumably maintain a constant probability of turning (x), since one can "turn" knowing only what was done on the previous step. E's direction of *travel can then be extrapolated for 2 steps* with probability $(1-x)^2$. With probability $x(1-x)$, E will be at the opposite extreme in 2 steps. This leaves a probability of $1 - (1-x)^2 - x(1-x) = x$ for the center point (no net movement). Thus, E can guarantee a kill probability no larger than

$$V = \min_{0 \le x \le 1} \max \{(1-x)^2, x, x(1-x)\}.$$

The minimum occurs when $x = (1-x)^2$, in which case $x = \frac{3 - \sqrt{5}}{2} = .382 = V$. This number turns out to be the value of the game, but the proof of the fact is not easy. Three different proofs can be found in [3], [4], and [5]. The proofs are made somewhat difficult

because the maximizing player has no optimal stragegy; he can guarantee a p_K arbitrarily close to V, but he can't guarantee V. Intuitively, A has no optimal strategy because it always pays to wait a little longer before firing, but it doesn't pay to wait "forever" before firing--the two rules are in conflict.

The results for the 2-step game have been extended somewhat. In [5], Ferguson showed that $V = \frac{3 - \sqrt{5}}{2}$ is a special case of $v_n = (n^2+2-n\sqrt{n^2+4})/2$, where $n + 1$ is the number of edges connected to each vertex in certain special graphs that he calls restricted n-graphs. For example, a lattice of hexagons is a restricted 2-graph, and the integers are a restricted 1-graph. Since a restricted n-graph may have no four-sided figures, the square lattice is not a restricted 3-graph, however.

In [5], Ferguson also mentions that the outstanding unsolved problem is the 3-step evasion game. He claims that the best E can do using the Markov Hypothesis is $(9\sqrt{3}-15)/2 = .294...$, but (in his words), "It is unknown whether or not this ... is optimal. In fact, it has been conjectured that no strategy with finite memory (that is, a strategy that depends only on the last m moves for some finite integer m) is optimal for the three-move lag problem." Thus, the Markov Hypothesis has not only not been proved (except for $n = 1,2$), but its truth is actually in doubt. Before exploring the reason for this doubt, however, we will first investigate the numbers V_n that are obtained as "game values" when the Markov Hypothesis is employed in the n-step evasion game.

According to the Markov Hypothesis E controls 2^{n-1} variables, each being the probability of going (say) to the right, given that the previous $n-1$ steps have been in a certain sequence. We can exploit the obvious right/left symmetry by considering only paths where the last step is to the (say) left, in which case there are only 2^{n-2} variables at E's disposal. Equivalently, the variables can be regarded as probabilities of turning given one of the 2^{n-2} possible previous sequences of turns. For each of these sequences, E must assure that the probability of being at any of the $n+1$ points reachable in n steps is $\leq V_n$, so that

$$V_n \leq \min_{\text{variables}} \max \left\{ \begin{matrix} (n+1)2^{n-2} \text{ functions} \\ \text{of } 2^{n-2} \text{ variables} \end{matrix} \right\}$$

In general, we can expect that about $2^{n-2}+1$ of the functions will be equal when the variables are optimum, since this establishes as many equations as unknowns, with the rest of the functions taking on smaller values. Evidently, the "controlling functions" (the ones that are equal) will become a smaller proportion of all functions as n grows large. Some idea of which functions will be controlling can be obtained by examining in detail the results for the 2, 3, and 4 step games, which are shown in Figures 5 and 6. The results for $n = 2$ are Isaacs'.*

*The first proof that $V_2 = .382...$ is actually due to Dubins [4].

The results for $n = 3$ are Ferguson's. The results for $n = 4$ are new. The probabilities shown in Figures 5 and 6 are the probabilities that E will be at the various points when using the associated optimal turn probabilities; $P(T|TN)$ is "the probability of turning if the previous move was a turn and the move before that was not." The controlling functions are the ones that are equal to the bound on the game value (underlined) in the optimal solution.

