HOW TO GAMBLE IF YOU'RE IN A HURRY

SHALOSH B. EKHAD, EVANGELOS GEORGIADIS, AND DORON ZEILBERGER

The excitement that a gambler feels when making a bet is equal to the amount he might win times the probability of winning it.
–Blaise Pascal.

Preamble

This article is a brief description of the Maple package `HIMURIM` downloadable from

`http://www.math.rutgers.edu/~zeilberg/tokhniot/HIMURIM` .

Sample input and output files can be obtained from the webpage:

`http://www.math.rutgers.edu/~zeilberg/mamarim/mamarimhtml/himurim.html` .

The Maple package `HIMURIM` is to be considered as the **main output** of this project, and the present article is to be considered as a short user's manual.

We also briefly describe another Maple package downloadable from

`http://www.math.rutgers.edu/~zeilberg/tokhniot/PURIM` .

How To Gamble If You Must

Suppose that you currently have x dollars, and you enter a casino with the hope of getting out with N dollars, (with, x and N, being positive integral values). The probability of winning *one* round is p $(0 < p < 1)$. You can stake any integral amount of dollars $s(x)$ (that must satisfy $1 \leq s(x) \leq min(x, N-x)$), until you either exit the casino with the hoped-for N dollars, or you become broke. Deciding the value of the stake $s(x)$, for each $1 \leq x < N$, constitutes your *strategy*. Naturally, the question of whether a strategy is *optimal* arises; the three main optimality criteria in gambling theory can be summarized as follows:

1 Maximizing the probability of reaching a specified goal (i.e., amount N), with no time limit.
2 Maximizing the probability of reaching a specified goal by a fixed time t_0.
3 Minimizing the expected time to reach a specified goal, subject to a pre-specified level of *risk-aversion*.

In their celebrated masterpiece, Dubins and Savage [4] proved that the optimal strategy (using the first criterion), if $p \leq \frac{1}{2}$, is the **bold** one taking $s(x) = min(x, N-x)$, always betting the maximum, and if $p \geq 1/2$, then an optimal strategy is the *timid* one, with $s(x) = 1$, always betting the minimum.

A beautiful, lucid, and accessible account of these results can be found in Kyle Siegrist's [8] on-line article.

Date: December 6, 2011.
E.G. would like to thank Ed Thorp, Sid Browne, Timothy Chow and Yuval Peres for background comments related to the Kelly strategy and its applications and remains greatly indebted to Dimitrios Kavvathas and Kostas Pantazopoulos of Goldman Sachs as well as Michael Sipser and John Tsitsiklis of MIT. E.G. was supported in part by a DK-GS-007 grant. D.Z. was supported in part by the USA National Science Foundation.

Alas, if you play timidly, i.e. according to the classical "gambler's ruin" problem ([5], p. 348, Eq. (3.4)) your *expected time* until exiting is (let $q := 1 - p$)

$$\begin{cases} \frac{x}{q-p} - \frac{N}{q-p}\frac{1-(q/p)^x}{1-(q/p)^N} & \text{if } p \neq \frac{1}{2}; \\ x(N-x), & \text{if } p = \frac{1}{2}, \end{cases}$$

and this may take a very long time. If $p > \frac{1}{2}$, but you're in a hurry, then you may decide to take a slightly higher chance of exiting as a loser if that will enable you to expect to leave the casino much sooner. It turns out that the bold strategy is way too risky. For example, if $p = 3/5$ and right now you have 100 dollars and the exit amount is 200 dollars, with the bold strategy, sure enough, you are guaranteed to exit the casino after just one round, but your chance of leaving as a winner is only 3/5.

As a compromise, we can employ a *deterministic fixed fractional* betting strategy, namely, the *Kelly strategy*[1], with factor f denoting a fixed fraction of our money. This is inspired by [6], however, in that paper, the underlying assumptions are: money is infinitely divisible, the game continues indefinitely, the game has even payoff, and the opponent is infinitely wealthy. Under those circumstances, Kelly recommends to take $f = 2p - 1$ for his agenda. Based on our set of assumptions – using integral values – we obtain,

$$K(f)(x) := min(\lfloor xf \rfloor + 1, N - x) \quad .$$

For example, the Kelly strategy with $f = 1/10$ (and still $p = 3/5$) enables you to exit as a winner with probability %99.98784517, but the expected duration is only 44.94509484 rounds, much sooner than the expected duration of 500 rounds (with a fat tail!) promised by the timid strategy.

