The Mathematics of Gambling

The Black and White Classic Invitational Backgammon Tournament

by Edward O. Thorp

Roulette will be continued next month so I can tell you about the Black and White Classic Invitational Backgammon Tournament in Pebble Beach. I was invited as one of several journalists, through the kind auspices of Sid Jackson, President of the American Backgammon Players' Association. Participation was by invitation only and the level of player was consequently very high.

The championship division included 100 players and in the intermediate division, 32. The entrance fees were $150 and $175 respectively, with total prize money of $17,500. Entrance fees were for the benefit of the Monterey Peninsula Boy's Club, and prize money was provided by Black and White. The eventual championship flight winner got $5,000 and also received a $1,000 backgammon table, and there also was a Gold and Diamond brooch ladies prize. In its sixth year, the tournament is hosted annually by the Pebble Beach Corporation.

This is the first backgammon tournament I've attended and I wanted to see what they were like. What impressed me the most weren't the details of play or who won, rather the fascinating people that were there and the ideas we exchanged. So that's what I'll write about.

Paul Magriel and I met for the first time and talked at length. I was surprised to learn that he has had a long time friendship with Mr. X and Mr. Y of Beat the Dealer (the men who financed my first test of my blackjack system in Nevada). I asked Paul what he thought the odds were that X would leave a blot. We are assuming that O has enough other men (not shown) to move so that he can hold the X one-point as long as necessary. Paul gave the odds as about 8:1 in favor of X leaving a blot. He also said that it was fairly accurately established that the odds were about 12½:1 in favor of an O win in Figure 2. Here O has one man on the bar and all other men off. But with two men on the bar, as in Figure 3, the odds shift to 2:1 in favor of X. Readers who have the patience can give these statements a rough test by playing out a hundred or so of each situation and keeping a tally of what happens.

Before the tournament, Sid Jackson had asked me what the odds were in situations like that of Figure 1, so I had already thought about how to solve it.

The way we have allocated the fifteen X men on X's five free home boards is only one of 19!/(14! 5!)=19 X 18 X 17 X 16 X 15/1 X 2 X 3 X 4 X 5=11,628 different ways to do this. As Sid pointed out, the answer to the question will in general vary for the 11,628 choices. Also, X will often have a choice of how to move his men, and this will affect the answer. The problem is so huge that modern high speed computers cannot solve it by a direct game tree analysis. However, I have invented a short cut procedure for solving all end game pure racing or no-contact positions, which I call "the fast recursion method." Even though this is not a pure racing position since men can still possibly be hit, it turns out that the fast recursion method works and yields a solution which can readily be found with currently available computers.

Figure 4 illustrates an easier problem along these lines, which I will work out for the ideas and techniques it shows.

Problem: Suppose O holds the X one-point and X had n men on the X two-point. Also X has borne off all his other men and O has enough men so far from home that he can hold the X one-point until the game ends. Suppose play begins with O's turn. What is the probability that X will leave a shot before he finishes bearing off?

In n=0, then all the X men are already off and the game is over so the probability O gets a shot is zero. In n=1, then X has a blot on the two-point, and the probability O has a shot is 1.00. This gives the first two entries in Table 1. Suppose n=2. Then O has no shot at first. If X rolls 1—1 on his next turn, he cannot move. If X rolls any 1-A with A greater than 1, he bears off only one man and O has a shot. This happens in 10 ways out of 36. If X makes any other roll, namely no ones, both men come off and O has no shot. This happens in 25 ways out of 36. Thus there are 35 (equally likely) ways (out of 36) that the situation is changed. Ten give a shot, 25 give no shot. Therefore, the probability is 10/35=2/7 that O will have a shot when n=2. The part of Figure 5 labelled "Case n=2" illustrates this schematically.

