Hoosier Backgammon Club

September-October 1995
Volume XII, No. 5

Rollout Techniques

by Fredrik Dahl

The methods and ideas presented in this article come from my work with the **Jellyfish Analyzer 1.0** backgammon computer program, but can be applied to any evaluation function based rollout program.

Cubeless Equities and Probabilities

A basic property of a position is **cubeless equity** as defined as the average number of points one side wins from a position. A single win counts as 1 point, a gammon loss counts as -2 points and so on. The definition assumes that both sides play optimally to maximize their cubeless equity on each turn. This may sound like a circular definition, using **optimal play** to define **cubeless equity** and vice versa. That is not a problem, however, and if one assumes that the probability of a game never ending to be zero, existence and uniqueness of the cubeless equity can readily be proven using game theory. It's non-trivial to prove that infinite games are not a problem, but we all feel it's not.

Every position has a **cubeless equity** somewhere between -3 and 3. Optimal (cubeless) play is defined by choosing the highest cubeless equity on each roll. It will be unique most of the time, but sometimes there may be several equivalent choices.

The **cubeless equity** can be broken down into probabilities. Let's assume we see the position from Blacks point of view. Probabilities are defined in the following way:

P(W) = probability of Black winning a single game, a gammon or a BG
P(G) = probability of Black winning a gammon or a BG
P(B) = probability of Black winning a BG
P(-W) = probability of Black losing a single game, a gammon or a BG
P(-G) = probability of Black losing a gammon or a BG
P(-B) = probability of Black losing a BG

This gives **P(W) + P(-W) = 1**.

And the **cubeless equity** is given by
EQ = P(W) + P(G) + P(B) - P(-B) - P(-G) - P(-W).

Two samples follows:

Sample #1 - Black on roll.

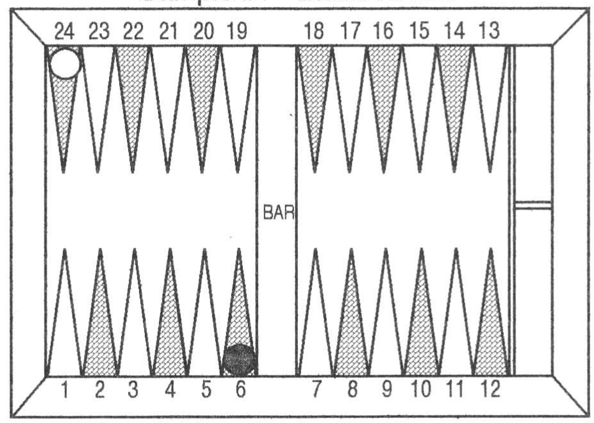

EQ = P(W) + P(G) + P(B) - P(-B) - P(-G) - P(-W)
+0.50 = 0.75 + 0.0 + 0.0 - 0.0 - 0.0 - 0.25

Sample #2
7-point match Black-3 White-3
Black to play 5-2?

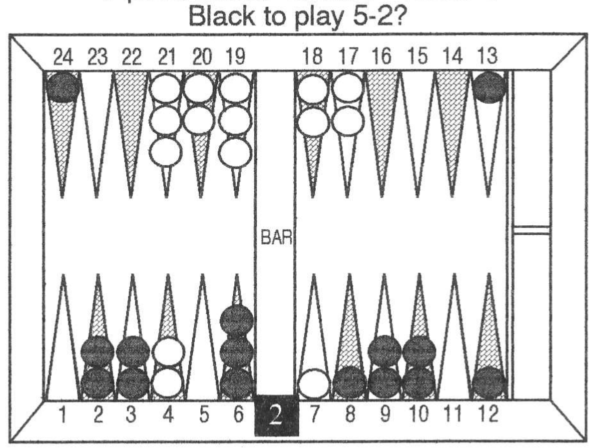

Play 1: 13-8, 12-10
-0.470 =0.315 + 0.047 + 0.001 - 0.685 - 0.144 - 0.005

Play 2: 12-7x, 9-7
-0.455 =0.346 + 0.097 + 0.003 - 0.654 - 0.233 - 0.013

Play 3: 13-8, 24-22
-0.577=0.303 + 0.058 + 0.002 - 0.697 - 0.234 - 0.008

Play 4: 12-7x, 24-22 (best play)
-0.401=0.378 + 0.100 + 0.003 - 0.622 - 0.247 - 0.013

...continues on Page 3...

