A-

AMAZING MATHEMATICAL TREATISE
I LIKE YOUR USE OF QUOTES FOR THE SUB-SECTIONS OF THIS PAPER. YOU COULD HAVE QUOTED & USED SOME REFERENCE WORKS PERTAINING TO "21" COUNTING SYSTEMS. I'M IMPRESSED WITH THE DEPTH OF ANALYSIS AND THE THOROUGHNESS OF THIS PAPER. YOU HAVE A STRONG ABILITY FOR TECHNICAL WRITING.

FASCINATING PAPER!

OUTLINE?
WORKS CITED?

THE THEORY OF RED DOG

by

Victor Aguilar

INTRODUCTION

> Get a new leash on life,
> Learn Red Dog,
> Easy to play card game.
>
> casino slogan

Red Dog, also known as acey-deucey, is a relatively recent addition to the sampling without replacement games played in Nevada casinos, though it was quite popular in the old west. What distinguishes these games from such independent trial games as craps or roulette is that the expectation is not a fixed constant but rather a function of the outcome of previous plays. Whereas in craps or roulette the expectation for any given bet is determined by the physical parameters of the game; vis., the six sides of each die or the 38 places on the roulette wheel, the expectation in Red Dog varies as cards are removed from the pack. While the top of the deck expectation is disadvantageous to the player (-2.19%), there may be certain subsets of the pack which are advantageous to the player, and if he can recognize these situations and act accordingly, he may improve his overall odds.

The pack of cards used in Red Dog consists of six regular 52 card decks and is dealt out of a box called a shoe. Hereafter, the term "shoe" will refer to the set of all cards yet to be dealt. After each player makes his wager, the dealer places two cards face up in front of her which we will call the "initial cards". If the initial cards are neither a pair nor consecutive, then the players may double their bets if they wish. If they are consecutive, then they are discarded and two more cards are dealt without any money changing hands (a "push"). Otherwise, a third card is dealt face up and if its face value (aces are high and are thus valued at 14, kings at 13, etc.) falls between the values of the initial cards, the player wins. If it equals either of the initial cards or is less than the smaller of them or greater than the larger of them, the player loses. If the initial cards were a pair, the outcome is determined differently. In this case, if the third card equals the pair and thus makes three of a kind, the player wins 11:1 odds. If it fails to equal the initial cards, the player keeps his money: a push. Hence, pairs are good for the player because he can't lose and can possibly win 11 dollars for each one he bet. Unfortunately, as was said earlier, the player is not allowed to double his bet when the initial cards are a pair.

Obviously, the greater the difference between the initial cards, the more likely it is that the third card will fall between them. An ace and a deuce can only be beaten if the third card is also an ace or a deuce, while a difference of two between the initial cards (i.e., 3,5) will only win if one particular card (in this case a 4) is dealt as the third card. Unfortunately, initial cards with a small difference between them appear more often than initial cards with a great difference between them but this is partially made up by the casinos increasing the pay rate for getting a third card in between closely spaced initial cards. The complete payoff schedule is as follows:

	Difference between initial cards												
	0	1	2	3	4	5	6	7	8	9	10	11	12
win:	11	0	5	4	2	1	1	1	1	1	1	1	1
loss:	0	0	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1

Certainly, Red Dog is a much simpler game than the more famous sampling without replacement game, blackjack. Unlike blackjack, the player's only decision has to do with the size of his wager; nothing he does can alter the outcome of each play. There are no complex charts for the basic (i.e., non-card counting) player to memorize (the basic tactic is: double whenever there are seven or more cards between the initial cards) and for the more ambitious who aspire to count the cards, they go by so slowly that even the most drunken gambler can handle the mental arithmetic. This is a good thing because card counting may defray some loses but the only real profit the Red Dog player can expect to make is the free drinks the cocktail waitress brings him. Nevertheless, Red Dog provides an interesting exercise in probability.

