



Optimal Doubling in Backgammon

E Emmett B. Keeler, Joel Spencer

Operations Research, Vol. 23, No. 6, (Nov. - Dec., 1975), pp. 1063-1071.

Stable URL:

[http://links.jstor.org/sici?sici=0013-788X\(197511\)23:6:1-L;1-0;FT1-2;957333A633C1063%3AADIB33F5-0;CXV33Z-Z](http://links.jstor.org/sici?sici=0013-788X(197511)23:6:1-L;1-0;FT1-2;957333A633C1063%3AADIB33F5-0;CXV33Z-Z)

Operations Research is currently published by INFORMS.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/informs.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Optimal Doubling in Backgammon

EMMETT B. KEELER

The Rand Corporation, Santa Monica, California

JOEL SPENCER

Massachusetts Institute of Technology, Cambridge, Massachusetts

(Received original July 5, 1974; final, April 9, 1975)

This paper shows that the optimal betting strategy for a continuous model of backgammon is to double when you have an 80 percent chance of winning. We discuss the differences with the published literature on the real game and the problem of infinite expectations. The optimal strategy for a simulation of the end game is computed by dynamic programming.

BACKGAMMON is a perfect information game, with both players' moves governed by alternating throws of dice. Thus, in theory, there is a best strategy for each player to maximize his chance of winning, and each player can compute his exact probability of winning at any time if the best strategies are followed. Much of the interest of backgammon lies in the betting. Before any of his turns, a player can *double* the stakes, whereupon his opponent can *fold* and pay the doubler one unit, or he can play on. If he plays on, only he has the right to *redouble*; that is, the players must alternate doubling. At the n th double, the nondoubler can fold and pay the doubler 2^{n-1} units or he may play on with the stakes at 2^n (and the right to make the next double).

In the first part of this paper, we describe the optimal betting strategy for a continuous model of the game. Later, we discuss how this strategy is modified in the real game. (Although it may upset backgammon aficionados, gammons and backgammons will be considered only as simple wins.)

1. CONTINUOUS GAMES

Consider a two-person game that always ends in some finite time. Let $p(t)$ be the probability at time t that player A will win (given that both players play optimally). If $p(t)$ is a continuous function of time, we call

the game continuous. In a noncontinuous game, single moves may cause $p(t)$ to jump (and sometimes drastically, as in Monopoly when Boardwalk with a hotel is hit, and backgammon, especially in the end game).

With any continuous game, we can associate a game (called the P -game) consisting of the unit interval $[0, 1]$ and an indicator. At time t , this indicator is at point $p(t)$. (Call that point the state of the P -game.) Player A wins if the indicator reaches 1 before it reaches 0. We now show that the P -game indicator for any continuous game moves as a continuous symmetric random walk on $[0, 1]$ although the time scale may be different.

LEMMA 1. Let the P -game be at state x , with $a, b > 0$ such that $0 \leq x - a < x < x + b \leq 1$. Let E be the event that $p(t)$ becomes $x + b$ before it becomes $x - a$. Then the probability of E is $a/(a + b)$.

Proof. The state $x = \text{prob}[A \text{ will win}] = P(E)P[A \text{ will win}|E] + P(\sim E)p[A \text{ will win}|\sim E]$. Since $p(t)$ is Markovian, however, $p[A \text{ will win}|p(t) \text{ becomes } x + b \text{ before } x - a] = P[A \text{ will win}|p(0) = x + b] = x + b$. Similarly, $P[A \text{ will win}|\sim E] = x - a$. (The event $\sim E$ is the event ' $p(t)$ becomes $x - a$ before it becomes $x + b$.' This is where continuity and a finite end come in.) Thus,

$$x = P(E)(x + b) + [1 - P(E)](x - a) \quad (1)$$

and therefore

$$P(E) = a/(a + b). \quad (2)$$

Example. Basketball is closer to being continuous than most major sports. If we decide a priori that UCLA has a 90-percent chance of beating MIT, how likely is it that at some point in the game, MIT will have a 50-percent chance of winning? By the lemma, the probability that $P = 50$ before $P = 100$ is $(100 - 90)/(100 - 90) + (90 - 50) = 1/2$.