Hoosier Backgammon Club's Newsletter for HBC members and subscribers.
Subscription rate: $12/year (Canada $14 and overseas $16). Let us know if your address changes.
Butch & Mary Ann Meese: (317) 845-8435. 7620 Kilmer Lane, Indianapolis, IN 46256-1634 E-Mail: hbc@ix.netcom.com

1995 **HOOSIER BACKGAMMON CLUB** Gammon Point Standings.
HBC **Player of the Month** for **July** was **Mary Ann Meese** with 202 gammon points.
HBC **Player of the Month** for **August** was **Larry Strommen** with 160 gammon points.

1)	Butch Meese	1134	Chuck Bower	258	Philip Degen	68	Richard Heinz	16
2)	Chuck Stimming	980	G.L. Harvey	168	Steve Perlman	64	Lance Jenkins	16
3)	Ellis Bray	900	Kevin McLeaster	150	Wendy Kaplan	60	Elijah Miller	16
4)	Don Woods	896	Peter Kalba	116	Dave Cardwell	50	Tom Helt	10
5)	Woody Woodworth	892	Rick Reahard	110	Jim Curtis	40	Dave Fey	10
6)	Larry Strommen	858	Scott Richardson	96	Bill Hodes	40	Alice Gerard	10
7)	Dave Groner	636	Brian Nelson	88	Reggie Porter	32	Dave Williams	10
8)	Gabe Stiasny	564	Karen Davis	84	Fred Badagnani	30	Stephen Maas	10
9)	Mary Ann Meese	562	John Brussel	82	Jim Bowman	20	Carol Falk	10
10)	Sean Garber	552	J.A. Miller	80	David Smith	20	Ralph Stowell	10
	Jan Gurvitz	408	Randall Witt	80	Bill Julian	20	Paul Ruterman	10
	Neil Ezell	290	Stan Gurvitz	76	Bob Cassell	20		
	Mick Dobratz	260	Bill Gheen	70	Jon Stephens	20		

Hoosier Pips: Bill Hodes appeared on Larry King Live August 28th in a debate with Alan Dershowitz about the Simpson Trial...Neil and Theresa Ezell are proud parents of Alexander Neil Ezell born August 9th...Wendy Kaplan and David Smith have moved to the Chicago area...Mick Dobratz has moved to take a new job in Madison, Wis...July and August visitors included Fred Badagnani, Peter Kalba, and Dave Cardwell, plus all the early arrivals for the Thursday evening before the **43rd Indiana Open**.

	July 6th	July 13th	July 20th	July 27th	July 30th
1st	*Larry Strommen*	*Butch Meese*	*Woody Woodworth*	*Butch Meese*	*Mary Ann Meese*
2nd	*Chuck Stimming*	*Chuck Stimming*	*Sean Garber*	*Don Woods*	*Peter Kalba*
2nd	*Mary Ann Meese*	...	*G. L. Harvey*	*Sean Garber*	*Dave Groner*

	August 3rd	August 10th	August 17th	August 23rd	August 30th
1st	*Larry Strommen*	*Mary Ann Meese*	*Chuck Stimming*	*Larry Strommen*	*(s) Karen Davis*
2nd	*Chuck Stimming*	*Butch Meese*	*Scott Richardson*	*Butch Meese*	*(s) John Brussel*
2nd	*Ellis Bray*	...	*G. L. Harvey*	...	*Randall Witt*
2nd	...	...	...	...	*Ellis Bray*

Backgammon Tournament Schedule
Oct 06-08..... **Nation's Capital Fall Championships**, Promenade, Bethesda, MD.......................... (301) 530-0604
Oct 11-15..... **5th Illinois Championships & America Cup**, Sheraton, Northbrook, IL.................. (708) 945-7801
Oct 27-29..... **Autumn Gran Prix**, Embassy Suite Hotel, La Jolla, CA...(619) 294-2007
Nov 3-5........ **New England BG Championships**, Oak n' Spruce Resort, South Lee, MA..............(603) 863-4711
Nov 10-12.... **NY/NJ BG CO-OP Big Apple Series**, Oritani Hotel, Hackensack, NJ.........................(201) 833-2915
Nov 24-25.... **1st Annual *No Turkey* Weekend Tournament**, Best Western, Akron, OH................ (216) 966-2811
Nov 24-26.... **Georgia Backgammon Championships**, DoubleTree Hotel, Atlanta, GA................. (770) 441-2074
Thursdays...... 7:00 PM at **SPATS** (842-3465) Castleton Square (between J.C.Penney's & L.S.Ayres).... 845-8435

Rollout Techniques
by Fredrik Dahl
...continues from Page 1...