Now consider the case n=3. If X rolls double twos, threes, fours, fives or sixes (5 ways), he is off, i.e. we advance to the n=0 case. If X rolls 1 — 1, we stay in the n=3 case (1 way). If X rolls 1 — A with A greater than 1, he bears off exactly one man and we advance to the n=2 case (10 ways). For all the other rolls, which are just those rolls where X rolls neither ones nor doubles, we advance to the n=1 case. There are 20 ways. Thus 5 ways (n=0) are safe, 20 ways (n=1) leave a shot and 10 ways (n=2) lead to a 2/7 chance that there will be a shot later. Hence the probability of X ever leaving a shot starting with n=3 is [20 + 10 X 2/7]/35=160/(7 X 35)=32/49.

Mathematical readers: The analysis for n=3 applies to the general case and gives the recursion relation $P_{n+3}=(10/35)P_{n+2} + (20/35)P_{n+1} + (5/35)P_n$ or $P_{n+3}=2P_{n+2}/7 + 4P_{n+1}/7 + P_n/7$. Thus, $P_4=64/343 + 8/49 + 1/7=169/343$. We could derive a general forumula for P_n. Table 1 is a computer calculation of the first 20 values. The values to five decimal places are 0.53846 thereafter. Proposition bet: for n=4, bet that X does not leave a shot. For n=5, bet that X does leave a shot.

It is fairly easy to use the same method to analyze the case where all X men are on the two and three points. The reason is that there is only one choice for how to play any roll. But when some men are on the 4, 5, or 6 point there may be more than one choice for playing a roll. The added complexity that results requires a more powerful method ("fast recursion") than the simple one I've given.

Next, we discussed a straight race, where (for simplicity) all pips for each roll count towards moving men off. Then there is no wastage of pips and no loss of part or all of a roll due to being blocked. Assume X has a pip count of n and he is on roll. For what O pip count should: X double? O accept? X redouble? O ac-