MATHEMATICAL EXPECTATION

> But what did [the odds] matter to me?...
> I wanted to astonish the spectators by
> taking senseless chances.
>
> Dostoevsky

Mathematical expectation or advantage is defined as the amount of money the player can expect to win in the long run for each dollar he risks. The best way for the player to keep track of this empirically is to bring a certain quantity of money to the table (called a bankroll) and play until it is exhausted, putting the winnings in a separate pile. If he then counts his winnings, devides the total by his bankroll and subtracts one from the quotient, he will have an estimate for expectation whose accuracy is a function of bankroll size, confidence approaching unity as bankroll approaches infinity. Be warned, however, variance is what determines how large the bankroll must be to get any real confidence (95% to 99%) on this estimate and most casino games are variant enough that relatively long winning or losing streaks are not that uncommon.

To determine expectation deductively, one must assemble all the possible courses the next play could take and multiply the probability of arriving at each point along a course, taking into account how previous points may have altered the probability of achieving later points, to determine the chance of taking that course to its termination and then multiply this figure by the payoff, either positive or negative, at the termination of each course and then sum up. Here, each course is by definition mutually exclusive, that is, it has a unique result not shared by any other course (i.e., a certain combination of cards, not a certain payoff - the size of the payoff may be shared by several courses). If two courses share the same result then the distinguishing mark is a superficiality and they are really one course whose probability is the sum of the two superficially distinct courses; vis., initial cards 3,6 are the same as initial cards 6,3 and the probability is the chance of getting a 3 and then a 6 (considering that the shoe has been depleted by one when drawing the first card) or the chance of getting a 6 and then a 3 (same consideration). When calculating probabilities, the chance of several events, each of which are essential to the outcome (a logical "and"), must be multiplied by each other to get the probability of the outcome. When there are several mutually exclusive courses which lead to an outcome (a logical "exclusive or"), then the probabilities of the courses must be summed up. When there are several courses which may happen simultaneously (a logical "inclusive or"), then their complements, the chances of their not happening, must be multiplied and the result complemented (subtracted from unity).

Recalling from the definition of expectation that antecedent points may have affected the probability of certain points along a course, the term "linearity" will now be defined. A play is linear if at no point along a course of that play does the outcome of a previous point affect the probability of achieving that point. This is not to say that the previous points are not essential to the course, the course being defined as the several points happening one after the other, but rather that the chance of a certain point is the same as it would have been in another course with different antecedent points. In the example above, the chance of drawing the 6 after the 3 is not the same as it would have been in every other course, for if the first card had been a 6 also, then the chance of drawing the second 6 would be different than if no 6 had yet been drawn. Having gotten this pair, the chance of success (drawing another 6) is different than the chance of success if the initial cards had been 5,7, even though winning in either case is dependent on the third card being a 6. This is the first and most important case of nonlinearity in Red Dog. The second case of nonlinearity has to do with losing a non-pair play by drawing a third card which is equal to one of the initial cards. If the initial cards were 5,7 and the play is lost by drawing a 5, this has a different probability than if it had been lost by drawing, say, a 3. While in this example there are twelve losing plays (2...5 and 7...14), only two of which are nonlinear, an acey-deucey (initial cards 2,14) has only two losing plays both of which are nonlinear. Thus nonlinearity becomes more important as the difference between initial cards increases. Of course the probability of getting initial cards very far apart is less than for getting initial cards close together.

AN UNBALANCED POINT COUNT

> A false balance is abomination to the Lord,
> A just weight is his delight.
>
> Proverbs 11:1

Why was the distinction made between linear and nonlinear plays? Because the chance of winning a linear play is solely a function of the relative proportions of the several denominations of cards in the shoe; vis., 30/150 kings is the same as 5/25 kings when determining the probability of winning if it's a king that one wants. Since a card count does not distinguish between 30/150 and 5/25 kings, it works best as a predictor for linear plays. Pairs are nonlinear and intuitively one can see that there is a difference between getting three kings in the first case where only 10% of the available kings are needed and getting three kings in the second case where one would have to deplete the stock of kings by 60%. The expectation of a Red Dog play with a rectangular distribution of card denominations at the several different levels of remaining cards, T, in the shoe is:

remaining cards	expectation
312	-2.19%
286	-2.29%
260	-2.41%
234	-2.55%
208	-2.72%
182	-2.94%
156	-3.23%
130	-3.61%
104	-4.19%
78	-5.06%

