

Distinguishing Gamblers from Investors at the Blackjack Table

David Wolfe

April 1, 2002

Abstract

We assess a player's long term expected winnings or losses at the game of blackjack without knowing the strategy employed. We do this by comparing the player's moves to those of a baseline strategy with known expected winnings. This allows an accurate estimate of the player's expectation to be found hundreds of times faster than the naive approach of using the average winning observed as an estimate.

1 Background

A skillful blackjack player, one who counts cards, maintains some information about the distribution of cards remaining in the deck at all times. The player adjusts both wager and play decisions based on this *count* information. Depending on the rules used by a particular casino, the skillful player may have a slight edge over the casino. Without knowing exactly what the player is counting, we would like to assess a blackjack player's playing skill.

There are two potential benefits from this research. First and foremost, this is related to the much harder problem of assessing the quality of decisions people make under uncertainty. For example, a pension fund manager tries to distinguish a good portfolio manager from a lucky one. Second, there are many gamblers who deceive themselves into thinking they are able to play blackjack well enough to beat the casino. In fact, casino blackjack revenues skyrocketed [Tho96] after Thorp published his landmark book, *Beat the Dealer* [Tho62], which explained how to effectively count cards.¹ Players who discover their true skill (usually very poor) will hopefully be deterred from gambling. (As an aside, the author suspects this sort of research is conducted by casinos who, due to their financial interests, are disinclined to publish results in the area.)

¹The actual history is a bit more complex: Casinos first response to Thorp's book was to panic and change the rules to remove any player advantage, at which point angered players stopped playing. Casinos then changed the rules back, and players (mostly poor ones) thinking they could exploit their *edge* returned to the tables in greater numbers than before.

1.1 Blackjack rules

Blackjack rules vary from casino to casino. We'll summarize one possible set of rules, and mention only some common variations.

Blackjack is played with a shuffled *shoe* consisting of one to eight standard 52-card decks. Aces are worth 1 or 11, face cards (and 10s) are worth 10, and all other cards are worth their face value. The player is playing only against the dealer (or house), and the dealer's strategy is fixed by casino rules.² Hence, if one discounts the second-order effects of other players playing at the same table and drawing from the same shoe, blackjack is a one player game. The goal is to come closer than the dealer to 21 without going over.

First the player places a wager. She is then dealt two cards face up, and the dealer is dealt one card up and one down. If the player (respectively, dealer) is dealt 21 or *blackjack* she immediately wins her wager times 1.5 (respectively, loses her wager). If both are dealt *blackjack*, the hand is a *push* or tie and the player keeps her wager. If neither has blackjack, play continues; achieving 21 later in play no longer pays 3:2.

The player then has three (or sometimes more) options:³

- *Stand* to accept no more cards.
- *Hit* to receive an additional card. The option to hit or stand is then repeated until she chooses to stand.
- *Double down*, to double her wager and receive exactly one additional card face up.
- *Split* if the two cards are of the same value. The player doubles her wager and the two cards are separated into two hands. Each hand receives one card. After splitting, doubling down is no longer permitted, blackjack no longer pays 3 : 2, but resplitting is usually allowed. Aces are often treated specially, with only one card dealt to each. Each hand split independently plays against the dealer.

If the player goes over 21 or *busts*, the dealer immediately wins her wager, giving the dealer his sole advantage.

The dealer then turns up his second card and draws cards until his total is 17 or higher.⁴

The player with the higher total then wins the amount of the wager, where a tie is a *push*.

²In this paper, the player is female and the dealer is male.

³There are many variations in casino rules. It may be that only some hands can be *doubled down* such as totals of 9, 10 or 11. Splitting may only be allowed up to four hands. The player may have an additional option to *surrender* half her bet, ending the hand. She may also be able to *purchase insurance* for half her wager to protect against the dealer's blackjack.

⁴In some casinos the dealer also draws a card when the total is a *soft 17*, i.e., a total of 17 with an ace counting as 11.

