Mathematical Recreations at the Casino

Alexander Nicholson
Department of Computer Science
California Institute of Technology
Pasadena, CA 91125

Malik Magdon-Ismail
Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180

Introduction

This presentation has been inspired by Connolly's article [1] on the Ultimate Strategy for casino gambling. Imagine a gambler playing a betting game. The gambler has control over what bets to make and how much to bet. Starting with an initial wealth c_0, the gambler's goal is to reach a target wealth T at which point she is satisfied and stops playing. Ofcourse, the catch is that she may lose all her initial wealth before attaining her target. This is clearly an undesirable event and one would like to minimize the likelihood that it occurs (and simultaneously maximize the likelihood of reaching the target). More specifically, what (optimal) betting strategy should the gambler use in order to maximize her probability of attaining her goal (wealth T), before being ruined (wealth 0)? When the bet size is constant, and the only outcomes are a win or a loss, then this is the well studied standard gambler's ruin problem [**?**]. Connolly [**?**] has studied a particular betting strategy where bet size is allowed to vary with the gamblers wealth and shows that this produces better results than the fixed bet size strategy. We will elaborate on his study by showing that his strategy is in fact optimal and, in addition, we will find all other optimal strategies. A closer analysis of these strategies reveals an unsuspected relationship to Knopp curves, curves that are continuous everywhere but no where differentiable, and in general possesing fractal properties. Finally we will consider games in which more than two outcomes are possible, and we will discuss the possibilities of spreading bets over multiple independent games.

1. A Simple Win/Loss Betting Game

As already mentioned, the gambler starts with some initial wealth c_0 and she will play until she either goes broke or attains her goal of T dollars. On each play, the gambler determines an amount to bet, b, which can depend on her current wealth. She either wins the bet, with probability p, in which case her wealth is increased by an amount b or she loses the bet with probability $1-p$, in which case her wealth will decrease by b. Depending on how much the gambler chooses to bet on each round, and on the outcomes of each bet (which are assumed to be independent), the wealth of the gambler will vary with time. Let the gambler's wealth after t plays be c_t.

The amount that the gambler bets on each play is determined by the gamblers betting strategy, B. Naturally, the gambler will use the betting strategy that will give her the best chance of reaching her goal. The betting strategy, B, should satisfy certain intuitive properties. The gambler can never bet more than her current wealth c_t, and it stands to reason that she will never bet more than what is necessary to reach her goal, $T - c_t$. Furthermore, it is of no relevence to the gambler how the wealth c_t was attained. All the gambler cares about in determining her future strategy is the value of her current wealth. In other words, her betting strategy will only depend on her current wealth. Thus, the gambler's entire betting strategy B can be sumarised by a function $B(c_t)$ that gives the bet b_t when the current wealth is c_t (the bet may also depend on T and p but we do not make these dependencies explicit). Thus, every time the gambler finds herself with wealth $c \in \{1, \ldots, T\}$, she will make the same bet, $B(c)$, with $B(c) \in \{1, \ldots, \min(c, T-c)\}$. A quick calculation shows that the set of all possible betting strategies has a size given by $\lfloor \frac{T}{2} \rfloor ! \lfloor \frac{T-1}{2} \rfloor !$, a large number. We consider a strategy to be optimal if it maximizes the probability that the gambler leaves the casino having achieved her target T. Of the large number of strategies available, the gambler must pick one, and a method for picking optimal strategies will be our concern here.

Using a betting strategy B, the gambler's worth is a random process that obeys the following update rule.

$$c_{t+1} = \begin{cases} c_t + B(c_t) & \text{with probability} \;\; p \\ c_t - B(c_t) & \text{with probability} \;\; 1-p \end{cases} .$$

where $B(c) \in \{1, \ldots, \min(c, T-c)\}$. We refer to the gambler's wealth c_t at any time as the *state* of the game, and the gambler's bet as the *action* to be taken. Thus the betting strategy B maps states to actions.