Inspired by Breiman [1] we can generalize the Kelly-type strategy, and "morph" it with the bold strategy, and play boldly once our capital is $\geq Nc$, in other words

$$B(f,c)(x) := \begin{cases} min(\lfloor xf \rfloor + 1, N - x), & \text{if } x < cN; \\ min(x, N - x), & \text{if } x \geq cN. \end{cases}$$

For example, taking $f = 1/10, c = 4/5$ (and still $p = 3/5$), your probability of exiting as a winner is %99.98721302, only slightly less than Kelly with $f = 1/10$, but your expected stay at the casino is about one round less (43.81842784). Paradoxically, lowering the c to 2/5 is not advisable, since your probability of winning is lower (%99.93836900) *and* you should expect to stay longer! (52.61769977 rounds). We observed, empirically, that for any f, lowering the c from 1 until a certain place $c_0(f)$ reduces the expected duration-until-winning (with a slightly higher risk of ultimate loss), but setting c below c_0 (i.e., playing boldly starting at cN) will not only lower your chance of ultimately winning, but would also prolong your agony of staying in the casino (unless you want to *maximize* your stay there, in which case you should play timidly).

Our question is: what is the optimal strategy according to Criterion 2 (i.e. maximizing the probability of reaching a specified goal by a fixed time t_0)? Borrowing the colorful yet gruesome language of [3], you owe N dollars to a loan shark who would kill you if you don't return the debt in $\leq T$ units time (rounds of gambling). Luckily, you are at a *superfair* casino, with the probability of winning a single round being p (superfair means $p \geq \frac{1}{2}$). If your current capital is i dollars (so you need to make $N - i$ additional dollars in $\leq T$ rounds to stay alive), if you want to *maximize* your chances of staying alive, how many dollars should you stake ?

Finding the Best Strategy If You're in a Hurry

So suppose that you currently have i dollars, and you need to have $N - i$ additional dollars (totaling up to N dollars) in $\leq T$ rounds of gambling, where at each round you can stake any amount between 1 and

[1]The Kelly strategy is also known as the *Kelly system* or the *Kelly criterion*; terms first coined by Ed Thorp in [9] and [10]. The theoretical underpinnings of this strategy were provided by Breiman in [1].

$min(i, N-i)$. You want to *maximize* your chance of success. How much should you stake, and what is the resulting probability, let's call it $f(i,T)$. The probability of winning a single round is p.

Obviously $f(i,T)$ satisfies the *dynamical programming* recurrence

$$f(i,T) = \max\{(1-p)f(i-x, T-1) + pf(i+x, T-1) \,:\, 1 \leq x \leq \min(i, N-i)\},$$

with the obvious *boundary conditions* $f(0,T) = 0, f(N,T) = 1$ and *initial conditions* $f(N,0) = 1$, and $f(x,0) = 0$ if $x < N$.

The set of x's that attain this max constitutes your *optimal strategy.* It is most convenient to take the largest x (in case there are ties).

By repeatedly computing $f(j,T)$ and the stake-amount x that realizes it, where j is the current capital and T is the steadily decreasing time left, the gambler can always know how much to stake in order to maximize his chances of staying alive, and also know the actual value of that probability.

Finding the Optimal Strategy If the Probability, p, is Unknown[2]