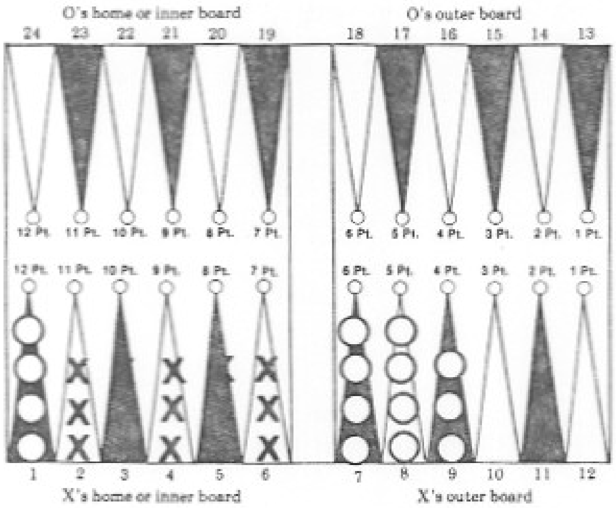

Figure 1

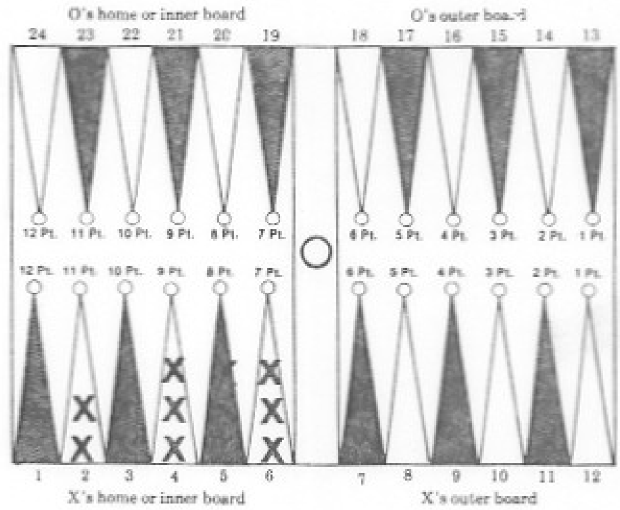

Figure 2

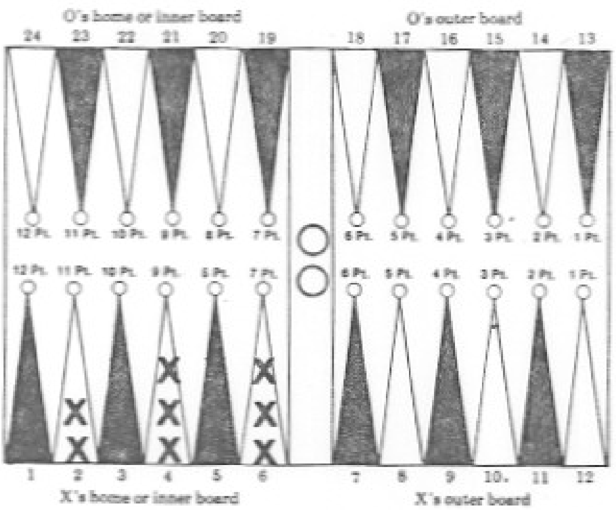

Figure 3

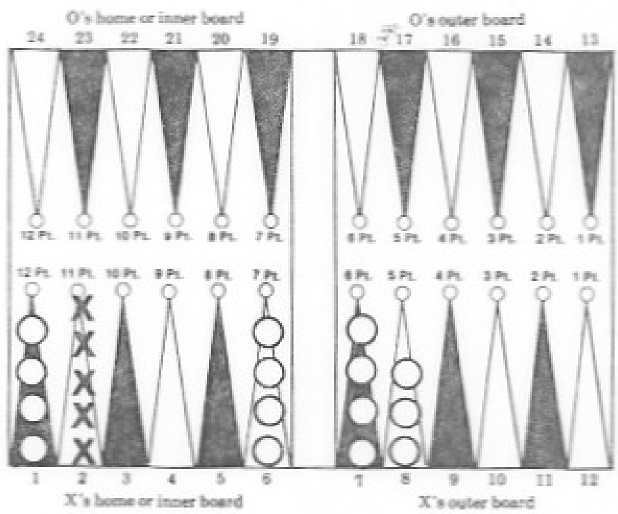

Figure 4

cept? I solved this exactly by fast recursion. Approximate solutions have been given by several authors. Schematically we might represent the situation by curves, as in Figure 6. Paul expressed the belief that there was an upper limit (or asymptote) both to d(n), the number of pips O needs to be behind, for X under best play to double, and f(n), the number of pips O needs to be behind in order for O with best play to fold.

I believe that d(n) and f(n) behave like n times a constant. Therefore, they have no upper limit as n increases. Here is a sketch of my argument: as n increases, this backgammon race is better and better approximated by the "continuous game" discussed by Keeler and Spencer in "Optimal Doubling in Backgammon." There they show that for this continuous approximation, X should double or redouble, and O should fold when the probability of X winning the race reaches 0.8.

Now each roll of the dice has a mean number of pips equal to 49/6 and a variance of 19. For a very large number of rolls, the total number n of pips rolled by X has approximately a normal distribution with mean 49R/6 and variance 19R. Then if X has a pip count of n it takes X about n ÷ (49/6)=R rolls to get off, so the game will last about this many rolls. The difference in the number of points rolled by X and O will have a variance after R rolls of 2 x 19=38R or about 38n ÷ 49/6=228n/49. This is a standard deviation of 228n/49=2.1571 n. Now if X starts ahead by a certain number of pips, he'll tend to be that far ahead at the end. But the "spread" of his lead will have a standard deviation of about 2.5171 n. For X to be ahead at the end with probability 0.80, he needs to start with a lead of about 0.842 standard deviations (from table of normal probability distribution), or 1.82 n. However, the player who has the turn has an effective lead of about 49/6 ÷ 2=4.08 so the actual pip count lead X needs is about d(n)=f(n)=1.82 n − 4.08.