1.2 Expected winnings under perfect play

There are a number of approaches listed in the literature for estimating a player's perfect strategy (from a given deck situation) and her expected winnings under that perfect strategy. (Here, winnings could be negative if she loses money.)

The first approach was given in Baldwin, Cantey, Maisel and McDermott in 1956 [BCMM56] and was honed by others [Hea75] [BH93] [Gor73]. The deck is assumed to be infinite and each card appears with probability $1/13$. In this case, dynamic programming can be employed to quickly and accurately calculate both perfect play and expected winnings.

A variant of this approach (and the one we use) is to assume the *remaining* deck is infinite, with the probability of drawing a card given by the fraction of its appearance in the current deck. Before the initial deal, this is identical to the first approach.

Thirdly, it is nearly possible to calculate expected winnings exactly by exploiting the fact that there are only, for instance, 2527 possible player hands (that haven't busted) when using a four deck shoe [MBG75] [Gri99]. This permits exact calculations when splitting is forbidden, and for all practical purposes is exact.

2 Assessing a player's skill

We wish to assess a player's skill by observing her play over a short period of time. Initially, we will also assume she plays with unit wagers. For simplicity of discussion, assume she is playing a fixed strategy. By *fixed strategy*, we mean that the player's wagers and decisions are completely determined by the cards she has seen since the last shuffle. While the strategy could easily be extended to be random, we assume it is independent of external factors such as her bankroll and how many complementary drinks she's had.

2.1 The challenge

One approach for assessing a player's skill would be to watch the player play for a while and report her total winnings divided by the number of hands she played. If we assume unit wagers, a single hand's winnings is a random variable, with variance typically exceeding 1.3. Her expected winnings, if she is good, are typically between $-.01$ and $.01$ assuming reasonable favorable rules and decent skill [Gri99] [Tho62]. Therefore, a reasonable goal is to be 95% confident that a measured mean is within, say, $.002$ of the true mean. This goal demands she play at least $1.3 \cdot (1.96 / .002)^2 \approx 1,250,000$ hands.⁵

In short, the challenge of assessing a player's skill is that her risk is so much higher than her expectation. The key idea to accelerating skill assessment is to consider the difference in expected winnings for a player's actual play (in each situation she is exposed to) and a known baseline strategy. This difference

⁵95% of a normal distribution is between -1.96σ and 1.96σ of the mean.

has much smaller variance than that of a single hand, and we can recover the measured mean by adding this difference to the expectation of the baseline.

We'll formalize this plan in the next two subsections.

2.2 Estimating skill with unit wagers

We first address the case when a player always wagers one unit. This section will be a bit informal in the hopes of building intuition.

The player's situation before any cards are dealt are determined by the order of the remaining cards in the deck. We'll separate this sample space into two components: C is the cards remaining (i.e., number of aces, 2s, 3s, ..., ten-cards), and O is the *order* of the cards remaining.

Again, a player's strategy S is a mapping of player situations to actions. The player's strategy cannot depend upon the order of unseen cards, though it may depend upon the distribution of unseen cards since that is completely determined by cards seen since the last shuffle.

We'll construct a reasonably good *baseline strategy*, B , to which the player's play will be compared. While the actual details of the baseline strategy are irrelevant to the discussion, the measurement will be most accurate if B is close to player's strategy S . In particular, since we expect the player to be pretty *good* (for very poor play is easy to diagnose), strategy B should implement pretty good — if not perfect — play.

