The Rational Gambler

Sahand Rabbani

Introduction

In the United States alone, the gaming industry generates some tens of billions of dollars of gross gambling revenue per year.[1] This money is at the expense of the gamblers who attend the casinos and racetracks and other gambling venues. While some extraordinary gamblers may use rigorous analysis and card counting to overcome the house advantage, the average gambler simply executes a passive strategy that subjects him to the negative expectation of the game. This raises the question, Why does a rational player ever participate in such games?

Some might argue that the average gambler is deluded, believing in his superstitions of finger crossing and dice blowing, but this assumption prematurely terminates the economic discussion of gambling and a gambler's incentives. Though some gamblers might believe that their expected payoff from gambling by their superstitions is positive, this is probably not the case for our average gambler. Rather, it is reasonable to assume that the average gambler accepts the negative expectation of the game for the thrill that it provides: the variance of the payoff, the chance of a nonzero payoff, or some other characteristic of the payoff's distribution other than its mean. That is, the negative expectation of the game is the price of the product that the casino offers. That product is thrill. Though this characterization may not seem controversial, I will use the framework along with reasonable assumptions to reach conclusions, some of which run contrary to commonly held beliefs about optimal gambling.

In this article, I first consider the problem of the rational gambler faced with placing a single bet. I define the gambler's problem with some restrictive assumptions. I relax these assumptions in the following section where I consider the gambler who places a series of bets in an allotted time. Based on the insights from this analysis, I draw conclusions about optimal gambling. Finally, I show that the rational gambler's problem and solution is analogous to that of the investor in Markowitz portfolio optimization though with some key differences. I conclude by discussing some of the limitations of this theory.

A Single Bet

For the sake of the analysis in this section, let us assume a model in which the player derives utility only from the variance in the payoff of the game. That is, the gambler is not concerned with the larger moments of the payoff's distribution. In order to advance with this analysis, we must first establish a somewhat rigorous framework in which we can consider bets, strategies, and payoffs.

Let us define a *strategy* as a complete set of actions for all possible states of a game. For example, a very simple strategy is to place a one-unit wager on the 0 on a roulette table and then to cease after the outcome of that wager is determined. In blackjack, a more sophisticated strategy is to place a one-unit wager on one game while defining the complete set of instructions of whether to stand, hit, double down, or split, based on the dealer's exposed card and the sum of one's own hand.

Let us define a *bet* as the act of applying the strategy once. In our two examples, a bet would last one spin of the roulette wheel or one hand of blackjack. Finally, let us define the payoff function π, which maps a strategy S into a stochastic payoff Π. The payoff function is determined by the rules of the game and the inherent probabilities of the underlying random events; e.g., dice rolls, wheel spins, and card drawings.

[1] "Gaming Revenue: 10-Year Trends." American Gaming Association [http://www.americangaming.org/].

Finally, let us establish one intuitive but critical requirement for a strategy. The payoff of a strategy must be linearly scalable. That is, for S to be a strategy, there must exist a strategy S' such that if the payoff of S is $\pi(S) \sim f(x)$, the payoff of S' is $a\pi(S) \sim \frac{1}{a} f\left(\frac{x}{a}\right)$ for any non-negative a. In most applications, this condition is trivial. Strategy S' exists and is the same as strategy S except that whenever S calls for a wager of x, S' calls for a wager of ax.

Having established this framework, we proceed to the problem of the rational gambler. Formally, the rational gambler attempts to maximize the dollar variance of his payoff for a given expected loss by selecting a strategy:

$$\max_S \operatorname{Var}[\pi(S)] \text{ subject to } \operatorname{E}[\pi(S)] = e.$$

Yet a more palatable formulation, yielding the same result, is a maximization of expected value (minimization of expected loss) subject to a fixed a variance:

$$\max_S \operatorname{E}[\pi(S)] \text{ subject to } \operatorname{Var}[\pi(S)] = \sigma^2.$$

It becomes clear, then, that the optimal strategy is some scaling of the strategy with the highest standard-deviation-to-expected-loss ratio. With this in mind, we define the "thrill factor" of a game with payoff Π as the ratio of its dollar standard deviation per bet to its expected dollar loss per bet: $-\frac{\operatorname{SD}[\Pi]}{\operatorname{E}[\Pi]}$. The thrill factor is a positive number.