If Red Dog were perfectly linear, it would not matter if there were six or a million decks left in the shoe (the expectation for a million deck shoe is actually -1.09%). As it were, the estimation will only be accurate for some particular T. Blackjack theorists would set this value of T midway through the dealt portion of the shoe and count the inaccuracies before and after that point as insignificant. We will find the effects of removal at the top of the shoe and then adjust for nonlinearity with a method peculiar to unbalanced counts and unknown in the theory of blackjack. Now let us determine how much the removal of one card of each denomination from the shoe affects the expectation and then correlate some integer values to these effects which will be easier to count.

card denomination	effect of removal	point count
2	+ 10.873 * 10^-4	+ 1
3	+ 2.805 "	+ 1
4	- 0.291 "	+ 1
5	- 2.998 "	0
6	- 4.158 "	- 1
7	- 4.158 "	- 1
8	- 4.158 "	- 1
9	- 4.158 "	- 1
10	- 4.158 "	- 1
11	- 2.998 "	0
12	- 0.291 "	+ 1
13	+ 2.805 "	+ 1
14	+ 10.873 "	+ 1

Why is the count symetrical and positive on the edges? The symetry is not hard to explain since there is nothing in the rules at all to distinguish between high and low cards. What the rules do distinguish is the difference between the initial cards. Middle cards simply cannot be a part of initial cards which are far apart nor can the draw of an edge card be a winner for very many sets of initial cards. Since an acey-deucey is good for the player and aces and deuces are essential for this set of initial cards to appear, it might seem that the player would miss the disappearance of edge cards such as aces and deuces and assign the observation of their having been dealt a negative number. On reflection, however, it is obvious that edge cards are primarily bad for the player for they do not often appear as initial cards and if dealt as the third card they will be a loser almost every time. Aces and deuces in particular will kill any non-pair hand if dealt as the third card.

A count composed only of the integers -1, 0, and +1 is called a "plus-minus" count and is easier to use than more complicated counts. This is because the player can simply visualize a number line with a pointer moving up and down it upon recognition of the cards rather than having to associate an integer with each denomination and perform addition or subtraction in his head. Thus the player need only remember the direction to move the pointer when the shoe is depleted of a certain type of card (middle or edge).

How well do these integers correlate with the actual effects of removal? Numerically, the correlation coefficient may be gotten by the sum of the products of the actual effect and their estimates (the "inner product") devided by the square root of the sum of the squares of the effects multiplied by the sum of the squares of their estimates (the entire denominator is under the root sign). For the above point count, the correlation is 76.19%. It is important to realize that correlation is a measure of how well the integers fit the effects of removing cards and not of how well these effects fit reality, that being determined by linearity.

The observant reader will have noticed that while the effects of removal add up almost exactly to zero (the difference is due to calculator error), the values of the point count add up to +1. Thus, because there are 24 sets of 2...14 in a six deck shoe, the count at the end of the shoe will be 24. This is what is meant by the term "unbalanced count". First, imagine a balanced count in a game like blackjack which has a basic strategy expectation of about zero. The player wishes to place minimum, waiting bets whenever his expectation is negative and then place extreme bets when the expectation is positive. Whenever the running count is positive, then so is the advantage and whenever it's negative, then the advantage is also negative. If that's all the player needs the count for, then he doesn't have to normalize it to the number of remaining cards in the shoe. If he wants to bet in proportion with the expectation then he'll have to multiply the running count (which is only accurate where T equals the full shoe minus one) by 51/T (assuming a 52 card pack; 311/T for Red Dog) to get his true count. This is because the removal of individual cards has a greater effect on advantage when there are few cards left in the shoe than when there are many. This can be represented graphically as such:

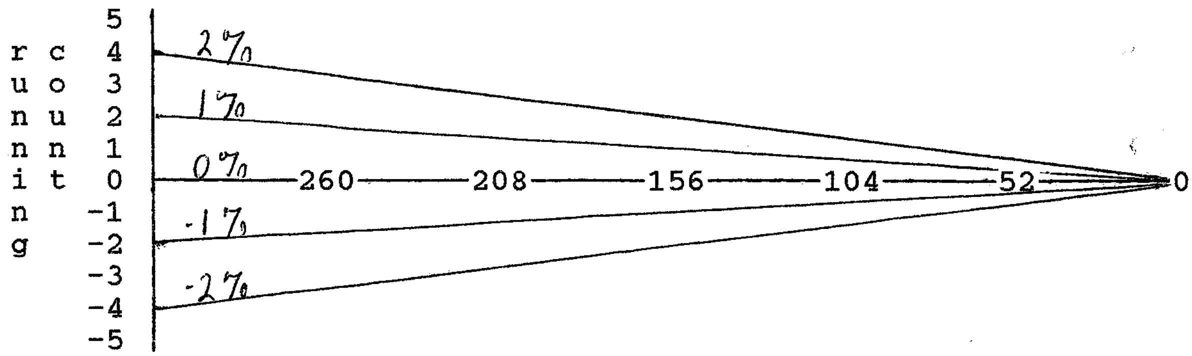

The horizontal axis is the number of cards remaining in the shoe. On the left, where most of the shoe remains to be dealt, it takes a great running count to reach a certain advantage represented by a diagonal line. Later on it takes less of a running count to reach the same advantage. If this same player wants only to make minimum and maximum bets as first hypothesized, but that he wants to reach for the black chips, not at zero, but at some positive expectation like 1%, perhaps because he distrusts the count due to nonlinearity, then he uses an unbalanced count such that the number it adds up to at the end of the shoe (called the "pivot") represents his desired expectation (1%) at the beginning of the shoe. The graph will be centered around 1% instead of 0% and the player need not normalize to find the true count for his purposes. Graphically, it looks like this:

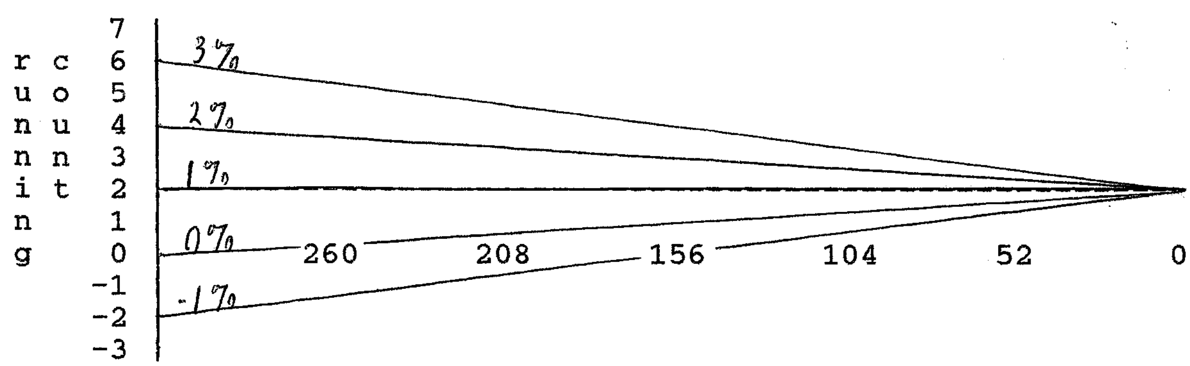

Back to Red Dog, which has a negative expectation for basic strategy play, one might think that if he set the pivot at zero expectation he would know when to plunge in with black chips without having to normalize running count to true count (most of the error in counting cards is due to this normalizing). This would work if the amount of deflection necessary to bring the advantage up to zero remained constant. Unfortunately, this quantity increases as the shoe is depleted due to nonlinearity (see the chart of advantage at the several T's). To approximate when the shoe has turned positive, we need a function which is inversely related to T. If the pivot of an unbalanced count is set less than 2.19% then the running count which represents a deflection in advantage of 2.19% at T = 311 will represent greater deflections at smaller T's. If the pivot represents a deflection of X% less than 2.19% then the diagonal line which crosses the point (T,2.19) will predict a percentage deflection of expectation according to the following formula:

$$2.19 - X + 311 * \arctan(X/T)$$

The optimal value for X is determined by trial and error. Letting X = 1%, a running count which represents 2.19% at T = 311 will represent the following deflections of expectation at the several T's:

remaining cards	expectation
312	-2.19%
286	-2.28%
260	-2.39%
234	-2.52%
208	-2.69%
182	-2.90%
156	-3.18%
130	-3.58%
104	-4.18%
78	-5.18%