Lastly, W_S , a random variable over sample space $C \times D$, is the winnings of the player playing strategy S . In summary, we have defined

C	$\stackrel{\text{def}}{=}$	Sample space of cards remaining
O	$\stackrel{\text{def}}{=}$	Sample space specifying ordering of the deck
W_S	$\stackrel{\text{def}}{=}$	Winnings playing strategy S
B	$\stackrel{\text{def}}{=}$	Any automated baseline strategy

Then, the expected winnings of the player, $\mathbf{E}_C \{ \mathbf{E}_O \{ W_S \} \}$, can be rewritten as follows:

$$\begin{aligned} \mathbf{E}_C \{ \mathbf{E}_O \{ W_S \} \} &= \mathbf{E}_C \{ \mathbf{E}_O \{ W_B \} \} \\ &\quad + \mathbf{E}_C \{ \mathbf{E}_O \{ W_S \} - \mathbf{E}_O \{ W_B \} \} \end{aligned}$$

Again, W_S typically has variance far exceeding its expectation. Therefore a large number of sample hands are required to accurately estimate $\mathbf{E}_C \{ \mathbf{E}_O \{ W_S \} \}$ if we estimate it by simply dividing actual winnings by number of hands played.

On the other hand, for any reasonable choice of baseline strategy B , $\mathbf{E}_C \{ \mathbf{E}_O \{ W_B \} \}$ can be calculated almost exactly or estimated by computer simulation. Since $\mathbf{E}_C \{ \mathbf{E}_O \{ W_S \} - \mathbf{E}_O \{ W_B \} \}$ typically has much smaller variance, the right hand side can be estimated by observing the player for far fewer hands.

2.3 Estimating skill with varied wagers

In a casino, a player can vary her wager between some minimum and maximum quantity. For instance, she may be able to bet any integer amount between \$2 and \$500.⁶

To estimate the player's skill, we'll watch her play. After she wagers, we'll estimate how much she *would* make on average if she were playing a known baseline strategy. When she makes a decision (whether to hit, stand, double, or split), we'll determine whether her decision differed from the baseline strategy. If it did, we'll credit her the difference between her expected winnings according to the decision she made (assuming she plays baseline for the rest of the hand), and the expected winnings under baseline.

It is plausible that two players are (in the long run) likely to encounter very nearly the same *situations* no matter their respective strategies. We'll make the simplifying assumption that the situations a player is exposed to are independent of her strategy. In particular, we are assuming the following:

- A player cannot walk away from the table if the odds have gone against her (though she can bet the minimum).
- The distribution of cards remaining are independent of her hit and stand choices on previous hands. While it's conceivable that a player might choose to make a suboptimal play on the current hand (for instance, by hitting rather than standing) in order to gain more information on future hands, it's hard to believe it could make much difference.
- The number of times a player has the choice to hit or stand with, say, 16 versus dealer's 10 is independent of her strategy. This is false, for a poor player might never hit with 12 versus dealer's 10, and thereby lose opportunities to get 16 versus dealer's 10. For a good player, this assumption is (hopefully) of little consequence.

Now a typical *situation* consists of the cards remaining, the current wager (if it's been made), the cards in front of the player and dealer, and the order of cards remaining in the deck. Focusing on the cards, let's separate the sample space consisting of the cards remaining in the deck and the cards in front of the players. Let C be the distribution of cards remaining in the deck before a hand is dealt. Let D be the initial deal and O the order of the deck after the deal. A strategy S now includes both how the cards should be played (when to hit, stand, split, or double) and how much to wager as a function of the cards remaining in the deck.

We'll define the random variable $W_{S_0 S_1 S_2 \dots}$ as the winnings of a player when wagering according to strategy S_0 , making her first decision of a hand (to split,

⁶It should be noted that if one is to maximize expected winnings, the only two reasonable wagers are the minimum or the maximum. In simulations, we considered much narrower ranges because, (1) wide bet variation results in far too high risk for most players, (2) casinos tend to evict good players with such high bet variation, and (3) rare events warranting large wagers can contribute more strongly to long term expected winnings, making measurement more difficult.

stand, double, or split) by strategy S_1 , her second decision according to S_2 , and so forth. Again, B is a baseline strategy. In summary,