The four step game shows that the following two propositions are false:

1) When E uses his optimal (Markov) strategy, A's choice of when to fire is immaterial (A must not fire after NN),

2) In E's optimal strategy, the probability of turning depends only on the time since the last turn ($P(T|TT) \neq P(T|TN)$).

On the other hand, one striking characteristic of the 4-step game is that the optimal turn probabilities are all very nearly equal, so that the approximation resulting from the assumption that they are equal might not be a bad one. At $n = 4$, the best single turn probability is .3015..., with the resulting bound being .2380... . For large n, assuming that there is a constant turn probability at each step, the time between turns would be approximately exponentially distributed, so that formula (1) should apply. However, formula (5) gives too large a kill probability (.32 at $n = 4$) because the $e^{-\alpha}$ term is free to A in the continuous game, whereas

2 - STEP

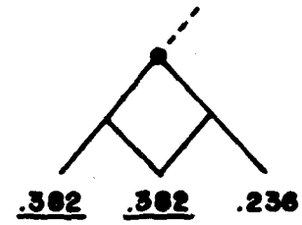

$P(T) = .382...$ $V_2 = .382...$

3-STEP

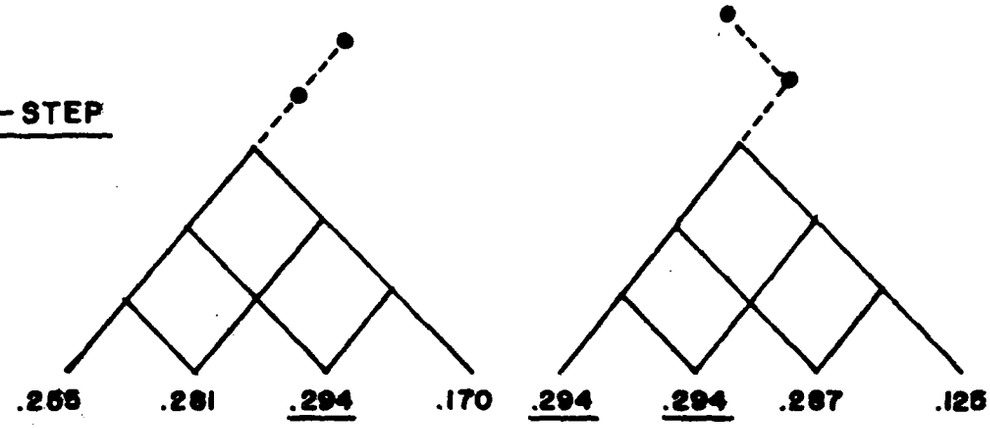

$P(T|T) = .267...$

$P(T|N) = .366...$

$.25 \leq V_3 \leq .294$

FIGURE 5.

THE TWO AND THREE STEP GAMES.