Most computer programs have an evaluation function tuned to approximate the cubeless equity. Note that the cubeless equity is not exactly what we need in normal play. If the cube is in play, adjusting to market losers will sometimes make a play with lower cubeless equity correct. And if the cube is dead in a match, the relative value of gammons and backgammons will often make a difference in checker play. Also the correct cube action cannot be completely inferred from the cubeless equity. For a take/drop decision, the future value of the cube must be taken into account. And for *no double/double* problems, the probability and size of market losers as well as future value of the cube are important. Nevertheless, having good estimates of cubeless equities is very useful.

General Properties of Rollouts

For any rollout, human or computer, there are two kinds of errors: *systematic* and *random*. *Systematic errors* come from a misplay of the position, or misevaluation of the final position if the games are truncated. The *random error* comes from the randomness of the dice rolls. Most computer random generators are not truly random, as they are the results of a deterministic, but very complex, calculation starting from some initial value. The initial value, or *seed*, is either given by the user or generated by the decimals of the computers clock or from some other external source. In general these *quasi random* numbers are random enough for the purpose of backgammon rollouts.

Random errors are easier to estimate than systematic ones. Random errors are easier to control, as they can be estimated from the rollout itself. There is a general law that says the random error is inversely proportional to the square root of the number of trials.

This can be illustrated by the following example: assume the position is a race, so that there are no gammons, and that in reality the sides win 50-50. Obviously the cubeless equity is 0.

From a single game **rollout** the random error is bound to be 1, as the result will be either -1 or 1. The quoted law says that if you do 100 games, the random error from your estimate of the equity (given by the average), is 1/sqrt(100) = 0.1. So even after a 100 games you may well be 0.1 off from the true value 0.

And if you want the random error to be less than 0.01, you must do 10,000 games! Variance reduction methods may reduce the size of the random error, but you must still do 100 times more games to reduce the error by a factor of 10.

Note that the length of the games does not enter into the calculations, and the random error is in general only dependent on the number of games and the probabilities of the outcomes. Although it may feel that way, the complexity of the games does not affect the random error of a rollout. But of course it affects the number of games one is able to play over a given time interval.

Variance Reduction Methods

There are various tricks one can do to reduce the random error of rollouts. It's possible to adjust the number of hits and misses of crucial shots, and to compensate the results for dice luck. That's quite complicated, and I will not go into that here, but rather concentrate on general dice roll stratification.

One way of doing this is what I call the **deck-of-cards** methods. Before the time of computer rollouts, I used this technique for hand rollouts. It works in this way: you have a number of decks of 36 cards each, one for every roll. So in each deck there is a card for 11, one for 51, one for 15, and so on. You shuffle the decks one by one, and then play a set of 36 games. In the first game you take the top cards of the first deck for the opening roll, the top card of the second deck to get the second roll of the game, etc. In the second game you take the second card of each deck, and so on.

The deck-of-cards method introduces no systematic errors, and reduces the random errors reasonably well in positions where the quality of a specific roll is relatively independent of the previous rolls. For example it works well in races, where a 66 on roll 5 is going to be good regardless of the value of the first 4 rolls. The deck-of-cards method guarantees that you get a fair cross section on each roll.

It also works quite well in prime positions where a given roll, like 55, may be a root number for a number of rolls into the game, and where a given number, like a 6, will be needed to jump the prime for some time.

This method does not accomplish much in most middle game positions where, for example, a 66 on roll 4 may either be a perfect roll or a dancing number, depending on the previous rolls.

If you do a multiple of 36 x 36 = 1296 games, you can cycle through every combination of first and second rolls. This cross section method can be used together with the deck-of-cards method. It is in general overrated, and only gains significantly if there is a very big swing on the first two shakes of the game. Doing this for a rollout of the opening position (with one side on roll, allowing doublets), for example, will only take away about 3% of the random error.

...continues on Page 4...

N-Ply Lookahead

An important concept in computer backgammon is **lookahead** or **ply**. Assume that you have a position, and want to use your evaluator function to estimate its cubeless equity. The evaluator function can be the output of a neural net, a hand-crafted function or something else. The obvious thing to do is to use the output of your evaluator function. This is called the **0-ply** evaluation, or **no lookahead**.

To calculate the **lookahead**, or **1-ply** evaluation, you do the following: for all 36 dice rolls you pick the play your evaluator function prefers, and take the average of the (0-ply) equity it gives for each of these 36 positions. Of course, you only have to calculate 21 rolls as a 62 is equivalent to a 26.

If each roll gives about 20 different legal plays, you have to evaluate more than 400 positions to find the lookahead evaluation. The lookahead evaluation is important because it improves evaluations, but it also has theoretical significance for the following reason. The correct evaluation function is the only function that gives the same result with and without lookahead for all backgammon positions, and evaluates final positions correctly.