Consider the optimal strategy with highest thrill factor t and any other strategy $\hat{S}$ with lower thrill factor $\hat{t} < t$. The scaling of the strategy $\hat{S}$ that achieves the expected loss $l = -e \geq 0$ has a variance of $l\hat{t}$. The scaling of the strategy S that achieves this same expected loss, however, has a variance of $lt > l\hat{t}$. Thus, the strategy with the highest thrill factor is the optimal strategy.

In the analogous framing, a scaling of S that achieves standard deviation σ has expected loss σ/t, whereas a scaling of $\hat{S}$ that achieves standard deviation σ has higher expected loss $\sigma/\hat{t} > \sigma/t$. Thus, a game with a higher thrill factor strictly dominates a game with a lower thrill factor, as any variance achieved in the dominated game can be achieved in the dominant game with lower expected loss.

Once the optimal game is selected, the gambler faces the additional problem of which scaling of the strategy to accept based on his degree of risk-loving and subject to his limits on expected loss. To compute the desired scaling of the optimal strategy requires further assumptions of the gambler's utility and is beyond the scope of this article.

Figure 1 compares the thrill factors for several craps and roulette strategies. These strategies are depicted in expected-loss–standard-deviation space, where the slope of the line joining the origin and the strategy's point is the thrill factor of that strategy. Any point, or combination of expected loss and standard deviation, along that line can be achieved by a simple linear scaling of the strategy.

In our analysis, the n-number roulette strategy is the one-time one-unit wager that any of n numbers will be achieved in a single spin. The m-odds craps strategy is a one-time one-unit wager on the pass-line along with a m-unit wager on the free odds bet for the point. The variable odds craps strategy, similarly, is a one-time one-unit wager on the pass-line along with a 3-, 4-, or 5-unit free odds bet on the point depending on the point's value.[2]

It is interesting to note that the roulette straight-up strategy, a bet on a single number, has the highest thrill ratio of any one-time wager roulette strategy at about 109.5. This is higher than the thrill ratio of the no-odds pass-line strategy in craps at 70.71. The single-odds thrill ratio at 133.9, however, exceeds that of all the roulette strategies. Craps strategies with larger free odds bets have corresponding larger thrill factors. This is because the craps free odds bet has zero expectation, contributing only to a higher variance. Some craps tables offer up to 100-times odds, scoring a thrill ratio of 7127.6 and by far dwarfing the thrill of any other passive casino game. In this sense, craps can offer the largest bang for the buck for the risk-loving gambler when free odds are available.

[2]The craps m-odds and variable odds bets are explained and quantified at length in S. Rabbani: "Craps: Computing the Distribution of the Pass-Line and Free Odds Bets."