We can say a few meaningful things about optimal strategies (*optimal*, here, refers to a strategy which attempts to maximize the probability of reaching a specified amount,N), when the winning probability, p is unknown.
In [2], Berry et al first consider the classical (continuous) red-and-black game where a gambler is allowed a sequence of bets. The gambler is permitted to stake any part of his current fortune x (where $0 < x < 1$) at even odds, where the winning probability, p, is known and fixed. The goal is to maximize the probability of reaching 1; if $p \to \frac{1}{2}$, then it is feasible to reach unity with probability 1. If we consider a subfair game (i.e., where $p \in [0, \frac{1}{2}]$), then the optimal strategy is to play boldly. In [2], Berry et al show that for $p \in (\frac{1}{2}, 1]$, the timid strategy where the stake is x^2 (for $x \in (0,1)$) is optimal. The mildly striking statement is that δ-bold strategies are ϵ-optimal for all p. The δ-bold strategy can be summarized as follows:

$$s(\delta)(x) := \begin{cases} x^2, & \text{if } 0 \leq x \leq \delta; \\ x-\delta, & \text{if } \delta < x < (1+\delta)/2; \\ 1-x, & \text{if } (1+\delta)/2 \leq x < 1; \\ 0, & \text{if } x \geq 1. \end{cases}$$

This new strategy is then shown to be also ϵ-optimal when p is unknown (yet a random variable whose probability distribution is known).

Far more relevant for us, is the problem in which the goal to be reached is a positive integer and the fortunes and allowed stakes are also integers. Berry et al, unfortunately only briefly touched this topic.
For fixed p, optimal strategies exist; for instance, for *subfair* games, the optimal strategy is to play boldly by always staking the largest available amount (this surprisingly enough follows from the analogous continuous case). For *superfair*, the timid strategy which always makes the least possible stakes is optimal.
Now, what happens when p is unknown ? In this discrete (integral) setting, unfortunately, the gambler (unlike in the continuous version) cannot safely ignore information on p. As a result, compelling and precise results about optimal strategies seem difficult without computer-assisted computation. Here is a known example which illuminates the problematics:

Assume that your goal is $N = 8$ and your current fortune is $x = 5$, the probability of win, p is either $\frac{1}{4}$ or $\frac{3}{4}$. Along with this information comes the prior probability of p, $P(p = 1/4) = P(p = 3/4) = 1/2$. You are not able to observe p directly; the only information are whether you win or lose. What is an optimal strategy here ? It turns out that the "unique" optimal initial stake has to be 3.

[2]We will consider this problem in greater depth in a follow-up paper.

The Maple Package HIMURIM

`HIMURIM` is downloadable, *free of charge*, from

`http://www.math.rutgers.edu/~zeilberg/tokhniot/HIMURIM` .

We will only briefly describe some of the more important procedures, leaving it to the readers to explore and experiment with the many features on their own, using the on-line help.

The most Important Procedures of HIMURIM

The most important procedure is `BestStake(p,i,N,T)`, that implements $f(i,T)$ with the given p and N.

If you want so see the *full* optimal strategy, a list of length $N-1$ whose i-th entry tells you how much to stake if you have i dollars, use procedure `BestStrat(p,N,T)`. See for example:

`http://www.math.rutgers.edu/~zeilberg/tokhniot/oHIMURIMk1` .

for the output of `BestStrat(11/20,1000,30);` .

Procedure `SimulateBSv(p,i,N,T)` simulates *one* game that follows the optimal strategy.

Finally, Procedure `BestStratStory(m0,N0,T0,K)` collects optimal strategies for various p's (all superfair), N's (exit capitals) and T (deadlines).

Other Procedures of HIMURIM

Procedure `ezraLA()` lists the procedures that use *Linear Algebra* to find the *exact* probabilities, expected duration, and probability generating functions for the random variables "duration" and "duration conditioned on ultimately winning" for a casino with exit capital N, probability of winning a round p (that may be either numeric or *symbolic*), and *all* possible initial incomes.

For example, `PrW(p,S)` inputs a probability `p` (a number between 0 and 1 or left as a symbol p) and a list `S`, of length $N-1$, say, where $S[i]$ tells you how much to stake if you have i dollars. It outputs the list, let's call it `L`, such that `L[i]` is the probability of ultimately winning (exiting with N dollars) if you currently have `i` dollars and always play according to strategy `S`.

It works by solving the system of $N-1$ equations for the $N-1$ unknowns $L[1],\dots,L[N-1]$

$$L[i]=(1-p)L[i-S[i]]+pL[i+S[i]] \quad ,1 \leq i \leq N-1 \quad ,$$

together with the boundary conditions $L[0]=0, L[N]=1$.

