Uncertainty in the Binomial Game

Philip Bretz

March 2020

1 Introduction

1.1 Goals

In this paper, I seek to answer three questions.

Probability of Ruin First, when betting the optimal amount in this game, what is the likelihood that a player will end up with less than 20% of their initial wealth?

In this game, there is an optimal amount to bet, but that amount depends on the probability of winning. If that probability is unknown, but estimated from data on previous games, that introduces an element of uncertainty. There is always a risk of ruin, even with exact knowledge of the probability of winning. However, this uncertainty will increase that risk, in a calculable way.

Managing Overestimation Second, what is the highest a player should bet to ensure, with a certain probability, that they are not betting more than the optimal amount?

This is a very different type of question than the first. Essentially, this question is about managing the risk introduced by uncertainty. The danger of this game is that betting more than the optimal amount introduces a greatly increased amount of risk.

Naturally, if the danger is overestimation, a good strategy is to bet less than we estimate. This question deals with how much less we should bet.

Managing Ruin Third, what is the highest a player should bet to ensure, with a certain probability, that they do not lose more than 20% of their wealth?

This question combines the first two in a natural manner. This is, of course the type of question a player would most like to know the answer to. The player wants to know the strategy that, with a certain probability, lets them avoid ruin.

1.2 Setup

Binomial Game In this simple betting game, each play is either a win or a loss. The player chooses an amount to bet. If the result of a play is a win, then the player's investment increases from 1 to $1 + b$. If the player loses a play, then the player's investment decreases from 1 to $1 - a$. Denote the player's wealth after n plays by X_n.

Kelly Criterion The Kelly criterion says that there is an optimal fraction of your current bankroll you should wager to maximize wealth in the long run. Let f denote the fraction chosen by the player. For the binomial game, the optimal fraction, $f = f^*$, is known. If the probability of a win is p and the probability of a loss is $q = 1 - p$, then the optimal fraction is

$$f^* = \frac{p}{a} - \frac{q}{b}.$$

Note that for this fraction to be positive, we need $pb > qa$. We interpret this to mean that if $pb \geq qa$, the game is not worth playing.

While the proportion f^* is optimal for growth in the long run, it can be risky in the short run. So, to manage risk, a player can choose a 'fractional Kelly' strategy. The player chooses a fixed fraction of the growth optimal f^*, i.e., they choose $f = \lambda f^*$ for some fixed λ. For example, a player might choose a half-Kelly strategy, where the fraction of their wealth they bet at each turn is $\frac{f^*}{2}$.

Alternate Parameters With the fractional Kelly strategy, a natural alternate parameter to replace f is λ. In some ways this parameter is more informative than f, since it describes the strategy of the player with respect to the optimal growth strategy.

In Appendix A, I showed that upon varying a and b, the results of the game stay fixed when a is proportional to b. So, we can eliminate these two parameters with a single parameter $k = \frac{b}{a}$ (as long as $a \neq 0$).

Alternatively, we can just fix $a = 1$ and let b vary and we will not lose any information. In this case, the condition $pb > qa$ has a very intuitive meaning.

This inequality can be manipulated to yield $p > \frac{1}{1+b}$. Now, in the binomial game where $a = 1$ (i.e., you lose everything you bet upon a loss), $\frac{1}{1+b}$ is exactly the probability of winning that is implied by the odds, which we will denote by p_i. So, you should only play when the underlying probability p is larger than the implied probability p_i. It is only in this scenario that the game is favorable to the player.

Ruin There are many different measures of ruin. For the binomial game, here is a typical way to define ruin. For this work, we will be using this measure of ruin.

For a given level $r \in [0, 1]$ and initial wealth X_0, the player is in a state of ruin upon the n-th play if the player's wealth X_n is at or below $r \cdot X_0$. Equivalently, ruin is when the player's relative wealth is below r:

$$\frac{X_n}{X_0} \leq r.$$

While the event ruin is defined naturally by the above inequality based on wealth (or relative wealth), we can give an equivalent definition with an inequality based on the number of wins. The player is in a state of ruin upon the n-th play if

$$W_n \leq C_{f,a,b,r} + D_{f,a,b}n$$

where W_n is the number of wins upon the n-th play and C and D are constants that depend on the parameters f, a, b, r and f, a, b, respectively. In Appendix B, I give the proof and show what these constants are, dependent on f, a, b, r.

In previous notes, I investigated a different measure of ruin. Ruin occurs by the n-th play if $X_k \leq r \cdot X_0$ for some $k \leq n$. I.e., rather than just looking at the final wealth, the player is ruined if their wealth ever drops below a specified level.

While this measure of ruin is interesting, the only way to compute the probability of ruin is recursively or to estimate by simulation. In comparison, the measure of ruin that we will be working with can be computed directly, using the cumulative distribution function that W_n follows.