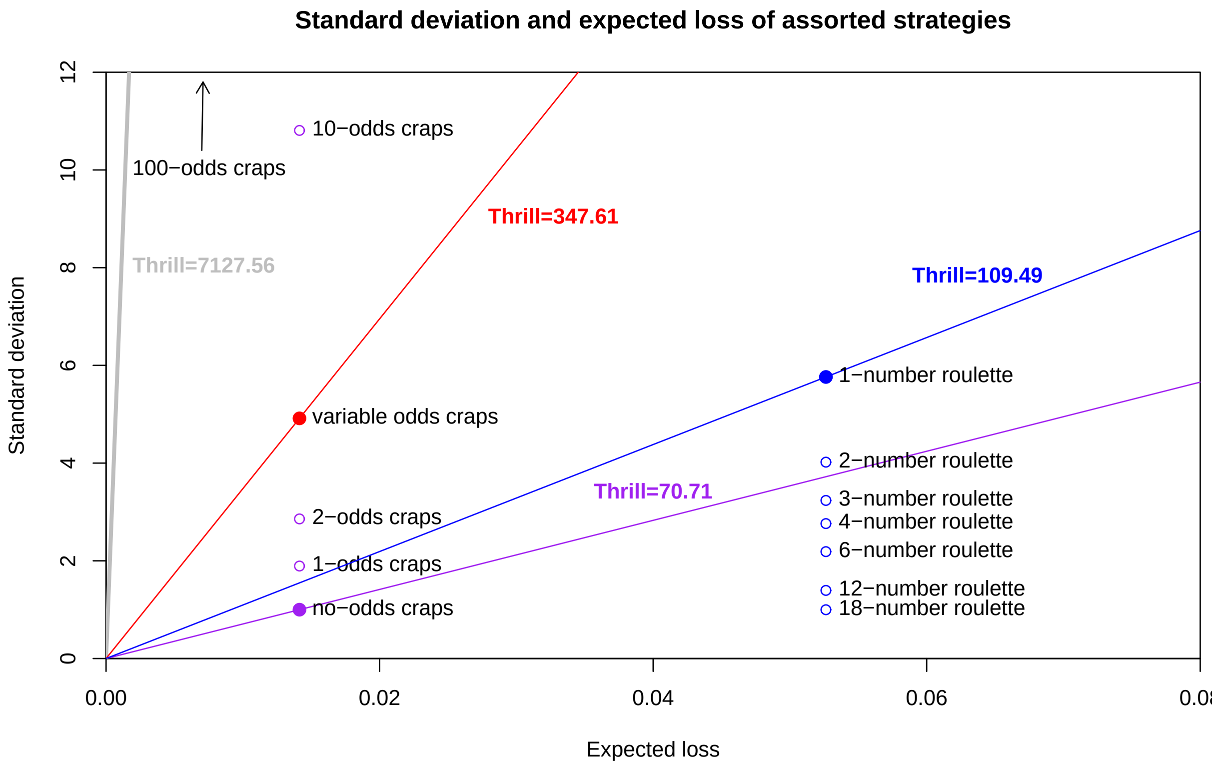

Figure 1: A comparison of the thrill factors of assorted strategies

The results of this analysis ring contrary to traditional gambling proscriptions. Gambling literature largely regards roulette as inferior to craps from the perspective of the gambler, citing the game's larger expected loss as a percentage of the amount wagered. Such a consideration of expected loss, however, is incomplete, for it ignores entirely the aspect of the game that provides the gambler with utility. If the gambler were only concerned with his expected loss, then he would be foolish to gamble at all. Though roulette strategies have an expected loss of 5.26% of the amount wagered while craps has an expected loss of 1.41% of the amount wagered, the straight-up bet in roulette offers much more variance for its expected loss than a simple pass-line bet in craps. By wagering less than one unit in roulette, a gambler can achieve the same variance as in a one-unit pass-line bet in craps but with a lower expected loss.

To remove all abstraction, let us consider a numerical example. A one-unit pass-line bet has a variance of about 0.9999 and an expected loss of about 0.0141. A 0.1735-unit wager on a single number in roulette, on the other hand, has the same variance of about 0.9999 but with an expected loss of only $0.0091 < 0.0141$. Not only does the straight-up roulette strategy beat the pass-line strategy in craps under our restrictive assumptions of the first two moments, it is also superior according to other features of the payoff distribution. The downside of the roulette strategy payoff is a loss of only 0.1735 while the downside of the craps strategy is a loss of 1. Similarly, the upside of the roulette strategy is 6.0730, much greater than the pass-line's 1 upside. On the other hand, the probability of a positive payoff in the roulette strategy is much smaller than in the pass-line strategy.

In the next section, we will consider the gambler who makes several bets. We take advantage of the central limit theorem to relax the assumption that the gambler only cares about the expected loss and standard deviation of a strategy. With sufficiently many bets, the payoffs of all strategies become approximately normally distributed and are only then distinguishable by their expected loss and standard deviation. The idiosyncrasies of single-bet distributions lose their relevance and the model of the rational gambler becomes more general.

Multiple Bets

There are many reasons to believe that a gambler would place several bets over a course of time rather than place a single bet. An initial consideration of the thrill factor, however, may not immediately reveal this.

Consider a gambler whose limit on expected loss is e and who has a preferred strategy S with thrill factor $t = \mathrm{SD}\left[\pi(S)\right] / \mathrm{E}\left[\pi(S)\right]$ and that is scaled to have expected loss e; i.e., $\mathrm{E}\left[\pi(S)\right] = e$. Let us consider another strategy S_n that consists of n bets of the scaled strategy $\frac{1}{n}S$. If the gambler were to space his wagers out over n bets, each with expected loss e/n, he would still face the aggregate expected loss e. However, the standard deviation of his strategy's payoff, rather than being $\mathrm{SD}\left[\pi(S)\right]$, would be

$$\mathrm{SD}\left[\pi(S_n)\right] = \mathrm{SD}\left[\sum_{i=1}^{n} \frac{1}{n}\pi(S)\right] = \frac{1}{\sqrt{n}}\mathrm{SD}\left[\pi(S)\right].$$