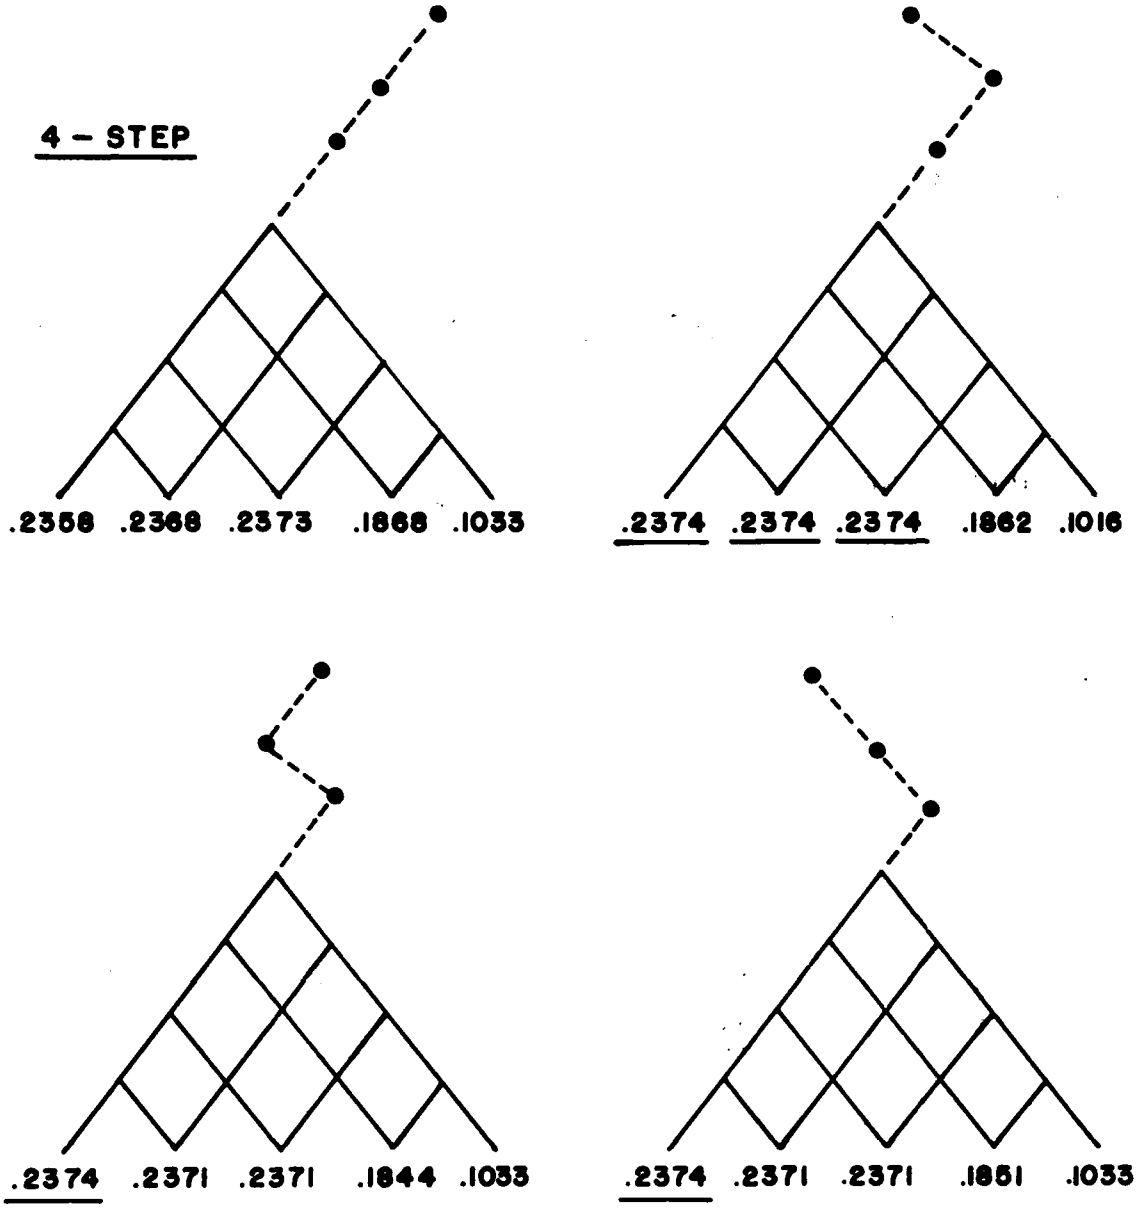

$P(T|TT) = .30318$
$P(T|NT) = .30318$
$P(T|NN) = .30318$
$P(T|TN) = .29833$

$.2 \leq V_4 \leq .23740$

FIGURE 6.