The *2-ply evaluation* is calculated in the same way, except that you calculate the 1-ply evaluation for each of the 36 positions and take their average. The 2-ply, or double lookahead evaluation, is the average evaluation two rolls into the game. In general n-ply evaluation (where n is some positive number) is the average evaluation n rolls into the future. It can be shown that when n goes to infinity, the n-ply evaluation converges to the same number as a rollout of infinitely many games.

Experience with Jellyfish and other neural nets has shown that going from 0- to 1-ply in evaluation gives a big improvement in playing strength. Deeper lookahead gives even better results, but quickly becomes too time consuming to calculate.

The difference between playing and evaluating can give confusion, as choosing from the list of positions can be considered an extra ply. Therefore, using 1-ply evaluation to choose the play, can be called 2-ply play.

Truncated Rollouts

The accuracy of neural net evaluation functions has made it possible to use truncated rollouts. They are done in the following way: the user specifies the **number of games** and the **horizon**. Each game is stopped after horizon number of rolls (or earlier, if the game ends before this happens). The evaluation of that last position is given as result of the game. If the horizon is n, the rollout will converge towards the n-ply evaluation, by definition.

A truncated rollout can be seen as a combination of a deep lookahead and a rollout. Truncated rollouts are both faster and less random than full game rollouts. They are faster because fewer moves have to be made in each game, and less random because the games don't have the time to diverge very much. If you specify a given tolerance for random error, you therefore need vastly more time for full rollout games than for truncated ones.

For checker play problems, truncated Jellyfish rollouts are very effective and accurate. In most positions truncated rollouts with horizon 7 have a bias of less than 0.05 compared to a full rollout. Here is a rollout example from one of my own games:

Money game. Black to play 6-4?

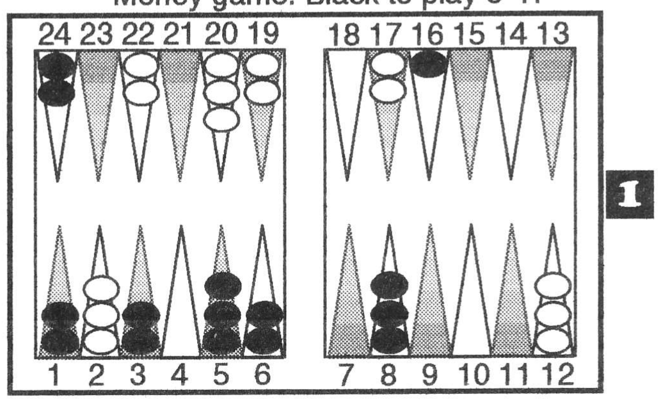

I misplayed this position in a game against Kit Woolsey on FIBS, making the obvious looking (to me) 16/6. The correct play is 24/14. My play is simply too shortsighted, and runs into trouble fast quite often. Jellyfish hated my play, and the truncated rollout did not change its mind. Here follows the equity estimates for the two plays as a function of the horizon, generated with 1296 games for both play:

Horizon	Play 24/14	Play 16/6
0	+0.005	-0.099
1	+0.003	-0.084
2	-0.037	-0.109
3	-0.026	-0.114
4	-0.029	-0.131
5	-0.037	-0.148
6	-0.050	-0.165
7	-0.055	-0.167
Full game rollouts with 5184 games	-0.077	-0.194

This case shows an example of a position where Jellyfish evaluates the difference between plays better than the absolute equities. This is typical with neural nets as well as humans. Each of the full game rollouts took about an hour, and each truncated rollout took only 3 minutes on a Pentium 90, and still the randomness is smaller for the truncated ones.

Rollout Techniques
by Fredrik Dahl
...continues from Page 4...

Full Game Rollouts with Cube Action

Truncated rollouts are very good for checker play problems, and in most cases they also give very good absolute cubeless equities. Sometimes, however, the truncated positions will be systematically misevaluated. For Jellyfish this may happen if it ends up evaluating deep backgames with good timing, which it systematically underestimates. This will not normally be a problem for a checker play rollout, where the choices mostly generate similar types of positions.

If the goal of the rollout is to find the correct cube action, however, bias of this kind will be harmful. This is why we still need full game rollouts, and at the same time we can get estimates of the value of the cube for both sides.

Obviously, the quality of rollouts with cube action is always extremely sensitive to cube errors in the games. This also applies to human rollouts, of course. But even if the program gives very accurate cubeless equities, it's not obvious how to use the cube in rollouts.