2. BACKGAMMON BETTING ON CONTINUOUS GAMES

Now suppose in a continuous game that only A has the right to double and his probability of winning is $p(t) = \alpha$. For which α should A double? Also, for which β should B fold if A doubles? There may be α at which A has equal expectations whether or not he doubles, and β where B has equal expectations whether he accepts or folds. Assume A doubles at such points α and B folds at such points β . The expected gain to player A is a monotone function of $p(t)$. Thus, by continuity considerations, there will be an α_0 such that A should double if $\alpha \geq \alpha_0$ and should not double if $\alpha < \alpha_0$. We shall call α_0 A 's *doubling point*. Similarly, B has a *folding point* β_0 , which is the lowest $p(t)$ at which he folds.

THEOREM 1. If both sides play perfectly, $\alpha_0 = \beta_0$.

Proof. Since B is assumed to play perfectly, A 's expectation is \leq the current value of the game. Hence if B would fold, A should double.

Therefore A should double at β_0 since B will fold (or do equally poorly if he accepts). So $\alpha_0 \leq \beta_0$. Next, suppose that $\alpha_0 < \gamma < \beta_0$. Let's compare A 's strategies of doubling at α_0 with doubling at γ . If A doubles at α_0 , either the game will eventually reach state γ or it will not. If it does, A could have waited, since B would have accepted anyway. If the game never reaches state γ , A loses, and loses more by his double at α_0 . Therefore to double at $\alpha_0 < \beta_0$ is premature; thus $\alpha_0 \geq \beta_0$ so $\alpha_0 = \beta_0$.

We can use Theorem 1 to solve for α_0 . By symmetry, an optimally playing B will double at state $1 - \alpha_0$. When A doubles at α_0 , B must have equal expectations whether he folds or accepts. Let K be the current stakes of the game. Then if B folds, he gets $-K$. Assume he accepts. Lemma 1 shows that A will win before B doubles with probability $(2\alpha - 1)/\alpha$. With probability $(1 - \alpha)/\alpha$, B doubles and, since it doesn't affect the expectations, we can assume A resigns. Thus,

$$-K = 2K[(1 - \alpha)/\alpha - (2\alpha - 1)/\alpha], \quad (3)$$

$$-\alpha = 2[2 - 3\alpha], \quad (4)$$

$$\alpha = 0.8 = 80 \text{ percent.} \quad (5)$$

Now suppose either player has the right to double. When player A doubles, the situation for B is no different from when only A had the right. So B 's folding point is still 80 percent. Again, the argument of Theorem 1 shows that A should not double until $\alpha = 80$ percent.

This result completes the strategy. A player doubles if he has at least an 80-percent chance of winning and accepts a double if he has at least a 20-percent chance of winning. If both sides play optimally, we can compute $V_A(p)$, the expected gain of player A who is at point p with the right to double. Since A wins when $p = 0.8$ and loses when $p = 0$, by Lemma 1 $V_A(p) = [(p - 0) - (0.8 - p)]/0.8 = (5p/2) - 1$ for $0 \leq p \leq 1/2$. Similarly, $V_B(p)$, the value to A of being at probability p when B has the right to double, is $(5p/2) - 3/2$ for $1/2 \leq p \leq 1$; and $V_C(p)$, the value to A of being at p when both players can double, is $5(2p - 1)/3$ for $1/2 \leq p \leq 3/4$. Thus, for example, $V_A(1/2) = 1/4$ is the value of having the exclusive right to double at the start of the game.

(In our paper⁽⁸⁾ we use similar methods to compute optimal betting strategies for a variety of other betting rules, including a fixed number of nonalternating raises and continuous raises.)