THE FOUR STEP GAME.

it is not in the discrete game. In the discrete game, (3) should be replaced by

$$p(\alpha) = \max \{e^{-\alpha}, \frac{2}{n+1} \sqrt{\frac{\alpha}{2\pi}}\} \tag{6}$$

which is minimized when the two quantities are equal. A formula for the minimum (p_K) as a function of n can be found by an analysis similar to that leading to (5). First, let $x = \frac{1}{n+1} \sqrt{2/\pi}$. Then $p_K = e^{-\alpha}$, $x = e^{-\alpha}/\sqrt{\alpha}$, $y \equiv P_K/x = \sqrt{\alpha}$, and consequently

$$y^2 + \log y = - \log x \tag{7}$$

For $n = 4$, (7) gives $p_K = .202$, which is too low. When n is very large, the $\log y$ term in (7) can be neglected, giving $p_K \approx x\sqrt{-\log x}$, or

$$p_K = \sqrt{\frac{2}{\pi}} \frac{1}{n+1} \sqrt{\log \frac{\pi}{2} + \log (n+1)}$$

$$= (.796) \frac{1}{n+1} \sqrt{.22 + \log (n+1)} \tag{8}$$

Compare (8) and (5) for $s = 1/(n+1)$.

Discrete Evasion Games with the number of simultaneous attacker shots being larger than 1 have not been investigated. E's strategy is sensitive to this number; in the 2-step game, for

instance, the game value when A has 2 shots (using the Markov Hypothesis) is

$$1 - \max_{x} \min\{(1-x)^2,\ x,\ x(1-x)\}.$$

The optimal x is .5, which is larger than the .382... that E uses when A only has 1 shot. Note that E turns more often against the stronger opponent; this is in contrast to his behaviour in the continuous version of the game if he is restricted to exponentially distributed times between turns.

It was mentioned earlier that there is some doubt that the Markov Hypothesis is true for $n \geq 3$. The next section should provide some insight into the reasons for this doubt.

3. Emission Prediction Games.

One of the things that complicates Evasion Games is the fact that there are many distinct sequences of moves that lead to each of the $n + 1$ points reachable in n steps, except for the 2 extreme points. This complication is missing in Emission Prediction Games; the Attacker (we will now call him the Predictor) wins unless the Emitter (Evader) emits a specific sequence of characters immediately after the Predictor "fires." Note that such games are trivial if the specific sequence is all the same character, since E will emit the character continually and win. E - P games were introduced by Blackwell [6], who observed that theorems of Wald and Karlin imply that slightly more general games have a value, and that E has a good strategy that is stationary in the following sense: Let A be the finite alphabet of characters, and let e_n be any sequence of n characters. Then E has an optimal strategy $P(\cdot)$ such that $P(e_n)$ is the probability that the first n characters will be e_n, and also $P(e_n) = \sum_{x \in A} P(x, e_n)$.* The last equation says that e_n is just as likely to be emitted at step 2 as it is at step 1, and the statement easily generalizes to step K.

Blackwell goes on to solve the EP10 game (the sequence is {1,0}), finding that the value of the game is $\frac{1}{4}$, and that an optimal strategy for E is therefore to choose each character by tossing a coin. Note that this strategy for E is Markov (no

*Blackwell's proof is for $A = \{0,1\}$, but the generalization is straightforward.

memory required) as well as stationary. P has only ϵ-optimal strategies.

In [7], Matula carried on the study of EP games. He found that the size of the alphabet is unimportant as long as the list of characters that P must predict will not occur is "self-disjoint" in the sense that no terminal segment is identical to any initial segment. For example, 1100 is self disjoint, but 101 is not. If there are $\lambda \geq 2$ characters in the self disjoint list, then the value of the game is $V_\lambda = \frac{1}{\lambda}\left(1 - \frac{1}{\lambda}\right)^{\lambda-1}$, which is asymptotic to $1/(\lambda e)$. It is interesting to note that the value of the game would be $1/\lambda$ if P were denied all information about E's emissions, so that the value of information can be exactly assessed in this case.