One could start by asking *how do human experts do it?* But this I really don't know! Often you see experts quoting hand rollouts without mentioning if and how the cube was used. This has always puzzled me.
To get a completely unbiased result one would have to imitate the cube actions of money play, where most cubes should be taken. This would be a foolish thing to do, as it increases the randomness a lot for little gain. What I assume they do is settle the games, somehow. One could play until the position is a cash, and then count it as 1, or one could wait till it's a double and count it as 1. The first way underestimates the value of the cube, as you lose your market more often than you should, thereby costing you equity. The other policy overestimates the cube value, as the opponent will be assumed to drop takeable positions. A better way would be to settle the game as 1 point when it's a strong double, but before it's a drop. Then you will give up some equity by losing your market too often, and gain some by having him drop some takes. One also needs some strategy for deciding if one should play on for a gammon or cash the game. Simply playing on if the equity is higher that 0.9 and gammon is possible is reasonable.

In Jellyfish Analyzer, the user can specify the settlement limit, in terms of cubeless equity estimates, for settling the games. In the program, a cubeless equity of 0.5 is given as a default value. In full rollouts equities can be estimated for all four states of the cube:

1) Cubeless
2) Central cube
3) Black holds the cube
4) White holds the cube

For all these states we assume the cube is on 1, although that would not have been the case in a game.

By playing all the games to the conclusion all these four equities can be estimated in parallel. You simply have to tabulate the results during the games. When a game is over, you will know who would have won with the cube in the center, with Black holding the cube, and with White holding the cube.

Here follows an example, where all the quoted eq's are estimated equities from Blacks point of view.

First eq=0.1. Then the equity moves up to 0.55 with Black on roll. If Black holds the cube (case 3), or it's in the center (case 2), Black gets 1 point. Then the equity falls down to -0.60 with White on roll. White wins the game where he holds the cube (case 4). Then the game turns around again, and Black eventually wins a gammon. Black therefore wins 2 points in case 1.

This example is typical in that gammons are much less common when the cube is used to settle games. It may be a good idea to use this procedure also in hand rollouts.
...continues next page...

Fredrik Dahl is 28 and lives close to Oslo Norway. He is married with one child and another on the way. He has a masters degree in computer science and works at the Norwegian Defence Research Establishment. His hobbies include playing backgammon and other games, drinking beer and programming neural nets.

The Value of the Cube

Experimenting with cube rollouts has led me to believe that many authors underestimate the value of the cube in gammonish positions. After discovering this, it turns out to be quite easy to support with mathematical arguments.

In a race you can never take a cube with less than 20% cubeless winning chances, except for a few late positions. This transforms into never being able to take with a cubeless equity of more than 0.6 for the opponent. So far so good, but this is not true for gammonish positions. The point is that the value of holding the cube is proportional to the probability of being able to use it. And that is roughly proportional to the cubeless winning probability of the position. If you compare a blitz position and a race, both with cubeless equity 0.61, then the underdog wins a lot more games in the blitz position, to make up for the gammons. Therefor the value of holding the cube is much higher in the blitz position, most likely making it a take.

The following classical opening blitz is a good example.

Money Game, Black on Roll, Cube Action?

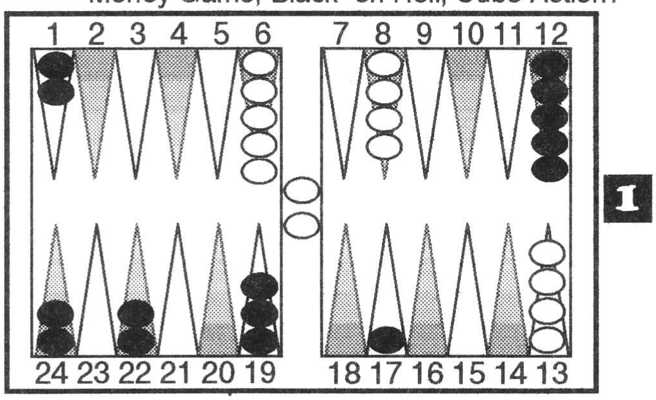

Jellyfish plays these attacks very well, and also evaluates the relevant redoubling positions well enough to be trusted, in my opinion. The conventional wisdom has been that this is a clear drop, and that the cubeless equity is in the mid 0.60s. Indeed, Jellyfish rollouts get the cubeless equity to be 0.65, but still has it to be 0.51 with the underdog holding the 1-cube, making it a very close take/pass.