Assume, for the sake of realism, that $p < \frac{1}{2}$, in other words, the game is not

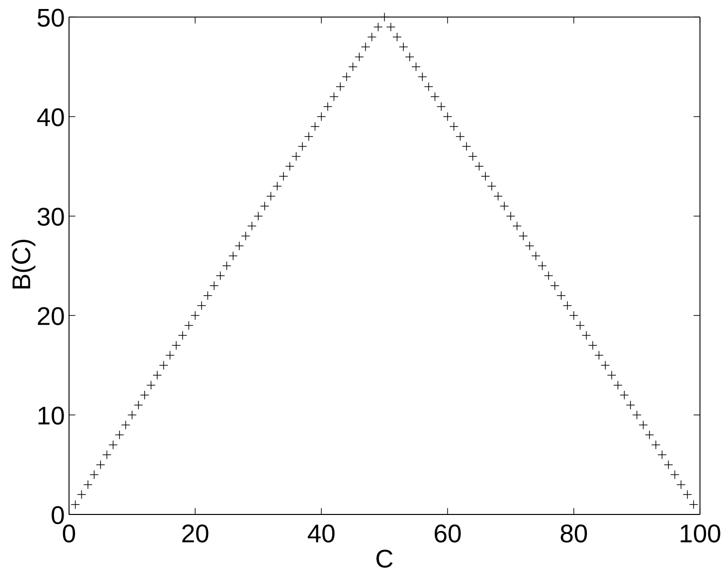

Figure 1: Ultimate betting strategy for $T = 100$.

fair, with the gambler expecting to lose. In [1] Connolly analyzes what he calls the 'ultimate betting strategy'. This strategy requires the bettor to bet as much as is needed to reach the target T or as much as he has, whichever is least. That is, $B(c) = \min(c, T-c)$ for all c. This betting strategy is illustrated in Figure 1 for a target worth of $T = 100$. We will use techniques borrowed from dynamic programming to show that Conolloy's ultimate strategy is in fact optimal for any T (in terms of maximizing the probability of attaining the goal), but that it is not necessarily unique. We will give a method for finding all the optimal betting strategies.

Evaluating Strategies In order to evaluate a strategy, introduce the function $Q_B(c)$ for a given betting strategy. $Q_B(c)$ is the probability of attaining the goal by following strategy B with starting wealth c (that is one bets $B(c_t)$ at each time step until either the goal T is reached or ruin occurs). $Q_B(c)$ is refered to as the *value* of the state c under strategy B. The function Q_B must satisfy the following consistency conditions.

$$\begin{aligned} Q_B(i) &= 0 \qquad \text{for } i \le 0 \\ Q_B(i) &= 1 \qquad \text{for } i \ge T \\ Q_B(i) &= pQ_B(i + B(i)) + (1-p)Q_B(i - B(i)) \qquad \text{for } 0 < i < T \end{aligned} \tag{1}$$

These consistency conditions can be written in a matrix form that simplifies their analysis.

$$A\vec{q} \;=\; \vec{b} \tag{2}$$

where

$$\begin{aligned}\vec{q} &= [Q_B(0), Q_B(1), Q_B(2), \ldots, Q_B(T-1), q_B(T)]^T \\ A_{ij} &= \begin{cases} p & i < T,\ i + B(i) = j \\ 1-p & i < T,\ i - B(i) = j \\ 1 & i = j = T \\ 0 & otherwise \end{cases} .\end{aligned}$$

and I is the identity matrix. Thus finding $\vec{q}$ amounts to finding a vector in the null space of $(A-I)$ subject to the condition $q_T = 1$. It can be shown that $\text{rank}(A-I) = T-1$, so there exists a unique solution.

This gives us a method of evaluating any betting strategy B. We can compute the probability of reaching the target for any initial worth. We illustrate with a simple example with $T = 5$. Let $B(1) = B(4) = 1$ and $B(2) = B(3) = 2$. Then

$$A = \begin{bmatrix} 0 & p & 0 & 0 & 0 \\ 0 & 0 & 0 & p & 0 \\ 1-p & 0 & 0 & 0 & p \\ 0 & 0 & 1-p & 0 & p \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

If $p = 9/19$, then $q = [0.173 \;\; 0.365 \;\; 0.565 \;\; 0.771 \;\; 1]^T$, that is, we have almost a 57% chance of reaching our target of 5 if we begin with 3.

Optimal Strategies Once we are able to evaluate a particular strategy, we can find the optimal strategy by a dynamic programming technique. It is important to note that in equation (1) each q_i is a monotonically nondecreasing function of all other q_j. Thus improving the strategy for any particular c cannot decrease q_j for any j. Therefore we can select some value for c, find the bet b that maximizes q_c and update the strategy so that $B(c) = b$. Since there are only finitely many strategies, if we iterate over c, we are guaranteed to converge to the optimal strategy in a finite number of steps.