The strategies used by P and E in guaranteeing V are of considerable interest. The Predictor ignores E's emissions except that he takes note of those instants of time t_j when E has just completed a transmission of the list, with $t = 0$ being such a time. At each such moment, if he has not yet made his prediction, P picks a random number X with $P(X=i) = p_i$, $i = 0,1,\ldots$. If the next complete transmission of the list has not occurred by time $t_j + X$, then E predicts that the next λ characters after $t_j + X$ will not be the list; otherwise, P repeats the procedure. Let $K = t_{j+1} - t_j$. Then $K \geq \lambda$, since the list is self disjoint. The probability that P will make a

prediction is $p_0 + \ldots + p_{K-1}$, and the probability that P will lose, given that he predicts, is $p_{K-\lambda}/(p_0+\ldots+p_{K-1})$. Therefore, if the probabilities p_i are chosen in such a manner that $p_{K-\lambda} \leq \beta(p_0+\ldots+p_{K-1})$ for all $K \geq \lambda$, then β is a bound on the game value enforceable by P. Matula shows that such "distribution bounded densities" exist if $\beta > V_\lambda$, but not if $\beta = V_\lambda$. In other words, these densities are ϵ-optimal strategies for P.

The game EP 1100 will serve to illustrate an optimal Emitter strategy. There are four characters in the list, which we re-lable a, b, c, and d. E first emits the character "a," and then generates a Markov chain with transition probabilities $P(b|a) = P(c|b) = P(d|c) = 3/4$, $P(a|x) = P(x|d) = 1/4$ for $x = a, b, c, d$. Note that no matter what E has just emitted, he will next emit the list abcd with probability $\frac{1}{4}(\frac{3}{4})^3 = V_4$, so that the given strategy for E (call it S) is optimal. Moreover, S is obviously Markov if E's emission is regarded as a sequence of letters, and has an ultimate stationary distribution. But consider what happens when the letters are translated into 0's and 1's. The generated sequence will still be ultimately stationary, but it will no longer be Markov! To show that this is the case, let p_n be the probability that the first n symbols will be "1." Then $p_1 = 1$, since the first letter is a, and $p_2 = 1$, since the second letter is either a or b. Moreover, p_n must satisfy $p_{n+2} = (\frac{1}{4})p_n + (\frac{1}{4})(\frac{3}{4})p_n$, and the only solution of this equation for which $p_1 = p_2 = 1$ is

$$p_n = \left(\frac{1}{2} + \frac{7/2}{\sqrt{13}}\right)\left(\frac{1+\sqrt{13}}{8}\right)^{n-1} + \left(\frac{1}{2} - \frac{7/2}{\sqrt{13}}\right)\left(\frac{1-\sqrt{13}}{8}\right)^{n-1}. \qquad (9)$$

Since p_{n+1}/p_n is not a constant for $n \geq n_0$ for some n_0, S is not Markov of any order. In other words, no finite memory will do for E if he attempts to remember only 0's and 1's. Furthermore, Matula proves that all of E's optimal strategies are non-Markov in the same sense. This is in contrast to the games EP10 and EP100, where E does have an optimal Markov strategy. Ferguson had these facts in mind when he made the quote on page 14, to the effect that the Markov Hypothesis may not hold for the 3-step evasion game.

Matula also solves certain classes of EP games where the list is not self-disjoint, but *some* EP games have not been solved.

4. The Filtering Approach.

"Optimal Filtering" is concerned with the prediction of a signal when it is accompanied by noise. The source of the noise is normally thought of as impartial, but there is no reason in the theory why the noise could not be deliberately created by E in an attempt to make A's job difficult. It seems natural, then, to attempt to apply the theory, which is substantial, to Evasion Games.