I have designed the following formula for the value of the cube in middle game positions:

Cube value = 0.35 * (P(W) + P(G)/2)

This measures unit cube equity, so it may be a bit confusing. Here is how it works: assume the cubeless equity is -0.6 for you, and that you win 30% of the games cubeless, whereof 10% are gammon wins. Then the formula gives

Cube value = 0.35 * (0.3 + 0.05) = 0.12.

So according to the formula the equity with you holding an 1-cube is -0.6 + 0.12 = -0.48. As this is less than 0.5, the formula says you should take.

With the high level of play by neural net based programs like Jellyfish, computer rollouts will be very important for the serious student of the game. Techniques that squeeze more information out of each second of computer time, like truncated rollouts, are indispensable if you want to analyze a large number of checker plays. The main task for backgammon players from now on may be to absorb the information presented by computer rollouts, and make the correct generalizations from the results. Rollouts only give numbers, not reasons, and a player who only tries to memorize the equities of positions will never succeed. Also remember that backgammon is a game, and the goal is not necessarily to make as few mistakes as possible, but to have the opponent make more and bigger ones. Against a weaker opponent it may often be correct to make a slightly inferior play to make the game more difficult. But to do this you need to know which plays are close, and which are not, and a computer rollout program like Jellyfish Analyzer is just what you need to get this knowledge.

Both program versions, **Tutor 1.2** and **Analyzer 1.0**, are available from Carol Joy Cole at (810) 232-9731 or email: carlcole@sils.umich.edu or Larry Strommen (317) 545-0224 or email: diceman@indy.net.

Annotated match
Kit Woolsey vs Jeremy Bagai
FIBS - 9 Point Match

In February 1994, **Kit Woolsey** and **Jeremy Bagai** played a match and then annotated it for **FIBS** (**F**irst **I**nternet **B**ackgammon **S**erver) players so they could see the thought process of the more experienced players. They played a fairly interesting match, logged it, and then annotated it independently. You will see reasons for their plays and cube decisions, as well as their second thoughts upon later analysis which often came to a different conclusion than their original choices.

Gerry Tesauro also volunteered **TD-Gammon's** valuable help. **TD-Gammon** analyzed the whole match and listed its top 3 choices for each play along with its estimated equities. These equities are always assuming a 1-cube and they do not take into account cube ownership. Thus on a pass-take decision an equity of -0.50 would be a break-even decision (not taking cube ownership into account -- that would probably make it a little higher), since that would translate to an equity of -0.100 on a 2-cube. **TD-Gammon** was also nice enough to comment on the game, giving its reasons behind its choices as well as getting in a few snide remarks about their mistakes. **Mark Damish** (MA), first formatted the commentary for the Internet.

Game 5 Continues...

Black (Kit) enters both checkers 51.

White (Jeremy) to play 52?

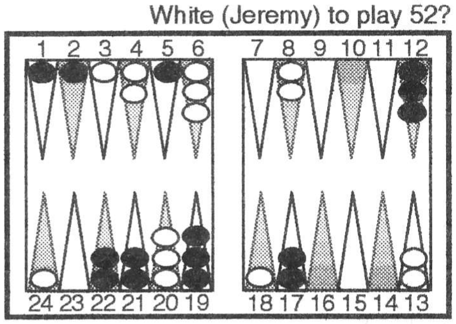

8/6 8/3

Kit: Not much choice. The alternative is 20/18, 8/3, but this would give me a double shot at the blot on his eight point and if I hit it he would be completely out of ammunition up front. Leaving the blot on my bar point isn't nearly as dangerous. Making my bar point isn't so important since he already owns my five point.

Jeremy: This looks better than anything else.

TD-Gammon: I'm afraid you guys missed the boat on this one. 20/18, 8/3 is much better. At least Kit found the play, but then rejected it because of the double shot on the eight point. That was silly. With the double anchor Jeremy would be well in the game whatever happened. The actual play piles a stack of checkers on the six point and gives Kit a free shot to attack. Note that the various plays which go to the ace point are better than Jeremy's actual play.

20/18, 8/3	-0.174
18/13, 3/1*	-0.182
6/1*, 3/1	-0.218
8/6, 8/3	-0.249

Black (Kit) to play 54?

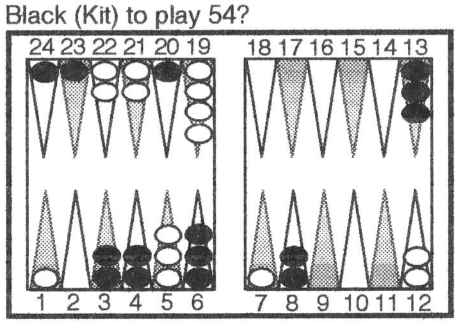

24/20 23/18

Kit: It looks more important to pump the checkers into the outfield rather than bring a checker down with 24/20, 13/8. The checker on his bar point is in little danger, and I need all the flexibility I can get in the outfield here.