The optimal strategy is unique, as it has a positive expectation against any other strategy. (However, see the Appendix for problems that may arise when the expectations do not exist.) Suppose x is the largest probability at which B doubles and y is the smallest probability at which B folds. The most interesting of the four cases to consider is when $x > 0.2$ and $y > 0.8$ (A and B both accept the double). By Lemma 1,

$$V_B(0.8) = (0.8 - x)/(1 - x) + (1 - 0.8)2V_A(x)/(1 - x). \quad (6)$$

The 2 in the second term on the right comes from B 's doubling at x and A 's accepting. Again by Lemma 1,

$$V_A(x) = 2xV_B(0.8)/0.8 - (0.8 - x)/0.8, \quad (7)$$

and

$$V_C(0.5) = 2[(0.5 - x)V_B(0.8)/(0.8 - x) + 0.3V_A(x)/(0.8 - x)]. \quad (8)$$

Combining (6), (7), and (8), we have $V_C(0.5) = (5x - 1)/4(1 - 2x)$, which is greater than zero for $0.2 < x < 0.5$. The other cases also show positive expectation when $x \neq 0.2, y \neq 0.8$.

While fold at 0.2, double at 0.8 has a positive expectation against any other strategy, it does not take full advantage of the opportunities an inferior opponent provides. Suppose B doubles at $x > 0.2$. A should definitely accept the double. If β_0 is B 's folding point, A can double at β_0 if he wants B to fold, or double at $\beta_0 - \epsilon$ if he wants B to accept. The value of $\beta_0, \beta_0 = y$, that makes A indifferent has

$$V_B(y) = \frac{1}{2}. \quad (9)$$

By Lemma 1,

$$V_B(y) = \{(y - x) + (1 - y)2V_A(x)\}/(1 - x), \quad (10)$$

and

$$V_A(x) = \{2V_B(y) - (y - x)\}/(y - 0). \quad (11)$$

Combining (9), (10), and (11), we obtain $6y^2 - 9xy - 5y + 8x = 0$. Hence $y = (9x + 5 \pm \{(9x + 5)^2 - 192x\}^{1/2})/12$. When $0.2 \leq x < \frac{1}{2}$, there are two roots. If $\beta_0 >$ the upper root, A doubles at $\beta_0 - \epsilon$ for an expected gain greater than 1. If $\beta_0 <$ upper root, A doubles at β_0 and B folds.

For $x > \frac{1}{2}$ one must be careful about the 'theoretical' problems described in the appendix. Let us consider A 's position if he sets $y = x + \frac{1}{3}$ as his doubling point. First assume B accepts this double. Then, by the notation of the appendix, $M = \frac{1}{3}, L = x, R = \frac{2}{3} - x$ so $3LR < M$. Thus the expectations are defined. Solving (10) and (11), we obtain

$$\begin{aligned} V_A(x) &= (x - y)(3x - 1)/\{y(3x - 1) - 2x(3y - 2)\}, & \text{(in general),} \\ &= 1/(3x - 1), & \text{(for } y = x + \frac{1}{3}\text{),} \\ &> 0 & \text{(as } x > \frac{1}{3}\text{).} \end{aligned} \quad (12)$$

If, on the other hand, B resigns if doubled at y , the value is

$$V_A(x) = [x - (y - x)]/y = (x - \frac{1}{3})/y > 0.$$

In either case A has positive expectation immediately after B 's double. Under such circumstances A can alter his play so as to make his expectation arbitrarily large. Suppose A 's expectation after B 's double is $\delta > 0$. When

B doubles, A should immediately redouble. Then B would redouble right back. Suppose A exchanges $2n$ doubles with B and then—after B 's n th double—switches to the doubling point $y = x + \frac{1}{3}$. The stakes have been increased by a factor of 2^n , and A 's expectation is now $2^n \delta$. By making the appropriate (large) choice of n , A may make his expectation as large as desired. It is a dangerous strategy, for if B suddenly 'wises up' (i.e., reads this paper) when the stakes are 2^{1000} , A will be sorry.