The total thrill of this strategy S_n is less than the thrill of the one-wager strategy S:

$$\mathrm{thrill}\left[S_n\right] = \frac{\mathrm{SD}\left[\pi(S_n)\right]}{\mathrm{E}\left[\pi(S_n)\right]} = \frac{(1/\sqrt{n})\,\mathrm{SD}\left[\pi(S)\right]}{\mathrm{E}\left[\pi(S)\right]} = \frac{t}{\sqrt{n}} < t.$$

Under our assumptions, the gambler should always choose S over S_n. However, our assumptions do not fully capture the utility of the gambler. The gambler, presumably, is also entertained while playing the game. Placing several bets over time exposes the gambler to more of this entertainment. A roulette spin may last only a minute, whereas a pass-line bet in craps may last several. The gambler, then, attempts to maximize thrill while filling his allotted time for gambling.

Considering that the rational gambler is playing over the course of time and placing several bets, we can compute his thrill on the total payoff over all of these bets. Let us consider the example of the variable odds craps strategy scaled to a one-unit pass-line wager. The one-time payoff distribution is hardly normal. After successive bets, however, and by result of the central limit theorem, the distribution tends to normality. Figure 2 shows the distribution of the payoff of the variable odds craps strategy applied for periods of varying length. By one hundred bets, the distribution of the payoff is nearly indistinguishable from the normal of corresponding mean and standard deviation, which is superimposed in red.

Because the payoff of any strategy applied sufficiently many times tends to a normal distribution and because this distribution is fully characterized by its mean and standard deviation, we no longer require the original assumption that the rational gambler is concerned only with the first two moments of his payoff. This assumption becomes endogenous to the model. We have, in a sense, traded this assumption with a more realistic one that the gambler places many bets to seek variance at the expense of expected loss while filling an allotted gambling time window.

The analysis suggests, therefore, that a craps strategy consisting of a pass-line bet with some reasonably high free odds is one of the best deals in the casino. Not only is its single-bet thrill factor higher than most games, the duration of the bet spans multiple rolls of dice in expectation. Boasting even higher thrill is a pass-line and free odds bet along with one or more come bets and corresponding odds. Though the thrill of this strategy is difficult to compute analytically, it certainly offers more variance than the pass-line and odds strategy as it strings several such wagers so that their payoffs become positively correlated.

Furthermore, the model of the rational gambler implies results that contradict traditional gambling literature, which typically determines the merits of a casino game based on its expected loss but without regard to its relative variance. The example of the pass-line craps strategy and the straight-up roulette strategy becomes even more compelling when we consider that their several-bet applications yield distributions that are very similar except for their mean and standard deviation. The merits of the straight-up roulette strategy over the simple pass-line craps strategy are now even more evident.

Next, we turn to the similarities between the model of the rational gambler presented here and the model of the rational investor in modern portfolio theory. Some familiarity with the Markowitz portfolio optimization problem is assumed. The next section is not required for a complete understanding of the theory of the rational gambler and is only included as an intriguing analogy.

Markowitz Analogy

Markowitz portfolio optimization deals in part with the efficient allocation of investments among risky assets to achieve high-return low-risk portfolios.[3] In the rational gambler's problem, the desirability of variance is reversed and the risky assets are the gambling strategies.

In portfolio optimization, individual assets and linear combinations of assets are depicted on a standard-deviation–expected-value plane, similar to that in Figure 1 except that the axes are swapped and expected gain is drawn rather than expected loss. The same effect exists in both cases, however, that whenever one point is located to the northwest of another, that point is strictly preferred.

Markowitz's risk-free asset is equivalent to the gambling strategy that calls for no betting, which has the riskless payoff of naught. The scaling of a strategy is merely the allocation of capital between that strategy and the riskless naught bet, the continuum of which is represented by the line passing through the strategy point and the origin. This is analogous to the capital allocation line, which represents allocation between a portfolio and the risk-free asset. The slope of that capital allocation line is the Sharpe ratio of the portfolio the same way that the slope of the line joining the strategy and the origin is the strategy's thrill factor.