Jeremy: 24/20, 13/8 looks more natural but I think Kit found a nice play. The key is that I'll be very ruluctant to hit from my stripped midpoint, isolating my rear checkers. I probably wouldn't have found this play.

TD-Gammon: After all the work I've been doing teaching the importance of getting checkers out from behind an enemy blockade and into the outfield in positional struggles, I would have been very disappointed if you got this one wrong. Actually, I'm surprised it came out as close as it did. Thanks for not failing me.

24/20, 23/18	+0.237
24/20, 13/8	+0.230
24/20, 6/1*	+0.229

White (Jeremy) to play 65?

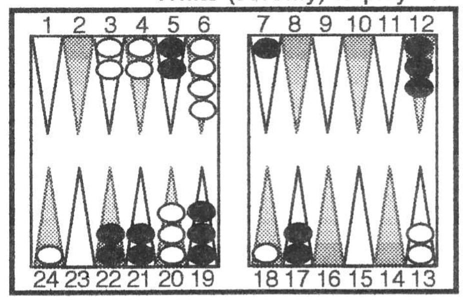

18/7*

TD-Gammon: This is correct, but it is not as automatic as it might seem. Making Kit's bar point and getting the back checker out of hock is very strong for Jeremy here.

18/7*	-0.194
24/18, 20/15	-0.220
24/13	-0.232

Black (Kit) dances with 63.

White (Jeremy) to play 21?

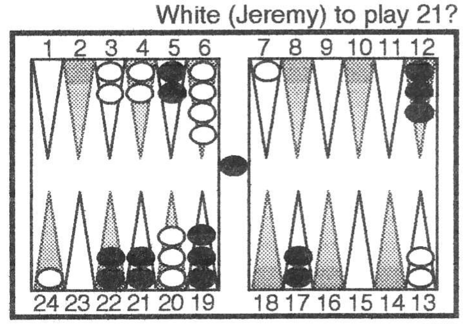

7/4

Jeremy: Leaving the bar point slotted in the hopes of making it looks like a mistake because there's nothing to cover it with.

TD-Gammon: Correct. Leaving the blot there can only lose.

7/4	-0.149
20/18, 7/6	-0.173
24/23, 6/4	-0.188

Black (Kit) to play 41?

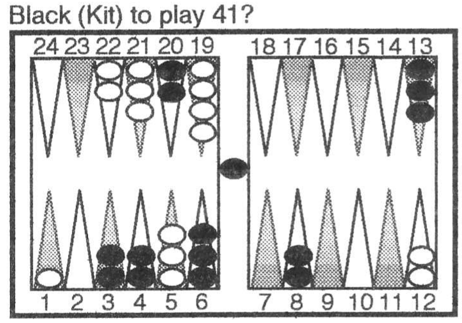

B/20

White (Jeremy) to play 54?

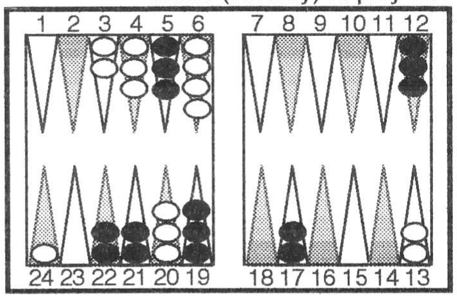

24/15

Kit: I agree. 24/20, 6/1 dumps another checker behind my anchor and leaves an ugly stack of four checkers on my five point. If Jeremy is hit he should have plenty of time to enter and try again, and if the blot is missed the extra checker in the outfield gives him some badly needed flexibility.

Jeremy: This feels better than 20/11, but I can't really say why. It does keep my spares together for return shots and point making.

TD-Gammon: It is way better, Jeremy. The real key is liberating the back checker so you will have some flexibility later on. That has overriding priority here.

24/15	-0.178
24/20, 6/1	-0.228
20/11	-0.252

Black (Kit) to play 41?

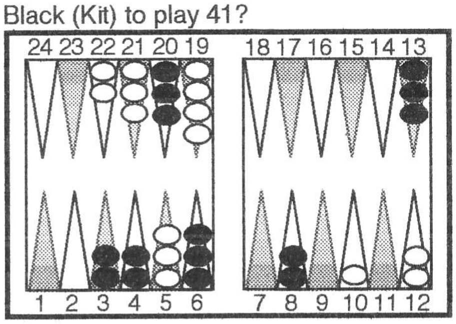

13/8

White (Jeremy) to play 61?