For example, if $N=3$, and the strategy S, being $[1,1]$ (the only possible strategy when $N=3$)

`PrW(1/3,[1,1]);`

would yield

`[1/7, 3/7]` ,

that means that if the probability of winning a round is $\frac{1}{3}$ and you exit the casino when you either reach 0 or 3 dollars, then your probability of exiting as a winner, if you currently have one dollar is $\frac{1}{7}$, and if you currently have 2 dollars, is $\frac{3}{7}$.

Procedure `ED(p,S)` inputs p and S as above and outputs the list, let's call it `L`, such that `L[i]` is the expected duration until getting out (either as a winner or loser) if you currently have `i` dollars and always play according to strategy `S`.

It works by solving the system of $N-1$ equations for the $N-1$ unknowns $L[1],\dots,L[N-1]$

$$L[i]=(1-p)L[i-S[i]]+pL[i+S[i]]+1 \quad , \quad 1 \le i \le N-1 \quad ,$$

together with the boundary conditions $L[0]=0, L[N]=0$.

For example, still with $N=3$ and $S=[1,1]$,

```
ED(1/3,[1,1]);
```

would yield

```
[12/7, 15/7] ,
```

that means that if the probability of winning a single round is $\frac{1}{3}$ and you exit the casino when you either reach 0 or 3 dollars, and you follow strategy $[1,1]$, then the expected remaining duration, if you currently have one dollar, is $\frac{12}{7}$, and if you currently have 2 dollars, it is $\frac{15}{7}$.

Procedure `EDw(p,S)` inputs p and S as above and outputs the list, let's call it `L`, such that `L[i]` is the expected duration until getting out (conditioned on being an ultimate winner!) if you currently have `i` dollars and always play according to strategy `S`.

For example, with the above (trivial) input

```
EDw(1/3,[1,1]);
```

would yield

```
[18/7, 11/7] ,
```

that means that if the probability of winning a round is $\frac{1}{3}$ and you exit the casino when you either reach 0 or 3 dollars, then the expected remaining duration until winning (assuming that you do win), if you currently have one dollar is $\frac{18}{7}$, and if you currently have 2 dollars is $\frac{11}{7}$.

Procedure `Dpgf(p,S,t)` inputs p and S as above, as well as a *symbol* (variable name) t, and outputs the list, let's call it `L`, such that `L[i]` is the probability generating function, in t, for the random variable "remaining duration" if you currently have i dollars, i.e. if you take the Maclaurin expansion of `L[i]` and extract the coefficient of t^j you would get the exact value of the probability that the game would last j more rounds.

For example, still with the same N and S,

```
lprint(Dpgf(1/3,[1,1],t));
```

would yield

```
[-t*(6+t)/(-9+2*t**2), -t*(3+4*t)/(-9+2*t**2)]   .
```

Typing "`taylor(-t*(6+t)/(-9+2*t**2),t=0,5);`" would yield

```
2/3*t+1/9*t**2+4/27*t**3+2/81*t**4+O(t**5)
```

meaning that if you play the above game with $p=1/3, N=3$ and you currently have one dollar, you would have probability $2/3$ of exiting after one round, probability $1/9$ of exiting after two rounds, probability $4/27$ of exiting after three rounds, and probability $2/81$ of exiting after four rounds.

Procedure `DpgfW(p,S,t)` is analogous to `Dpgf(p,S,t)` but the duration is conditioned on the fortunate event of exiting as an ultimate winner. For example, for the timid strategy, and $N = 3$,

```
lprint(DpgfW(1/3,[1,1],t));
```

would yield

```
[-t**2/(-9+2*t**2), -3*t/(-9+2*t**2)]   .
```

Typing `taylor(-t**2/(-9+2*t**2),t=0,5);` would yield

```
1/9*t**2+2/81*t**4+O(t**6)
```

meaning that if you play the above game with $p = 1/3, N = 3$ and you currently have one dollar, and you are destined to leave as a winner, you would have probability of 0 of exiting after one round, probability of 1/9 of exiting after two rounds, probability of 0 of exiting after three rounds, probability of 2/81 of exiting after four rounds.