Consider the Evasion Game in one dimension, and let X_t be a stochastic process representing E's position at time t. The probability law governing the stochastic process is determined by E, subject to some constraints. At time t, A will have observed E's motion up to t, and must construct a set $S(X_u,t)$ with Lebesgue measure a in the continuous problem or with n points in the discrete problem, where S does not depend on X_u for $u > t$. If A chooses to fire at time t and if $X_{t+\tau} \in S(X_u,t)$, then A wins, so that the interpretation of S is "the set of points at which A would fire if he had to fire at time t." The payoff would then be $\sup_t P(X_{t+\tau} \in S(X_u,t))$.

Let us explore some of the difficulties in applying filtering theory to the above problem. The most obvious is that A is choosing a set, rather than the single number $X^*(t+\tau)$ that is the "optimal estimate" in filtering theory. However, we can get around this problem by restricting A to firing at an "interval" with $X^*(t+\tau)$ being (say) the midpoint. This will be no real

restriction if E nonetheless behaves in such a manner that the probability law for $X_{t+\tau}$ given X_u, $u \leq t$ is unimodal.

There are other difficulties. If A fires at time t, the payoff is supposed to be $P(|X^*(t+\tau)-X(t+\tau)| \leq \frac{a}{2})$, with an analogous formula in the discrete case. However, the payoff at t in most of filtering theory is taken to be the variance $\sigma_t^2 \equiv \mathrm{Var}(X^*(t+\tau)-X(t+\tau))$. Furthermore, the usual object to be minimized in filtering theory is $E(\sigma_t^2)$, whereas we have in mind a maximizing operation, corresponding to the idea that A picks the time of firing. The expected value criterion would seem to apply better if an impartial referee were to randomly pick the time of firing, rather than A.

All of the above difficulties disappear if E is restricted to using a stationary Gaussian process. Since the distribution of $X_{t+\tau}$, given X_u for $u \leq t$, is normal, A's best strategy is always to straddle the mean. Furthermore, since the variance of that normal distribution is independent of X_u for $u \leq t$, σ_t^2 is actually a constant, so that the average and supremum operations give the same result. Since the kill probability at any time is a given, decreasing function of σ^2, the variance itself can serve as a payoff function, with E maximizing and A minimizing. In other words, if E is constrained to using stationary, Gaussian processes, then filtering theory is applicable to the Evasion Game.

Unfortunately, the rules of the Evasion Game make it impossible for E to move according to a stationary, Gaussian process. E's velocity is supposed to be bounded. We have two choices:

1) Ignore the problem, and replace the bound on the velocity with a bound on the variance of the velocity. In other words, change the game.

2) Look for E strategies for which the velocity is bounded.

We shall explore both possibilities.

4.1 Constrained Variance.

References [8] and [9] are both applicable. We will follow [8], wherein the discrete and continuous problems are both solved completely. In the discrete case, the velocity Y_n is assumed to hold from time $n-1$ to time n, so that, if X_K is the position at time K, $X_K = X_0 + \sum_{n=1}^{K} Y_n$. The stationary, normal process Y_n is assumed to be filtered noise: $Y_n = \sum_{j=0}^{\infty} \alpha_j U_{n-j}$, where the U_j are independent, normal random variables with 0 mean and unit variance. The variance of Y_n must never be larger than (say) 1, so $\mathrm{Var}(Y_n) = \sum_{j=0}^{\infty} \alpha_j^2 \leq 1$. Thus, a strategy for E is a potentially infinite dimensional vector α with a constrained norm.

We can assume that A fires at time 0, since the firing time is immaterial. If the delay is K steps, corresponding to the K step Evasion Game, then the payoff is $\mathrm{Var}(X_K - \phi(Y_0, Y_{-1}, \ldots))$, where ϕ is A's prediction of X_K based on what he knew at time 0. It is known that the optimal predictor ϕ is just the

conditional expected value [10]. This fact does not depend on the U_j being normal. However, the fact that the U_j are normal implies that the conditional expected value is a linear function of $Y_0, Y_{-1}, \ldots$, so that it is no restriction to assume that $\phi(Y_0, Y_{-1}, \ldots) = X_0 + \sum_{j=0}^{\infty} \gamma_j Y_j$. Thus, a strategy for A is a vector γ. Bram goes on to express the payoff as a function $J(\alpha,\gamma)$, subject to a mild restriction on γ, and to show that

$$\min_{\gamma} \max_{\alpha} = \max_{\alpha} \min_{\gamma} J(\alpha,\gamma) = \frac{1}{2\left(1-\cos\left(\frac{\pi}{2K+1}\right)\right)} . \qquad (10)$$