For those who doubt this argument, I admit that (a) all the details are not absolutely "tight," and (b), it only applies to "large" n and I haven't said how "large" n needs to be to make this formula good. One might expect it to have to be hundreds or even millions. Now, prepare to be amazed. Table 2 compares the exact computed values for various n with the values from my formula. They are remarkably close even for n=30. The values used for P, the probability the first player wins, in Table 2 have been obtained from my four place table of P values for all pip counts up to 168 for either player. The similar table published by Zadeh is too abbreviated, and the similar table published by Keeler and Spencer is both too approximate and too abbreviated. The last two columns show that 1.82 n − 3.83 is a better fit in the 40 to 150 pip range but appears to be losing ground to 1.82 n − 4.08 as n increases beyond 130 pips.

Table 1

N	P(N)
0	0
1	1
2	0.28571
3	0.65306
4	0.49271
5	0.55476
6	0.53334
7	0.53978
8	0.53824
9	0.53842
10	0.53851
11	0.53842
12	0.53847
13	0.53845
14	0.53846
15 or more	0.53846

Table 2

X pip count n	$1.82\sqrt{n}$ −4.08	pips 0 behind for P=0.8	$1.82\sqrt{n}$ −3.83	Error
10	1.68	0.41		
20	4.06	3.91		
30	5.89	5.92		
40	7.43	7.61	7.68	+.07
50	8.79	9.00	9.04	+.04
60	10.02	10.25	10.27	+.02
70	11.15	11.40	11.40	0
80	12.20	12.45	12.45	0
90	13.19	13.44	13.44	0
100	14.12	14.38	14.37	−.01
110	15.01	15.26	15.26	0
120	15.86	16.11	16.11	0
130	16.67	16.90	16.92	+.02
140	17.45	17.68	17.70	+.02
150	18.21	18.42	18.46	+.04