C	$\stackrel{\text{def}}{=}$	Sample space of cards remaining
D	$\stackrel{\text{def}}{=}$	Sample space specifying ordering of initial deals
O	$\stackrel{\text{def}}{=}$	Sample space specifying ordering of the deck after deal
$W_{S_0 S_1 S_2 \dots}$	$\stackrel{\text{def}}{=}$	Winnings using wagering strategy S_0 and playing strategy S_i for the i^{th} decision after the deal.
B	$\stackrel{\text{def}}{=}$	Any automated baseline strategy

The player's long term winnings per hand are then given by $\mathbf{E}_C \{ \mathbf{E}_D \{ \mathbf{E}_O \{ W_{SSSSS\dots} \} \} \}$. We'll estimate this quantity using the equation:

$$\begin{aligned} & \mathbf{E}_C \{ \mathbf{E}_D \{ \mathbf{E}_O \{ W_{SSSSS\dots} \} \} \} \\ &= \mathbf{E}_C \left\{ \begin{array}{l} \mathbf{E}_D \{ \mathbf{E}_O \{ W_{SBBBB\dots} \} \} \\ + \mathbf{E}_D \{ \mathbf{E}_O \{ W_{SSBBB\dots} \} - \mathbf{E}_O \{ W_{SBBBB\dots} \} \} \\ + \mathbf{E}_D \{ \mathbf{E}_O \{ W_{SSSBB\dots} \} - \mathbf{E}_O \{ W_{SSBBB\dots} \} \} \\ \vdots \end{array} \right\} \end{aligned}$$

To estimate a player's average winnings, we'll watch her play. Before each deal, she fixes a wager. Once the wager is fixed, $\mathbf{E}_D \{ \mathbf{E}_O \{ W_{SB} \} \}$ depends only on B and the cards remaining, and can be calculated or estimated through simulation with high precision. Once she decides what to do first (hit, stand, split or double down), we can calculate the difference $\mathbf{E}_D \{ \mathbf{E}_O \{ W_{SSBBB\dots} \} - \mathbf{E}_O \{ W_{SBBB\dots} \} \}$, and so forth. The sum terminates when she no longer is taking cards.

Simulation study

We simulated a six deck shoe with *random penetration* varying from 0 to 5/6, i.e., we assume that before each hand a fraction of the cards between 0 and 5/6 of the cards have been removed. The player knows what cards have been discarded.

In the simulation, the *basic* baseline strategy consists of playing perfectly assuming a standard infinite deck (1/13th of the cards are aces, etc.) This strategy does not depend upon what cards have been seen.

The player's strategy is the *high-low* count system as described by Stanford Wong in [Won94]. Wong attributes the system to Harvey Dubner in 1963. The player's *count* is the number of 10s and Aces which minus the number of twos through sixes which have been seen in the deck. (This is a *balanced count*, meaning it has 0 expected value.) The player's *count per deck* is the ratio of the count and the number of decks not yet seen. This *count per deck* is truncated, i.e., rounded to the integer nearest 0. The player wagers 1 unit at a count per deck which is negative, 2.5 units at 0 or 1, 5 units at 2, 7.5 units at 3, and 10 units at 4 or greater. The reader can refer to [Won94] to see how the

count affects play. Wong predicts the player’s long-term advantage is around $0.014 \pm .002$ units per hand.

Hands Simulated	Average Winnings	Assessed Expectation
500	-0.0810	0.0275
1,000	-0.1145	0.0170
2,000	-0.1318	0.0188
5,000	-0.0742	0.0174
10,000	-0.1204	0.0108
20,000	-0.0925	0.0102
50,000	-0.0331	0.0149
100,000	-0.0077	0.0138
200,000	-0.0022	0.0158
500,000	-0.0049	0.0154
1,000,000	0.0032	
2,000,000	-0.0010	
5,000,000	0.0017	
10,000,000	0.0037	
Variance	20	0.83

Observe that using the techniques in this paper, the predicted long-term winnings converge much faster than the player’s actual average winnings. To explain the high variance, recall that while typical winnings are \pm wager, the actual winnings could be larger or smaller if, say, the hand is blackjack, or has been doubled or split.