The theories of the rational gambler and portfolio optimization diverge, however, when we consider that the maximization and minimization of risk are asymmetric problems. The Markowitz optimal portfolio is a linear combination of several portfolios because the total variance achieved by this combination is lower than that of any individual asset. This is the effect of diversification. In the rational gambler's problem, the optimal strategy is not a blend of multiple strategies. All strategies stand alone and are statistically independent of each other. As such, any combination of strategies will cause an unwanted reduction in total variance. Thus, is it not fruitful to consider blends of strategies.

Perhaps the most compelling difference between the two theories concerns their assumptions. Markowitz portfolio optimization assumes the normality, constancy, and knowledge of the underlying distributions of asset returns, whereas in reality the distribution of asset returns are fat-tailed, variable throughout market conditions, and difficult to estimate. These complications, in fact, are largely responsible for the scale of the financial sector. In the theory of the rational gambler, however, the distributions of strategy payoffs are exact functions of the well-defined underlying distributions of the random processes: dice rolls, wheel spins, and card drawings. As such, the distributions are perfectly known and static. Normality, furthermore, is derived from the central limit theorem as explained in the previous section. Thus, what portfolio optimization assumes about the distribution of asset returns, the theory of the rational gambler finds endogenous to the model. This is a prime source of the theory's strength. The theory is not, of course, without its limitations.

Limitations

While I attempt to be general in developing the theory of the rational gambler, certain simplifying assumptions are limitations of the theory.

The previous analysis assumes that the gambler has an infinite bankroll. In practice, the rational gambler faces a finite bankroll, and optimal behavior in the presence of this constraint may vary from the thrill-maximizing approach considered in this article. For example, the gambler must face the additional problem of weighing the optimal thrill of placing large odds bets in craps against the risk of bankruptcy. This analysis may yield a strategy that does not necessarily call for pursuing the maximum-thrill strategy. Optimal behavior

[3] Markowitz first published his theory of portfolio optimization in his 1952 paper "Portfolio Selection," *The Journal of Finance* **7** (1): 77-91.

in the presence of bankroll constraints relates to the gambler's ruin problem. Quantifying a more general theory with this consideration is beyond the scope of this article.

Further limitations of the theory developed here relate to the scalability of strategies with infinite precision. In reality, a minimum denomination always exists and payoffs that are not an integer multiple of this denomination are often rounded down. Table minimums and maximums also present constraints to the scalability of a strategy.

The theory of the rational gambler, however, may be a useful starting point for the development of a more thorough economic theory. The concept that variance is a product offered by the casino frames the problem in a manner that is more amenable to classical economics, which would otherwise identify gambling as irrational and stymie the formulation of a general theory.