Of course, if B 's doubling point is $x > \frac{1}{2}$, the analysis is simpler and more intuitively clear. A merely exchanges $2n$ doubles with B and then 'throws away' the doubling die.

3. DOUBLING IN BACKGAMMON

If their books can be believed, backgammon experts do not wait for a 0.8 chance of winning to double. For example, Jacoby and Crawford^[1] suggest a redouble at '9 to 5,' which is about 0.65, and Goren^[2] gives "2 to 1" or 0.67 as the proper odds. All suggest that there is considerable difference between the first double and later redoubles and recommend accepting doubles only when $p > 0.25$. Is this bad advice, or is our continuous game a poor model for the real game of backgammon?

One possibility is that the advice is practical—if the great mass of players is doubling at around 0.6, a player does better by doubling at 0.7 (rather than 0.8) and keeping such opponents in the game than he does by playing minimax. As we have seen, the optimal strategy is to double just below the novice's folding point, but perhaps because of uncertainty about the novice's evaluation of the game and about his folding point, 0.67 may be safer. (A more cynical view is that experts can make more money by playing their readers than by selling books.)

More important is the fact that backgammon is not really continuous, especially near the end of the game. A player may then know that his next throw is decisive—that after the throw the outcome will be certain. Since the game effectively ends on that throw, a player should double whenever $p > 0.5$. At no time should a player decline a double if $p \geq 0.25$, since $2[p - (1 - p)] \geq -1$ at that point. The farther from the end of the game, the less the probability of winning changes with each move, and the closer to each other the doubling point and the folding point become.

To find out exactly how the probability of winning jumps and what effect this has on doubling strategy, we computed the probability of winning and optimal strategies for a simplified version of the later part of a backgammon game. This part of the game is essentially a race to move all one's men to the end of the board. With each turn players throw dice to determine how far they can move. Expert players keep track of the 'count,' which is the sum of the distances remaining for all their men. If it were not for the

rule on doublets, we could use the standard statistical results on the distribution of the number of rolls necessary to make $x_1 + \dots + x_n \geq S$, where the x_i are independently drawn from a normal distribution with mean 7 and variance 5%. (A doublet throw of the dice counts 4 times in backgammon. Thus a pair of fives counts 20 instead of 10.) Instead we 'rolled' (on a computer) the dice until they added up to 125, keeping track of when their sum exceeded each number less than 125. We repeated this procedure 1,000 times and used the resulting distribution to compute Table I. The numbers in Table I give the probability that a player will win when it is his turn and his opponent's count is at various distances from his own. Only a large doublet makes a really big shift in the probability of winning. For example, when the counts are 100 and 90, a before-the-throw prob-

TABLE I
PROBABILITY OF WINNING THE RUNNING GAME WHEN IT IS YOUR TURN

Opponent's count	Your count							
	30	40	50	60	70	80	90	100
Opponent's - 10	0.27	0.31	0.33	0.36	0.37	0.38	0.38	0.39
-5	0.47	0.47	0.47	0.48	0.48	0.48	0.48	0.48
Even	0.65	0.62	0.61	0.60	0.59	0.59	0.58	0.58
5	0.79	0.75	0.73	0.70	0.69	0.68	0.67	0.67
10	0.87	0.84	0.81	0.79	0.78	0.77	0.75	0.74
15	0.93	0.91	0.88	0.86	0.84	0.83	0.82	0.81
20	0.97	0.95	0.93	0.91	0.90	0.89	0.87	0.86

ability of winning of 0.39 goes to 0.33 with a 5, 0.42 with a 10, and 0.62 with a 20.