Furthermore, $\alpha_i^* = 0$ for $i \geq K$ in E's optimal strategy. At first glance, this seems to confirm the Markov Hypothesis for these games, since E needs to know only $U_{n-K-1}, \ldots, U_n$ in order to determine Y_n. However, a finite memory will suffice only if E remembers U's; no finite memory will suffice if E remembers Y's, contrary to the Markov Hypothesis. This situation is very like the one encountered in Emission-Prediction games; E's finite memory is for something that isn't observable by A. This constitutes additional evidence that the Markov Hypothesis does not hold in general for the n-step Evasion Game.

Even though A has an optimal strategy γ^*, the results for Bram's game are similar to the results for other evasion-type games in the sense that γ^* had infinitely many non-zero components. If A's strategies had been restricted to vectors with only finitely

many components, corresponding to the idea that A can wait a long time but not forever before firing, then A would have only ϵ-optimal strategies.

It was mentioned earlier that the chief reason for restricting E to normal processes is the fact that the time of firing is chosen by A, rather than at random. In some applications, the time may actually be chosen at random; for instance, in the conventional anti-aircraft problem. It is reasonable in that case to take $E(\mathrm{Var}(X_K - \phi(\bar{Y})))$ as a payoff, even though it may differ from $\inf_{\bar{y}} \mathrm{Var}((X_K - \phi(\bar{Y}))|\bar{Y} = \bar{y})$. Since Bram's arguments do not rely on the normality or commonness of the distribution of the U's, it follows that

1) The independent U's can have arbitrary distributions with unit mean and variance.

2) A can remain restricted to linear functions of the Y's.

3) The value of the game will be (10).

Bram also solves the continuous version of the above game. If T is the time delay, then the value of the game turns out to be $\frac{4T^2}{\pi^2}$. It is interesting to compare this result with a similar result of Grenander [9], who solves a similar game with $T = 1$ where A is restricted to using an "extrapolation" prediction rule:

$$X^*(t+1) = X(t) + k(X(t)-X(t-1)). \qquad (11)$$

Grenander finds that the value of the game is 2.1 if k = 1, or .82 if A sets k at .12, which is the optimizing value of k for A. Evidently, the game value is sensitive to changes in A's set of prediction strategies; Grenander finds an improvement by a factor of 2.5 if A is allowed to use values of k other than 1, and Bram has proved that another factor of 2 is possible ($4/\pi^2 \approx .4$) if A's prediction strategies are unrestricted.

As mentioned earlier, the chief objection to the constrained variance approach is that E's velocity is not necessarily bounded. Bounded velocity results are scarcer and less satisfying than constrained variance results, but still deserve mention.

4.2 Bounded Velocity.

In one dimension, the natural assumption to make is that E travels back and forth at top speed, with the times between direction changes being independent random variables with common D.F. F(x). The assumption is "natural" more out of force of habit than through any inherent logic, since there is no good reason to suppose that E's current decision should be independent of past decisions, but it nonetheless offers some hope of analysis. We earlier investigated the consequences of making F(x) exponential when the payoff was a kill probability (p. 6). In [9] Grenander finds the optimal D.F. under the assumptions that A is restricted to the extrapolation rule $X^*(t+1) = 2X(t) - X(t-1)$, and that the value of the game is to be the (average) variance of $X(t+1) - X^*(t+1)$.