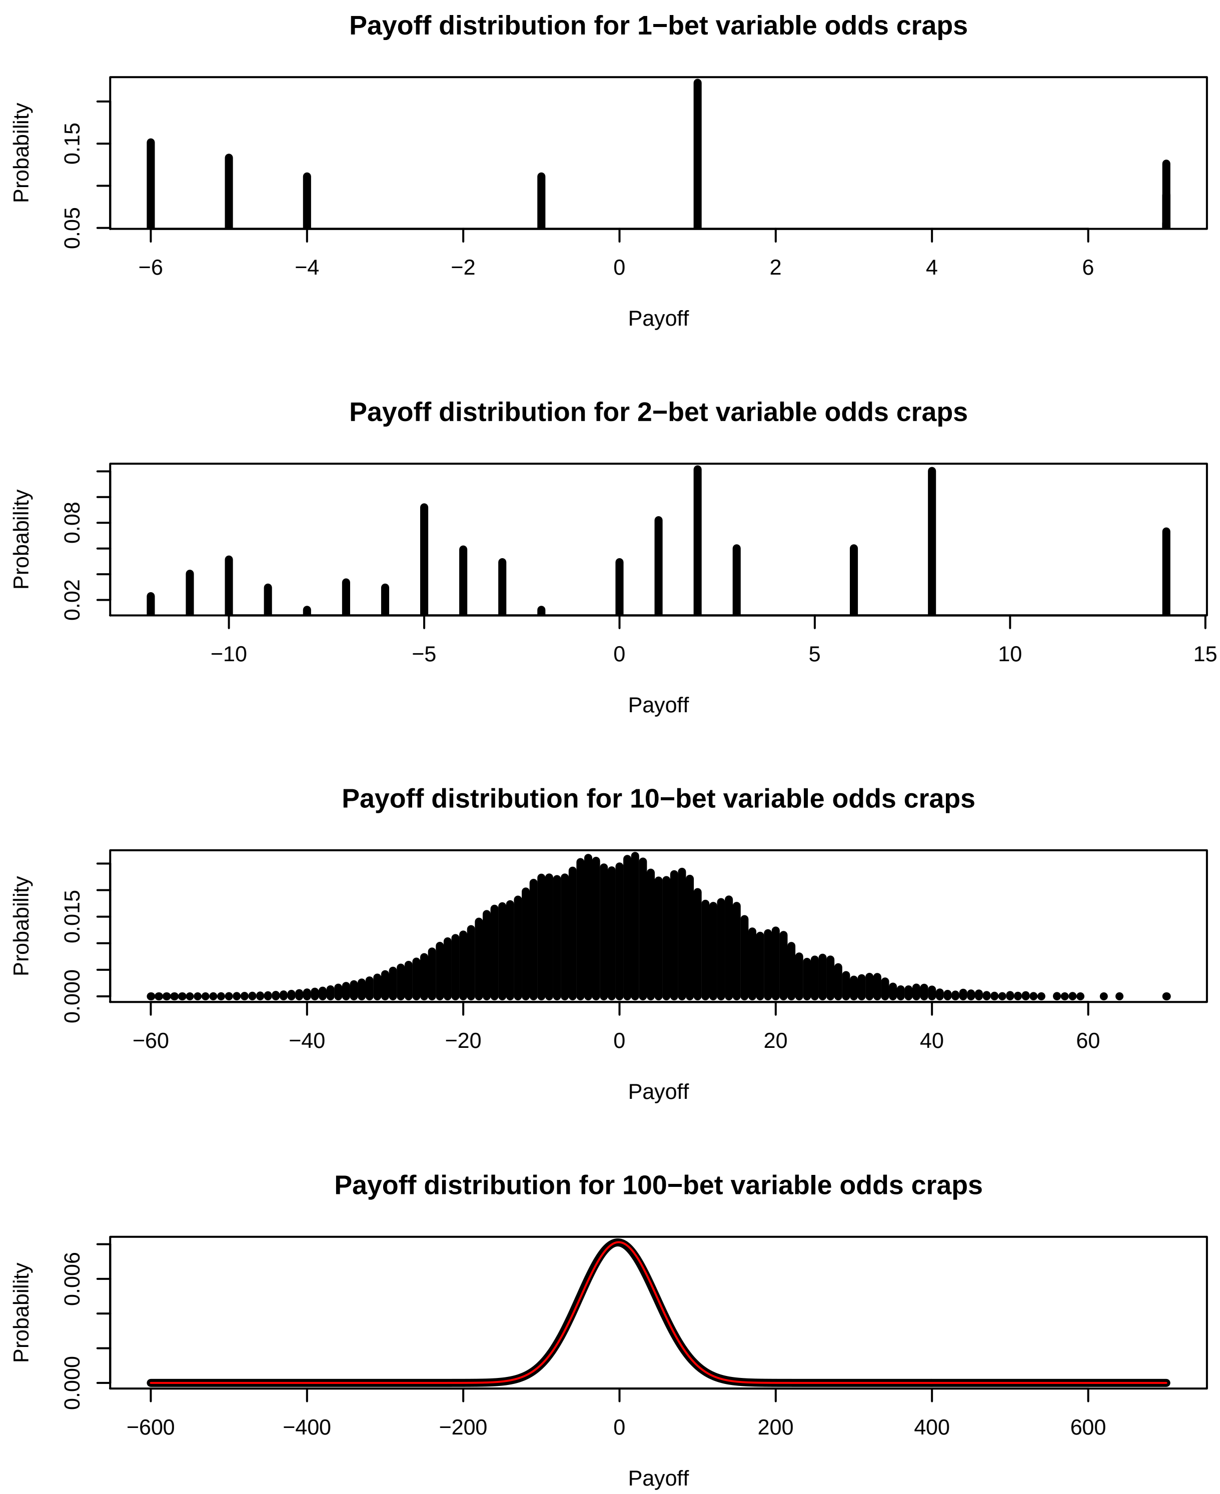

Figure 2: Convergence of multiple-bet payoff to normality