We next computed the optimal doubling strategy for a simplified form of the game in which the dice could show either 5 with probability 0.55, 10 with probability 0.36, and 20 with probability 0.09. (The 5 was supposed to replace 3, 4, 5, 6, 7, the 10 to replace 8, 9, 10, 11, 12 and the 20 to replace 16, 20, and 24. The probabilities are set to make the mean $\sim 8\frac{1}{2}$ and the variance $\sim 18\frac{1}{2}$, which are the real backgammon values.) Such a three-valued die is qualitatively very similar to the actual backgammon dice. Indeed, the probabilities of winning shown in Table II are close to those in Table I in the critical range for doubling. We can compute the optimal actions by dynamic programming, working backwards from the end, and we see that the doubling point approaches 0.8 rather slowly as we move away from the end of the game. Table III gives the probabilities for making the first double, for redoubling, and for folding at various points in the game. (Because utility is assumed to be linear, the size of the stake

TABLE II
A SIMPLIFIED RUNNING GAME

Opponent's count	Your count							Probability of winning if it is your turn
	30	40	50	60	70	80	90	
Even	0.64	0.62	0.61	0.60	0.59	0.58	0.58	0.57
+5	0.80	0.76	0.73	0.71	0.70	0.68	0.67	0.66
+10	0.91	0.87	0.84	0.81	0.79	0.78	0.76	0.75

Optimal Action ^(a)								
	F	R	R	D	D	F	R	R
	F	F	F	F	F	F	R	R
Even								
+5								
+10								

	Value								
Even	0.48	0.45	0.42	0.40	0.38	0.38	0.38	0.38	0.37
+5	1.0	0.92	0.78	0.68	0.65	0.63	0.62	0.60	0.60
+10	1.0	1.0	1.0	1.0	1.0	0.97	0.90	0.84	0.84

(a) D = You can make first double; opponent should accept.
 R = You can also redouble; opponent should accept.
 F = You can also redouble; opponent should fold.

is unimportant, and the redoubling point is the same as the re-redoubling and all further doubling points.) By examining Table I, we see that a simple formula approximates all these late-game actions. If your count is x , and your opponent's is y :

you double if $y \geq (11/10)x - 1$; he folds if $y \geq (11/10)x + 2$.

[Another peculiarity in the rules (bearing off) causes a slight wastage of pips on the last few rolls of the dice. This means that the count we have called 50 corresponds to an on-the-board count of 45-50.]

TABLE III
CUT-OFF POINTS FOR OPTIMAL STRATEGY

	Your count						
	30	50	70	90	110		
First double	0.65	0.68	0.70	0.72	0.73	0.73	0.73
Redouble	0.69	0.71	0.72	0.73	0.74	0.74	0.74
Fold	0.76	0.77	0.78	0.79	0.79	0.79	0.79

APPENDIX

GREAT EXPECTATIONS

There is a theoretical flaw in our treatment of backgammon doubling points. We have tacitly assumed that the expected value of the game, given doubling strategies for A and B exists. It need not. Suppose A 's doubling point is 0.8, which we consider optimal. Now imagine that B sets his doubling point at 0.8. (That is, when A has a 0.8 chance of winning. In this Appendix all doubling points will be expressed in terms of A 's probability of winning.) When the game reaches 0.8, the two players alternate doubles, ad infinitum. The game, like life, becomes an absurdity.

The problem of infinite expectations may appear in more subtle guises. Suppose A 's doubling point is 0.8 and B 's doubling point is 0.7. Let us calculate the expected value to A , beginning at 0.7 with A having the doubling die.