1.3 Uncertainty

In the binomial game, if we know the true underlying probability of a win, p, we can easily compute the optimal for growth Kelly fraction f^*. However, in practice, we may not know p.

In a more realistic scenario, we need to use an estimate for p, p_e, given knowledge of the outcomes of previous games. For example, if we know that in the past, a player won 2 games out of 10, we would estimate $p_e = 0.20$. Based on p_e, we would then compute our estimated Kelly fraction f_e.

Since f_e is dependent on p_e, the wrong estimate can yield drastic consequences. If we choose our Kelly fraction to be too large, the probability of ruin dramatically increases.

The question is, how can we analyze this uncertainty in p?

2 Probability of Ruin

We can consider this question in a Bayesian framework. Essentially, we assign the parameter p a distribution that expresses the probability of values of p. We then quantify how this distribution changes based on data.

2.1 Bayes' Rule

Bayes' rule states that

$$p(\theta|y) \propto p(y|\theta) \cdot p(\theta).$$

The term $p(\theta)$ represents our prior belief about the distribution of the parameter θ. In our example, this unknown parameter θ is the probability of a win, p. The term $p(y|\theta)$ represents the likelihood of the data y, given the parameter θ. Finally, $p(\theta|y)$ represents our posterior belief about the distribution of the parameter θ, given the data y.

2.2 Beta Binomial Conjugate Pair

In the binomial game, $W_n \sim Binomial(n, \theta)$ where θ is the underlying probability of winning. So, the likelihood of the data y (where y is the number of wins in n games) is distributed $y \sim Binomial(n, \theta)$.

If we choose the prior distribution for θ to be $\theta \sim Beta(\alpha, \beta)$, then we have what is known as a 'conjugate prior'. This means that, using Bayes' rule, the posterior will also follow a beta distribution with different hyper parameters. Since θ is the main parameter of interest, 'hyperparameters' refer to α and β. In this case,

$$\theta|y \sim Beta(\alpha', \beta')$$

where $\alpha' = \alpha + y$ and $\beta' = \beta + n - y$.

Let's say that we do indeed have data from a previous game where a player had y wins after n plays. The natural choice for our prior would be the uniform distribution. The uniform distribution can be expressed as the beta distribution $Beta(1, 1)$. So, the posterior will be

$$\theta|y \sim Beta(1 + y, 1 + n - y).$$

2.3 Posterior Predictive

In the Bayesian framework, we can also construct a distribution for predicted values given data. This is known as the posterior predictive distribution, $p(y'|y)$. For the beta binomial conjugate pair, the posterior predictive follows a distribution known as the beta-binomial:

$$y'|y \sim BetaBin(\alpha', \beta')$$

where $\alpha' = \alpha + y$ and $\beta' = \beta + n - k$.

So, if we choose $\alpha = 1, \beta = 1$,

$$y'|y \sim BetaBin(1 + y, 1 + n - y).$$

So, given data from a previous game y, we have a probability density function for the number of wins in a future game W_n:

$$W_n|y \sim BetaBin(1 + y, 1 + n - y).$$

This means that we can exactly compute the probability of ruin

$$p(W_n \leq C + Dn).$$

One thing to note is that C and D depend on fixed parameters r, a, b and our choice of f. In the alternate parameterization, these constants depend on r, k and our choice of λ coupled with the estimate p_e. Given data y, our estimate p_e will be the expected value

$$p_e = E[\theta|y] = \frac{y}{n}.$$

Of course, this matches up with our intuitive estimate for p_e: the number of wins over the number of plays.

So, using the posterior predictive, we can compute the conditional probability of ruin.

2.4 Examples

For these examples I use my alternate parameter $\lambda = \frac{f}{f^*}$ (the fraction of the optimal Kelly fraction) and set $a, b = 1$. I set $\lambda = 1$ (full Kelly). The data is simulated with a win probability of $p = 0.75$. Note that this is larger than the implied win probability $p_i = \frac{1}{1+b} = 0.5$ (the game is worth playing).

With those parameters set, I examined how the ruin probability changed in three different scenarios. The first scenario contains 10 simulated data points, the next 50, and the third 100. Lastly, we can compute the exact probability of ruin if p is known (to be 0.75 in this case).