The Best Breiman-Kelly Strategies If You are in A Hurry

Procedure `BestBKdd(p,N,i,T,h)` tells you the best Breiman-Kelly Strategy if the probability of winning a round is `p`, you have `i` dollars and you must win `N` dollars in <= `T` rounds, and you're using resolution `h`. It also returns the expected duration until exit (either as a winner or loser).

To get the story for various initial capitals, and various deadlines, try out procedure `BestBKddStory(p,N,h,t0,MaxF,M0)`. See the on-line help, and the sample input and output files in

`http://www.math.rutgers.edu/~zeilberg/mamarim/mamarimhtml/himurim.html` .

The Best Kelly Factor With a Given Level of Risk-Aversion

Try out procedure `KellyContestx(p,N,x,h,conf)`. For example see:

`http://www.math.rutgers.edu/~zeilberg/tokhniot/oHIMURIM8a` ,

`http://www.math.rutgers.edu/~zeilberg/tokhniot/oHIMURIM8b` .

The Maple package `PURIM`

For an "umbral" approach, see our Maple package `PURIM` that tells you much more. It explores the whole "tree" of possibilities. See the package itself and the input and output files

`http://www.math.rutgers.edu/~zeilberg/tokhniot/inPURIM2` and

`http://www.math.rutgers.edu/~zeilberg/tokhniot/oPURIM2` for an example.

Further Work

There are many possible generalizations and extensions. See for example the interesting article [3].

Conclusion

We provided a playful yet algorithmic glimpse to a field that has up till recently flourished on the Kolmogorov, *measure-theoretic* paradigm [as evidenced by the work of Dubins and Savage [4] (see [7] for more recent developments]. The advent and omnipresence of computers, however, ushered an era of *symbol crunching* and *number crunching*, where a few lines of code can give rise to powerful *algorithms*. And it is the ouput of algorithms that usually provides insight (and inspiration) for conjectures and theorems. Those, in turn, can then be proven in their respective measure-theoretic settings. Additionally, a computational approach lends itself easily to more complex scenarios that would otherwise be considered pathological phenomena (and would be fiendishly time-consuming to prove – even for immortals like Kolmogorov and von Neumann).

References

[1] L. Breiman, *Optimal gambling systems for favorable games*, Fourth Berkley Symposium on Math. Stat. and Prob. **1** (1961), 65-78.

[2] D. A. Berry and D. C. Heath and W. D. Sudderth, *Red-and-black with unknown win probability*, Ann. Statist. 2 (1974), 602608.

[3] R.W. Chen, L.A. Shepp, Y.-C. Yao, and C.-H. Zhang, *On optimality of bold play for primitive casinos in the presence of inflation*, J. Appl. Prob. **42** (2005), 121-137.

[4] Lester E. Dubins and Leonard J. Savage, *Inequalities for Stochastic Processes* (How to Gamble If You Must), Dover, 1976.

[5] William Feller, "*An Introduction to Probability Theory and Its Application*", Wiley,volume 1, third edition, 1968.

[6] J.L. Kelly, Jr, *A new interpretation of information rate*, Bell System Tech. J. **35** (1956), 917-926.

[7] A.P. Maitra and W.D. Sudderth, *"Discrete Gambling and Stochastic Games"*, Springer, 1996.

[8] Kyle Siegrist, *How to Gamble if you must*, JOMA v. 8, `http://www.maa.org/joma/Volume8/Siegrist/RedBlack.pdf`

[9] E.O. Thorp, *"Beat the Dealer"*, 2nd Ed., Vintage, New York 1966.

[10] E.O. Thorp and W. Walden, *"A winning bet in Nevada baccarat."*, J. Amer. Statist. Assoc., 61 Part I, (1966) 313-328.

Shalosh B. Ekhad.

E-mail address: `zeilberg@math.rutgers.edu`

Evangelos Georgiadis.

E-mail address: `egeorg@mit.edu`

c/o Doron Zeilberger, Department of Mathematics, Rutgers University, Hill Center, 110 Frelinghuysen Road, Piscataway, NJ 08854-8019 USA.

E-mail address: `c/o zeilberg@math.rutgers.edu`