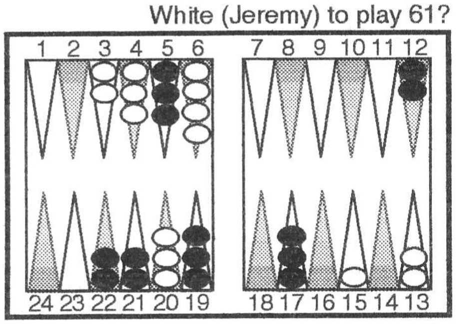

20/14 15/14

Black (Kit) to play 42?

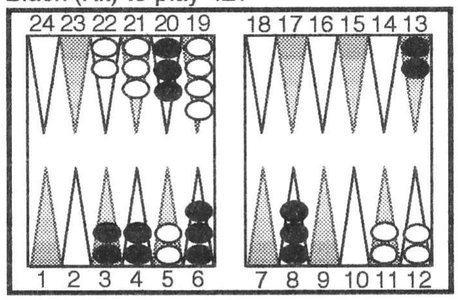

8/2

TD-Gammon: This is correct, but it is not as automatic as it might appear. Any time you are diving behind your opponent's anchor, it is well worth searching for alternatives.

8/2	+0.005
20/14	-0.013
20/18, 8/4	-0.023

White (Jeremy) to play 66?

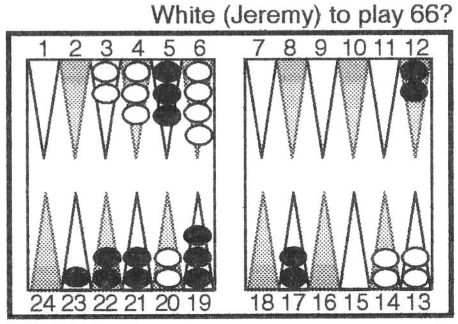

20/14 20/14 13/7 13/7

Kit: This looks better than 20/8(2). The latter play would leave a stripped position with two outfield points to clear, so he would quickly have some problems. His actual play leaves only one outfield point to clear and four checkers on the 14 point to give him some maneuvering room.

Jeremy: The pip count is even before the roll so I'm definitely moving my back anchor to disengage. 20/8(2) keeps all my checkers communicating but leaves me with three stripped points. My play isolates the checkers on the 14 point a bit, but gives me two spares to play with. It looks right.

TD-Gammon: I couldn't have said it better myself.

20/14(2), 13/7(2)	+0.386
20/8(2)	+0.371
20/14, 20/2	+0.254

Black (Kit) to play 42?

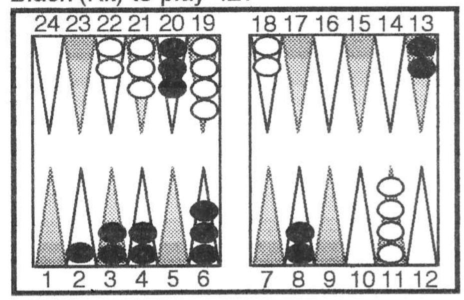

8/2

TD-Gammon: Only barely correct. Moving the spare on the 20 point has a lot going for it. Remember that Kit is behind in the race, so he doesn't mind provoking contact.

8/2	-0.402
20/14	-0.405
8/4, 8/6	-0.418

White Doubles?

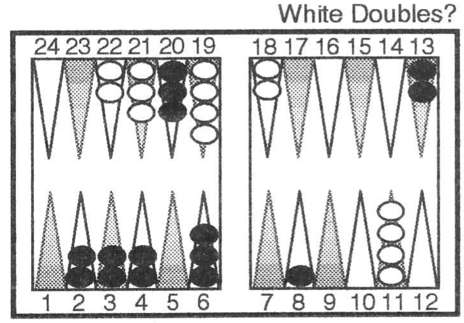

Kit: Close. Jeremy has a useful lead in the race, but it is not overwhelming. I figure to have plenty of shot potential as he tries to clear his outfield point, and my board is decent. One problem I may have is finding a safe play in the next roll or two; I'll probably have to move the spare off his five point even if it means leaving a direct shot in order to avoid wrecking my board. My take is easy, and he can lose his market only by rolling doubles. I can't say for sure the double is wrong, but I would have held off.

Jeremy: I'm up 18 pips and on roll. Very strong double. If the distance of the 14 point bothers you remember that Kit still has his five point open and might not be able to make it soon. Clear double.

TD-Gammon: Equity of +0.402, moderate volatility. Just barely worth a double, but it is there.

Game 5

...continues next issue...