This has a 99.93% correlation with the actual values. It was to achieve this correlation that the integers of the point count were chosen as they were. Because the average absolute deflection in advantage for removing one card from the full shoe is 0.0421% and the average absolute deflection in the given point count is 0.846 and by deviding the latter into the former we find that each unit of true count represents an average deflection of 0.0498% in the advantage, then 24 units (the pivot) of true count represent 1.19% (vis., 2.19% - 1%) and 44 units represent 2.19%. Thus, whenever the running count equals or exceeds 44, the shoe has a positive expectation for the player regardless of the number of remaining cards in the shoe, and he should bet as much as possible. Actually, it's unusual for the running count to reach 44; most of the time the game has a negative expectation for the player. He can cut his loses considerably, however, by leaving the table whenever the true count has not reached, say, 19 after about a fifth of the shoe is dealt.

DOUBLING THE BET

There are strange things done in the Vegas sun
By the men who moil for gold;
The Nevada trails have their secret tales
That would make your blood run cold;
The casino lights have seen queer sights;
But the queerest they ever did see
Was the night of the show at Lake Tahoe
That I doubled on five and three.

Apologies to Robert Service

Determining the size of the initial wager is not the player's only decision. While hardly inundated by the choices he must make, there is one more decision to consider: when to double the bet. With a rectangular distribution of the remaining cards, the expectation after the initial cards have been dealt are as follows for the several possible differences between those initial cards (the difference is one more than the spread the little "red dog" denotes at the table):

difference	expectation
2	-53.90
3	-22.73
4	-30.52
5	-38.31
6	-22.73
7	- 7.14
8	8.44
9	24.03
10	39.61
11	55.19
12	70.78

It is from this chart that we derived the basic tactic for Red Dog: Double whenever the difference between the initial cards is 8 or more. Now, for differences of 7 and 8, the absolute value of the advantage is relatively small and might be easily swayed by changes in the composition of the shoe. For large point counts, one will want to double more often and for small point counts, one will want to double less often. However, since the player will quit the game if the true count does not equal at least 19, it is not necessary to be concerned with when not to double. To see how well the point count correlates to the change in expectation on drawing to initial cards with a difference of 7 caused by the removal of individual cards of each denomination, we must repeat the calculation we used to find the point count except in regards to this more limited situation. However, just doing it on initial cards with a difference of 7 is not limited enough for the correlation will be more accurate when it is middle cards that are needed to win than when it is transition cards (3,4,12,13) that are needed. The following chart gives the several effects of removal (each effect is multiplied by 10^-4 but that is omitted to condense the printing) for three of the possible sets of initial cards.

denomination	5,12	6,13	7,14	count
2	30.247	30.247	30.247	+ 1
3	30.247	30.247	30.247	+ 1
4	30.247	30.247	30.247	+ 1
5	30.247	30.247	30.247	0
6	-34.900	30.247	30.247	- 1
7	-34.900	-34.900	30.247	- 1
8	-34.900	-34.900	-34.900	- 1
9	-34.900	-34.900	-34.900	- 1
10	-34.900	-34.900	-34.900	- 1
11	-34.900	-34.900	-34.900	0
12	30.247	-34.900	-34.900	+ 1
13	30.247	30.247	-34.900	+ 1
14	30.247	30.247	30.247	+ 1

The other three possible sets of initial cards will have corresponding correlations as they are equally far from the middle as these. The correlations (in percentages) are:

2,9	3,10	4,11	5,12	6,13	7,14
24.60	58.27	91.94	91.94	58.27	24.60

This is quite a good correlation for 4,11 and 5,12, a marginal one for 3,10 and 6,13, and a next to useless one for 2,9 and 7,14. This lack of accuracy speaks of the need for a multi-parameter count. Hence, we will have three counts, one for when the player needs low cards to win (the L-count), one for when he needs middle cards (the M-count), and one for when he needs high cards (the H-count). Considering the complicated counts people are using in blackjack where the cards go by much quicker, it should not be inconceivable for someone to keep three counts in his head. However, if brains were computers, the human one would not be considered a very efficient parallel processor and players might want to work in teams of three, each person keeping one of the counts. They are:

card denomination	L-count	M-count	H-count
2	+ 1	+ 1	+ 1
3	- 1	+ 1	+ 1
4	- 1	+ 1	+ 1
5	- 1	0	+ 1
6	- 1	- 1	+ 1
7	- 1	- 1	0
8	0	- 1	0
9	0	- 1	- 1
10	+ 1	- 1	- 1
11	+ 1	0	- 1
12	+ 1	+ 1	- 1
13	+ 1	+ 1	- 1
14	+ 1	+ 1	+ 1