He finds that the times between direction changes should be geometric: $P(T_1 = n) = (\frac{1}{2})^n$, $n = 1,2,\ldots$, which is the same as saying that E flips a coin every time unit to decide which way to go next. If the top speed is 1, the value of the game is 1. It should be emphasized that this game value is an <u>average</u> variance:

$$V = E(\sigma_t^2) = E(\mathrm{Var}(X(t+1)-X^*(t+1))).$$

For the proposed strategy for E, σ_t^2 turns out to be $\frac{1}{2} + 6(t-\frac{1}{2})^2$ for $0 \le t \le 1$, with period 1. It is true that $\int_0^1 \sigma_t^2\, dt = 1$, but $\min_t \sigma_t^2 = 1/2$, with the latter quantity being the more appropriate payoff if A picks the time of firing. In other words, Grenander implicitly assumes that a neutral agent picks the time of firing when he adopts $E(\sigma_t^2)$ for a payoff.

Grenander also formulates several games where A is not restricted to simple extrapolation. In one of them, A is free to use the conditional expected value as a predictor, and E moves in 2 dimensions at unit speed in a direction ϕ_t that is filtered, Gaussian noise. He solves the problem in case $V \equiv \mathrm{Var}(\phi_t)$ is small, finding that the game value (variance of prediction error) is $(4T^2/\pi^2)V$. Note the similarity of this result to Bram's in one dimension; the optimal filter for E is also similar.

5. Summary.

The mathematical attack on the Bomber vs. Battleship game has been proceeded by the two time honored tactics of approximately solving the actual problem and actually solving the approximate problem. Unfortunately, neither approach has worked very well; the upper and lower bounds on the actual problem are not very close to each other, and the approximate games that have been solved do not resemble the Bomber vs. Battleship game very closely. A great deal has been learned, but the really practical problems have yet to be solved.

One of the most intriguing questions that remains unanswered concerns the Markov Hypothesis for the n-step Evasion Game. Results for similar games indicate that the hypothesis as stated earlier probably does not hold, but that E nonetheless has an optimal strategy wherein he needs to remember only $n - 1$ quantities. The question is, "What quantities?, and what is the mapping from these quantities to the actions that E must take?"

There are other questions that need to be answered. What happens, for instance, if the Attacker's endurance is not large compared to the time lag, or if his observations are in some way noisy? The potential for useful research in the field is far from exhausted.

REFERENCES

[1] Washburn, A. R., "Infinite Track Approximation to the Rail-Mobile System," Coordination Sheet No. 2-1301-0-008, The Boeing Co , 26 January 1966.

[2] Washburn, A. R., "Probability Density of a Moving Particle," ORSA Vol. 17 (1969), pp. 861-871.

[3] Isaacs, R., "A Game of Aiming and Evasion: General Discussion and the Marksman's Strategies," RAND Memo RM-1385, 24 November 1954.

[4] Dubins, L. E., "A Discrete Evasion Game," Inst. Air Weapons Research Tech. Note 2.

[5] Ferguson, T. S., "On Discrete Evasion Games with a Two-Move Information Tag," Proc. of the Fifth Berkeley Symposium on Math. Stat. and Prob., Vol. 1, University of California Press, 1967.

[6] Blackwell, D., "The Prediction of Sequences," RAND Memo RM-1570, 12 October 1955.

[7] Matula, D., "Games of Sequence Prediction," O.R. Center, U. of California, Berkeley, ORC 66-3, February 1966.

[8] Bram, J., "Minimax Prediction and an Evasion Game," OEG Memo IRM-39, 1963.

[9] Grenander, U., "A Tactical Study of Evasive Maneuvers," FOA Reports, Vol. 2, No. 4, Research Inst. of National Defense, Stockholm 80, Sweden, 1968.

[10] Doob, J. L. Stochastic Processes, Wiley, New York, 1953.

[11] Karlin, S., "Operator Treatment of the Minimax Principle," Cont. to the Theory of Games, Ann. Math., 24 (1950).