We concern ourselves only with the values 0, 0.7, 0.8, 1, and the transition probabilities as indicated in Fig. 1. For $n \geq 0$ with probability $\frac{1}{8}(\frac{7}{8})^n$, the game




Figure 1

has $2n$ doubles and then B wins 4^n . With probability $\frac{1}{8}(\frac{7}{8})^n$, the game has $2n+1$ doubles and then A wins $2 \cdot 4^n$. Thus the expected value to A is

$$\sum_{n=0}^{\infty} \frac{7}{8} \left[\frac{7}{8} \left(\frac{2}{3} \right) \right]^n 2 \cdot 4^n - \sum_{n=0}^{\infty} \frac{1}{8} \left[\frac{7}{8} \left(\frac{2}{3} \right) \right]^n 4^n.$$

However, both of these sums are *infinite*. A 's expectation has both infinite positive and infinite negative components and hence must be regarded as undefined. We easily generalize this case by

THEOREM 2. Let B 's doubling point be L and A 's doubling point be $L+M$. Set $R = 1 - L - M$. (Assume $0 < L, M, R$.) Suppose both players accept all doubles. Then the game has undefined expectation if and only if

$$3LR \geq M. \quad (13)$$

Proof. The situation is exactly as in the example, with terms of the form

$$[LR/(L+M)(R+M)]^n 4^n.$$

Under assumption (13), $4LR/(L+M)(R+M) \geq 1$; hence the expectation becomes undefined.

As an example, suppose $L = R = 0.35$, $M = 0.3$. Then (13) holds and the expecta-

tion is undefined. Intriguingly, this is the doubling strategy advocated by many backgammon experts. It certainly will lead to an exciting game.

What effect does all this have on the practical player? Our answer is: 'essentially none.' There are some critical differences between the 'pure' game and any real game.

(a) The real game is not continuous.

(b) In a real game, assets are limited to some finite amount.

Suppose that the number of doubles by each player is limited to n . Since this limits the stakes to 4^n , infinite expectations cannot occur. We have shown in our paper⁽⁶⁾ that the appropriate first doubling point is then $\alpha(n) = 0.8 + (-1)^n / 5.4^n$, which rapidly approaches 0.8.

What are the practical implications of Theorem 2 if, for example, $L = R = 0.35$? Since the actual game is not continuous, the expectations are finite. In the pure game the expected value of the stake (that is, the amount that changes hands) is infinite. The closer a real game may be mirrored by a continuous game, the larger the expected stake will be. If, for example, the game consists of watching a random walk on $\{0, 1\}$ with steps of 10^{-4} , the expected stake will be large indeed. We do not see how to estimate the expected stake in backgammon itself, where—as a very practical problem—it is nearly impossible to calculate the probability of winning with any precision except in the end game.

In the pure game, if A 's doubling point $= L + M = 0.8$, there will be undefined expectations if B 's doubling point $= L \geq 0.5$. Even more subtle possibilities may occur. If B 's n th doubling point is $L_n = 0.5 - \epsilon_n$, where ϵ_n rapidly approaches zero, then, with all doubles accepted, there will be undefined expectations. However, there are 'reactive' strategies for A that avoid infinite expectations. Essentially, one accepts 'absurd' doubles but does not redouble. The following strategy gives finite nonnegative expectation for A against any strategy for B .

Reactive Strategy: Initially, A 's doubling point is 0.8. If B doubles at p :

- (i) If $p < 0.2$, resign.
- (ii) If $0.2 \leq p < 0.49$, accept and make 0.8 the doubling point.
- (iii) If $0.49 \leq p < 0.51$, accept and make 0.9 the doubling point.
- (iv) If $0.51 \leq p$, accept and throw away the doubling die.

Of course, one may do better against poor players by fine-tuning the doubling point.

ACKNOWLEDGMENT

We would like to thank Thomas Cover for his many helpful suggestions and encouragement.

REFERENCES

1. C. GOREN, *Goren's Modern Backgammon Complete*, Doubleday, New York, 1974.
2. O. JACOBY AND J. R. CHAWFORD, *The Backgammon Book*, Viking, New York, 1970.
3. E. KEESLER AND J. SPENCER, "Proper Raising Points in a Generalization of Backgammon," The Rand Corporation, P-4078, May 1969.