The six possible sets of initial cards with differences of 7 have the following correlations:

L-count		M-count		H-count	
2,9	3,10	4,11	5,12	6,13	7,14
91.94	75.10	91.94	91.94	75.10	91.94

These are much better correlations than before and if the player wants to improve the 75% on 3,10 or 6,13 for the L-count or the H-count when the count in question only slightly exceeds the index, he can also consider whether the M-count exceeds the other count before doubling. What index represents the 7.14% deflection in advantage necessary to make doubling on a difference of 7 a good bet? Because the average absolute deflection in advantage for removing one card from the full shoe is 0.324% and the average absolute deflection in the point count is 0.846 and by deviding the latter into the former we find that each unit of true count represents an average deflection of 0.383% in the advantage when drawing to initial cards with a difference of 7, then 19 of these units represent a deflection of slightly more than the 7.14% needed. Recalling the graph with the diagonal lines, a running count of 19 at the beginning of the shoe represents a true count of 19, as does a running count of 20 a fifth of the way through the shoe, a running count of 21 two fifths of the way through the shoe and so on. Thus the index for doubling on differences of 7 is 19 for the first fifth, 20 for the second fifth, 21 for the third fifth, and 22 for the fourth fifth (the last fifth is not dealt).

This is the beauty of using unbalanced counts: Very often the pivot can be set exactly at the index so that no normalization of running count to true count is needed. When it can't, it is close enough that the player only needs to estimate T, the remaining cards, within a fifth of the full shoe (62 cards) to perform this adjustment. Because unbalanced counts are rarely understood, balanced counts are often used (in blackjack - no other Red Dog system exists at this time) when the index is far from the balanced count's pivot of zero requiring the player to estimate T within a quarter of a deck (13 cards) to normalize running count to true count. As it is difficult for the player to estimate T with this accuracy, errors occur. Furthermore, the technique of using a constant running count index other than the pivot (as was done to predict initial advantage) to compensate for nonlinearity is a method unknown to blackjack theorists.

As for the other decisions concerning doubling, a difference of 3 and 6 require a deflection of 22.73%. Using calculations similar to those used for determining when to double on a difference of 7, we find that the true count indexes for differences of 3 and 6 are 44 and 62 respectively (average absolute deflection in advantage is inversely related to the difference between initial cards). 62 is too much to hope for but 44 is marginally possible. Because 44 is quite far from the pivot, the player will have to estimate T within a tenth of the full shoe (32 cards) to adjust running count to true count. The rule is: For each tenth of the shoe that has been dealt, two units must be subtracted from the index which running count must exceed. Thus for the first tenth of the shoe the index is 44, for the second tenth the index is 42, for the third tenth the index is 40, and so on. The player should use the L-count for determining when to double on 2,5 through 4,7, the M-count for 5,9 through 8,11, and the H-count for 9,12 through 11,14. All of these estimations will have a correlation of 50.10%.

CONCLUSION

Start at the beginning,
~~And finish at the end.~~
go until you reach the end,
then stop.
Lewis Carroll

Probability and statistics are often lumped together indiscriminantly but are actually distinct subjects related to each other very much as tactics and strategy are related in a military campaign. Probability tells one how to maximize the expectation of winning each play and statistics tell one how to win as much money as possible in some time period given a limited bankroll and a certain propensity to take risks. Up to this point, ~~our~~ THIS paper has discussed only the tactics of Red Dog. While any game in which it is not impossible for the player to gain an advantage can be beaten by betting heavily when an advantage is perceived, Red Dog is rarely advantageous to the player. The win rate in Red Dog relative to the bankroll needed to support the required betting variances compares unfavorably with the return on a savings deposit. Strategy then is simply not to play. When it's relevant, the discussion of strategy usually takes up a large part of the theory. As it were, our paper is complete.