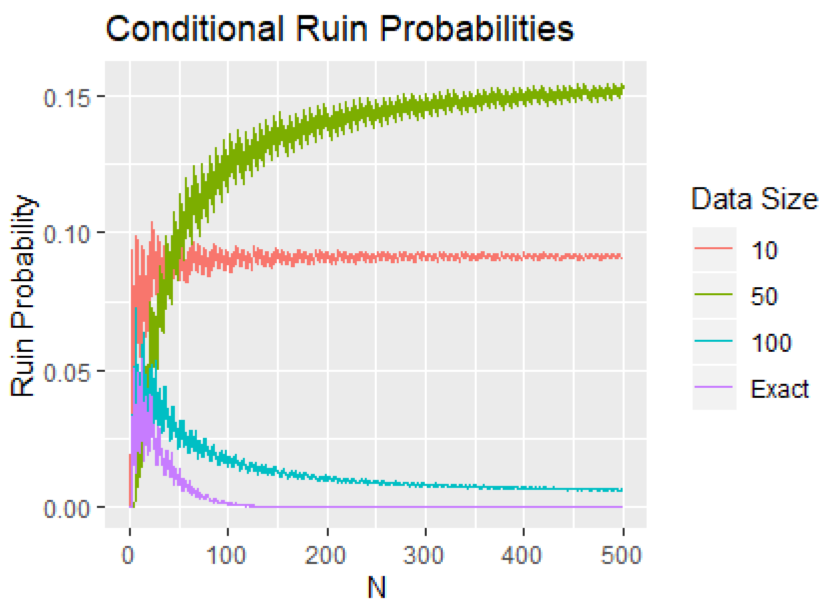

Figure 1: Conditional Ruin Probabilities

The more previous data we have, the more certainty we have in our estimate p_e. We see that with more data, the probability of ruin is noticeably smaller for essentially all values of N. Moreover, if we have very little data, the uncertainty about p can even lead to a near constant probability of ruin as time passes.

Additionally, for more previous data, the probability of ruin converges to the probability of ruin given exact knowledge of p. This is to be expected: for large values of α and β, the beta-binomial distribution converges to the binomial distribution with parameter $\frac{\alpha}{\alpha+\beta}$. With more data, we increase the value of the hyperparameters in the beta-binomial posterior predictive distribution, causing it to converge to the exact binomial.

3 Managing Overestimation

3.1 Choosing λ

In this game, our only choice is how much of our wealth to bet. So, to manage risk, we need to understand how uncertainty in p affects the chosen Kelly fraction. Again, for the following analysis, only consider $a = 1$ (i.e., you lose everything you bet).

If our estimate p_e differs from the true p by ϵ ($p_e = p + \epsilon$), then our chosen Kelly fraction f_e differs from the actual Kelly fraction f^* by

$$f_e = f^* + \epsilon((b+1)/b).$$

Note: we consider ϵ both positive and negative. Positive ϵ correspond to an overestimate of p and $\epsilon < 0$ corresponds to an underestimate.

Now, if we allow fractional Kelly with fraction λ, then the partial Kelly fraction f_e^λ is

$$f_e^\lambda = \lambda[f^* + \epsilon((b+1)/b].$$

A natural way to manage risk would be to choose $\lambda \geq 0$ that guarantees $f_e^\lambda \leq f^*$. I.e., we want

$$\lambda[f^* + \epsilon((b+1)/b)] \leq f^*.$$

This equality is equivalent to

$$\lambda[p_e - p_i] \leq p - p_i$$

where $p_i = \frac{1}{1+b}$ is the implied probability from the odds b.

Now, if $p_e - p_i \leq 0$, then $f_e = 0$; i.e., we will always choose not to play. So, let's only consider $p_e - p_i > 0$. Next, if $p - p_i \leq 0$ then $f^* = 0$, i.e., we need to choose $\lambda = 0$. So, we must choose λ such that

$$\lambda \leq \frac{p - p_i}{p_e - p_i}$$

when $p > p_i$ and $\lambda = 0$ otherwise.

So, if we think about the chosen λ as a function of p, $\lambda = g(p)$, then the natural choice for f is

$$\lambda = g(p) = \begin{cases} \frac{p-p_i}{p_e-p_i} & \text{if } p > p_i \\ 0 & \text{otherwise.} \end{cases}$$

3.2 Distribution of λ

If λ is a function of p, a natural problem arises. Namely, the whole point of this process is that we do not know p; we only have our estimate p_e. And if p is unknown, then $\lambda = g(p)$ is also unknown.

However, while the precise value might be unknown, we do have information about p. Namely, given previous data of y successes from n trials, and a uniform prior distribution,

$$p|y \sim Beta(1 + y, 1 + n - y).$$

Since λ is a function of p, then λ should have a distribution as well.

Discrete-Continuous Distribution If $g(p)$ was strictly monotone, then we could apply the transformation of random variables theorem and find the probability density function

$$f_\Lambda(\lambda) = f_\Theta(g^{-1}(\lambda)) \left| \frac{d}{d\lambda} g^{-1}(\lambda) \right|.$$

Unfortunately, $g(p)$ is not a strictly monotone function. The problem is that $g(p)$ maps all $p \le p_i$ to 0. So, the point $\lambda = 0$ has non-zero probability mass, which is incompatible with continuous distributions.