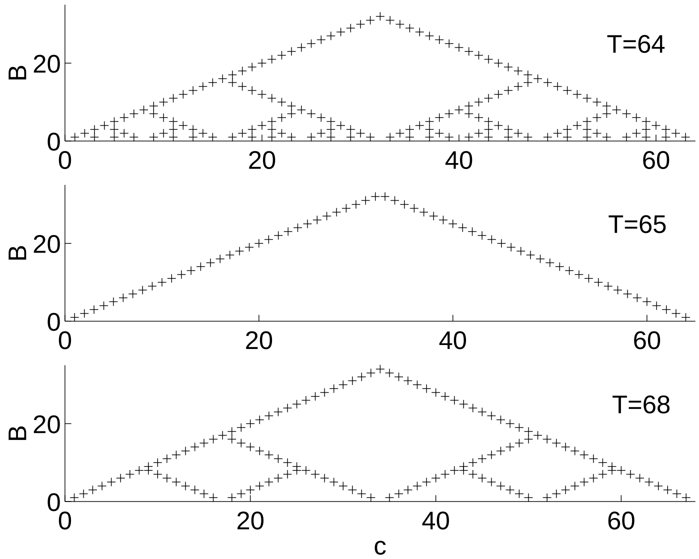

Figure 2: Optimal betting strategies for T = 64,65,68.

2. Expected Losses

Now that we have discovered the optimal wagering strategy, we investigate the gambler's expected losses. We have shown that there may be several optimal betting strategies. Thus a gambler may be able to choose between several different bets to add variety without increasing the risk. We note that Connolly's ultimate strategy is an optimal strategy for any target T, so we will use it for the purposes of the following analysis. Furthermore, we normalize with respect to T and take the limiting case $T \to \infty$, so that C and b are real numbers in $[0, 1]$. Then the ultimate strategy is defined by the function $b : [0, 1] \to [0, 1]$

$$b(c) = \begin{cases} c & 0 \leq c \leq \frac{1}{2} \\ 1 - c & \frac{1}{2} \leq c \leq 1 \end{cases}$$

Using the ultimate strategy with probability of success $p = \frac{1}{2} - \delta$, the expected loss given starting capital c obey the following relations. When $0 \leq c \leq \frac{1}{2}$

$$\begin{aligned} L(c) &= (1-p)c - p\ c + p\ L(2c) \\ &= 2\delta\ c + p\ L(2c) \end{aligned}$$

since the loss is c if the first bet is lost, otherwise we have a gain of c and the

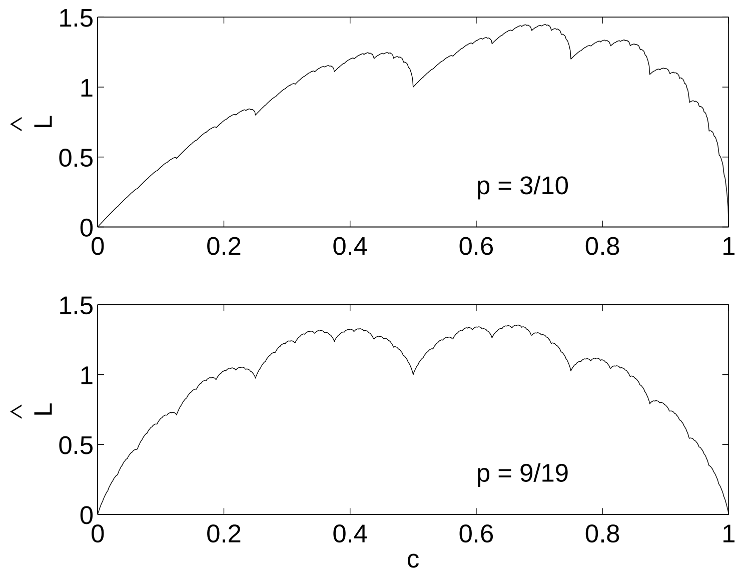

Figure 3: Expected losses for $p = 3/10$ and $p = 9/19$.

expected losses from the new state. When $\frac{1}{2} \leq c \leq 1$

$$\begin{aligned} L(c) &= -p(1-c) + (1-p)(1-c) + (1-p)L(2c-1) \\ &= 2\delta\ (1-c) + (1-p)L(2c-1) \end{aligned}$$

For the sake of comparison, we normalize with respect to δ, defining

$$\widehat{L}(c) = \frac{L(c)}{\delta}.$$

Then the normalized losses obey the functional relations

$$\widehat{L}(c) = \begin{cases} 0 & c \in \{0,1\} \\ 2c + p\ \widehat{L}(2c) & 0 \leq c \leq \frac{1}{2} \\ 2(1-c) + (1-p)\widehat{L}(2c-1) & \frac{1}{2} \leq c \leq 1 \end{cases} . \tag{7}$$