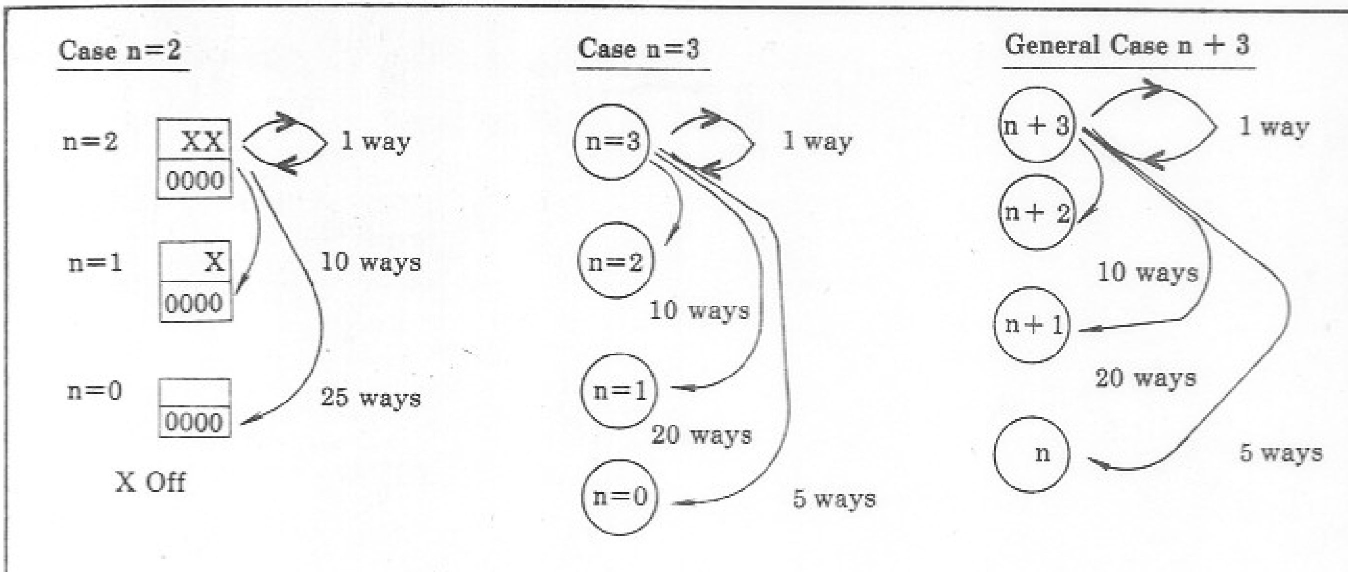

Figure 5

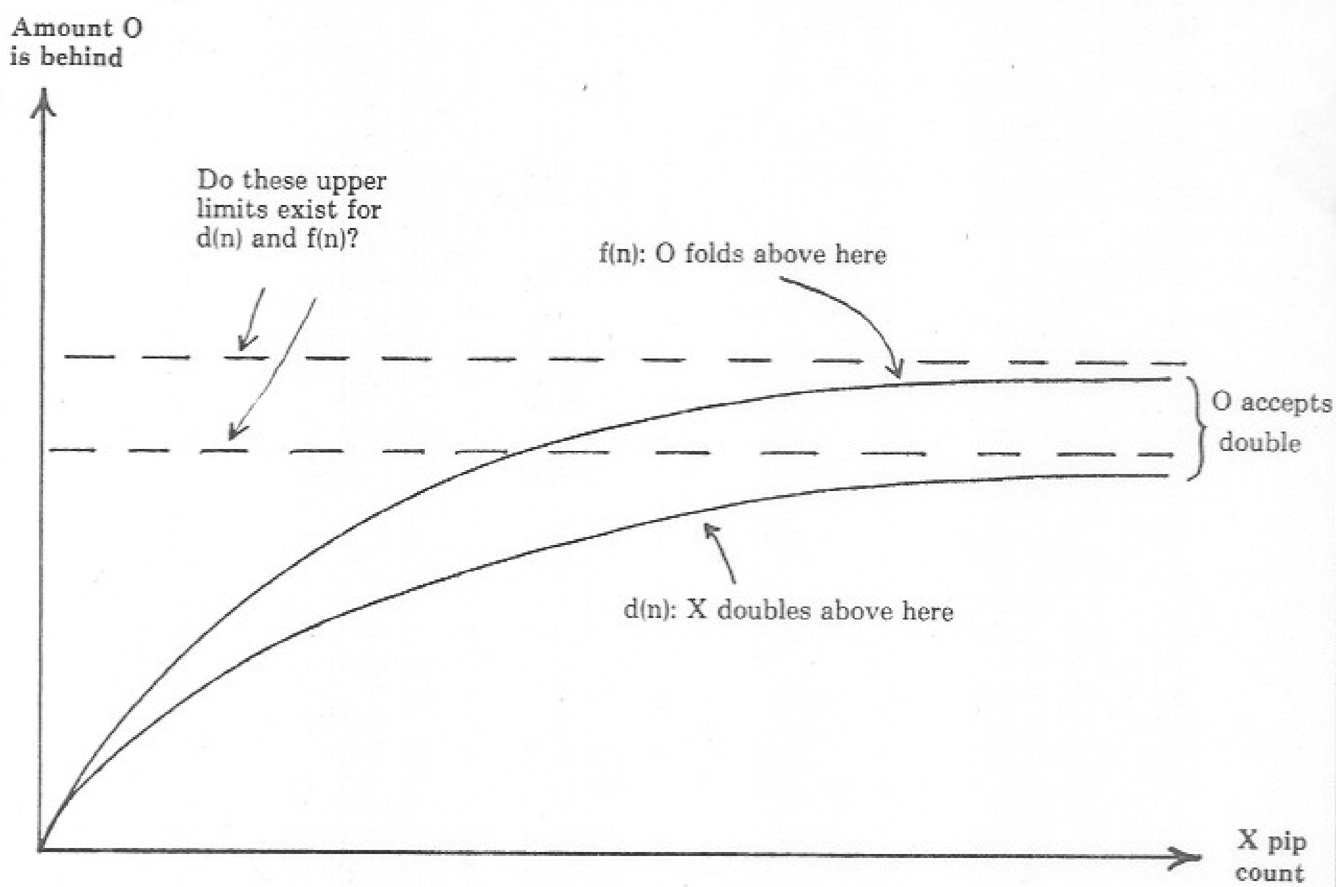

Figure 6