3 Further research

In addition to dispensing more formally with some of the assumptions mentioned in this paper, a number of other (perhaps more interesting) questions remain.

Although our technique allows the player’s expectation to be measured hundreds of times faster than the naive technique, it still requires observing upwards of a thousand hands. An experienced player can play about one hundred hands per hour and has practiced for hundreds or thousands of hours to become proficient, so playing a few thousand hands to benchmark her progress is not out of the question. (It should be quicker to assess lack of skill in a poor player.) Further improvement may be possible if we can limit which situations the player is exposed to. Since most decisions a player makes are rote, it should be possible to test the player on only *challenging* hands and more rapidly assess the player’s skill.

In this paper we have not addressed the affect of multiple hands, either played by the same player against one dealer, or by different players sitting at the same table against the dealer.

The most serious blackjack players take into account risk as well as expected winnings. One popular approach is the *Kelly criterion* [Kel56] [Tho] which seeks

to maximize the expected logarithm of the player's bankroll. It's therefore of interest to measure a player's skill who is trying to safely increase her bankroll rather than merely maximize her expected winnings. In addition, it might then be more reasonable to assess a player's skill who is permitted wide bet variation without some of the concerns expressed in the footnote on page 5.

There are a number of assumptions which we are less optimistic of addressing. In particular, the real world just doesn't fit the model. In real casinos, the deck's distribution is not uniformly random and depends greatly on the shuffling methods used. Real players don't use a fixed strategy, but change their strategy both deliberately (for instance, to avoid getting spotted as a card-counter by the casino) and inadvertently (due to distractions, exhaustion, etc.).

Acknowledgements

The author would like to thank Brian Kleinke, whose undergraduate research project confirmed the feasibility of the computer calculations necessary to complete this work, and Max Hailperin, whose helpful advice improved this presentation. The author is also indebted to a blackjack *investor* (who will remain anonymous) for educating the author about card counting issues.

References

- [BCMM56] Roger Baldwin, Wilbert Cantey, Herbert Maisel, and James McDermott. The optimum strategy in blackjack. *Journal of the American Statistical Association*, 51(275):429–439, September 1956.
- [BH93] Art T. Benjamin and Eric Huggins. Optimal strategy with “lucky bucks”. *The Journal of Undergraduate Mathematics and its Applications*, 14(4):309–318, Winter 1993.
- [Gor73] Edward Gordon. Optimal strategy in blackjack. *Claremont Economic Papers*, January 1973.
- [Gri99] Peter A. Griffin. *The Theory of Blackjack*. Huntington Press, Las Vegas, Nevada, 6th edition, 1999.
- [Hea75] David Heath. Algorithms for computations of blackjack strategies. Presented at the Second Conference on Gambling, sponsored by the University of Nevada. Available at <http://www.bjmath.com>, 1975.
- [Kel56] J. L. Kelly, Jr. A new interpretation of information rate. *IRE Transactions on Information Theory*, 2(3):185–189, September 1956.
- [MBG75] A. R. Manson, A. J. Barr, and J. H. Goodnight. Optimum zero-memory strategy and exact probabilities for 4-deck blackjack. *The American Statistician*, 29(2):84–88, May 1975.

- [Tho] Edward O. Thorp. The Kelly criterion in blackjack, sports betting, and the stock market. Paper presented at The 10th International Conference on Gambling and Risk Taking, June 1997. Revised May, 1998 in <http://www.bjmath.com/bjmath/thorp/paper.htm>.
- [Tho62] Edward O. Thorp. *Beat the Dealer*. Random House, 1962.
- [Tho96] Walter Thomason. *The Ultimate Blackjack Book: Basic Strategies, Money Management, and More*. Carol Publishing Group, Secaucus, New Jersey, 1996.
- [Won94] Stanford Wong. *Professional Blackjack*. Pi Yee Press, 7910 Evanhoe, PMB #34, La Jolla, CA 92037-4511, 1994.