We show a plot of the normalized expected losses for $p = 3/10 (\delta = 1/5)$ and $p = 9/19$. It is interesting to note that the relation in equation (7) gives rise to a fractal-like curve with apparent self-similarity. In fact, we can write the solution to

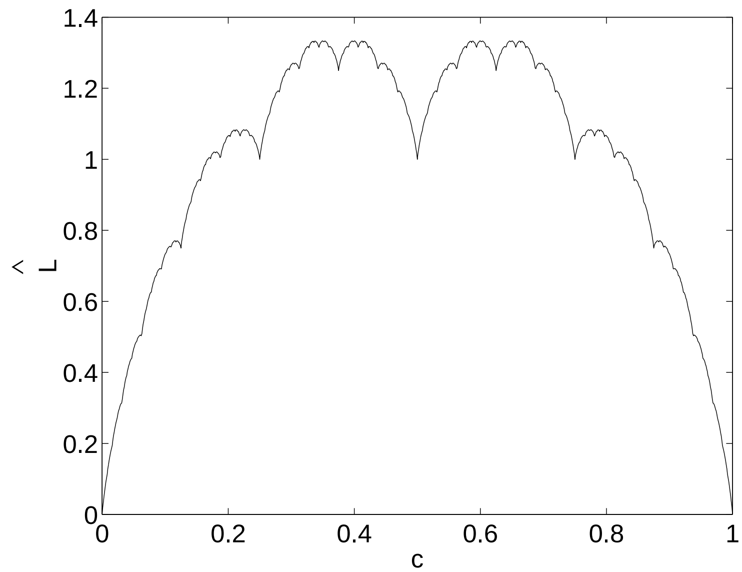

Figure 4: The Knopp function.

(7) as an infinite series.

$$\begin{aligned} g_0(x) &= \begin{cases} 2x & 0 \le x \le \frac{1}{2} \\ 2-2x & \frac{1}{2} \le x \le 1 \end{cases} \\ g_{i+1}(x) &= \begin{cases} p\ g_i(2x) & 0 \le x \le \frac{1}{2} \\ (1-p)g_i(2x-1) & \frac{1}{2} \le x \le 1 \end{cases} \\ \widehat{L}(c) &= \sum_{i=0}^{\infty} g_i(c). \end{aligned}$$

$\widehat{L}$ is continuous but differentiable nowhere in $[0,1]$. In the limit as $p \to \frac{1}{2}$, $\hat{L}$ approaches the Knopp function, illustrated in figure 4.

3. Changing The Odds

In many casino games, several different bets may be made, each having its own probability of success and odds paid by the house. Often there are several possible outcomes, and the payoff may take several values. It is straightforward to include these scenarios in the procedure for determining the optimal betting strategy given in section 1. We begin with the simplest extension, that of a single bet with different odds.

Uneven Odds We consider the case in which each game has only 2 outcomes, but the bet pays $M : 1$ odds, with $M \neq 1$. That is, for a bet b, the gambler wins $M \cdot b$ with probability p, and forfeits the bet with probability $1 - p$.

We illustrate the optimal strategy for $M = 2$, but it should be evident that the approach of section 1 can be used by discretizing allowable bets (as is typically the case in casinos) for any rational M.

While the single bet, uneven odds scenario is typical of many sports parimutual wagers, many casino games allow the player to make several different types of bet.

Multiple Betting Options Returning to the game of roulette, we now consider placing any single table bet. These bets are listed in Table 1.

Table 1: Single Roulette bets

1 : 1	Even, Odd, Red, Black, 1-18 or 19-36
2 : 1	1-12, 13-24, 25-36, 3n, 3n+1 or 3n+2
8 : 1	Any four adjacent numbers
17 : 1	Any two adjacent numbers
35 : 1	Any single number

For any state c, the action is now represented by an ordered pair (b, β), where b is the amount to be bet as before, and β represents which type of bet is to be placed.

Further Extensions We have not addressed scenarios that involve placing multiple simultaneous bets (for example, betting several outcomes at roulette). There are also some cases in which the probabilities or amounts of bets may change before they are settled (for example, splitting at blackjack or placing new bets at craps after the point is established).

4. Conclusion

We have shown that the ultimate betting strategy as described by Connolly is an optimal strategy for achieving a fixed target. This strategy is not necessarily unique in its optimality, though, and we have illustrated the equivalence class of optimal bets. We have also discovered that, using an optimal strategy, the expected loss for a single bet gambling game is a complex and nonintuitive function of the initial wealth.

References

[1] D. Connolly, Casino Gambling, The Ultimate Strategy, *College Math. J.* 30 (1999), 279–278