One solution is to have λ follow a mixed discrete-continuous distribution. At $\lambda = 0$, f_Λ is a discrete distribution with mass $P(p \le p_i)$ and for $\lambda > 0$, f_Λ is a continuous distribution with p.d.f according to the transformation of random variables theorem.

While this will certainly work, there is a simpler method.

Alternate Transformation We assigned λ the piecewise function

$$\lambda = g(p) = \begin{cases} \frac{p-p_i}{p_e-p_i} & \text{if } p > p_i \\ 0 & \text{otherwise} \end{cases}$$

because we interpreted $\lambda < 0$ to mean we should not play the game, and consequently set $\lambda = 0$ when that occurred. This is the natural interpretation.

However, we could allow $\lambda < 0$ without subsequently setting $\lambda = 0$. While negative λ values seem somewhat artificial, there is no problem with interpreting them. We simply retain our view that $\lambda < 0$ means we should not play, while still allowing for such values.

With this in mind, we can assign λ to be a slightly different function of p, namely,

$$\lambda = h(p) = \frac{p - p_i}{p_e - p_i}$$

for all values of p. Since $h(p)$ is strictly increasing, we can now apply the transformation theorem.

P.D.F. of λ To apply the transformation theorem, we first need to compute $h^{-1}(\lambda)$. This is simply

$$h^{-1}(\lambda) = \lambda(p_e - p_i) + p_i.$$

Additionally,

$$\left| \frac{d}{d\lambda} h^{-1}(\lambda) \right| = p_e - p_i.$$

So, given $p|y \sim Beta(1 + y, 1 + n - y)$ the conditional p.d.f. of $\lambda|y$ is

$$f_{\Lambda|Y}(\lambda|y) = \frac{[\lambda(p_e - p_i) + p_i]^y [1 - (\lambda(p_e - p_i) + p_i)]^{n-y}}{B(1 + y, 1 + n - y)} (p_e - p_i).$$

Mean, Variance, and Quantiles It is important to note that $h(p)$ is a particularly nice transformation. Namely, it is linear. With this in mind, the mean and variance of λ are simple:

$$E[\lambda] = \frac{E[p] - p_i}{p_e - p_i}$$

and

$$Var(\lambda) = \frac{Var(p)}{(p_e - p_i)^2}.$$

Additionally, we would typically estimate p_e from the distribution of p. I.e., we will normally choose $p_e = E[p]$. When we do this, we get

$$E[\lambda] = 1.$$

Another thing to note is that for any strictly monotone transformation like $\lambda = h(p)$, probability mass is preserved, so $P(\lambda \leq \lambda_0) = P(p \leq h^{-1}(\lambda_0))$ or, equivalently,

$$P(p \leq p_0) = P(\lambda \leq h(p_0)).$$

I give the proof in Appendix C.

So, for a given probability c, we might want to find λ_c such that the probability we did not overestimate the true λ is exactly c ($P(\lambda > \lambda_c) = c$). We can do this by finding p_c such that $P(p > p_c) = c$ and then $\lambda_c = h(p_c)$. Essentially, the transformation (and any monotone transformation) preserves quantiles.

3.3 Examples

Setup For these examples, set $a, b = 1$ and simulate data from underlying $p = 0.75$. For this game, $p_i = \frac{1}{1+b} = 0.5$.

Given the data, let's find the value λ_c we should pick so that the probability the true λ is above λ_c is c. I.e., find λ_c such that

$$P(\lambda > \lambda_c) = c.$$

We will do this by finding p_c such that $P(p > p_c) = c$ and then $\lambda_c = h(p_c)$.

Why do we want the true value of λ to be above our chosen λ_c? Well, if this is the case, then by underestimating $lambda$, we will be betting less of our wealth. So, by choosing λ_c, we ensure that with probability c, we are incurring less risk than we would if we chose the actual best value of λ.

Varying Data Size First, let's examine how λ_c changes when we fix c, but vary the size of the simulated data. We should expect that with more data, we have more certainty about our estimate p_e, so λ_c should tend towards 1. I.e., when we have more certainty about p_e, we should tend closer to a full Kelly strategy, the strategy we choose when we know the true value of p.

Let's use previous data with size 10, 100, and 1000 with c fixed at 0.95. I.e., we are finding $\lambda_{0.95}$ such that, with probability 0.95, we are underestimating the true value of λ.

	n	lambda
1	10	-0.32
2	100	0.6
3	1000	0.91

Figure 2: Varying Data Size

As expected, with very little data, we have to choose a small value of $\lambda_{0.95}$ to manage the risk of overestimating p. But as we grow more certain about the underlying probability of winning, we have larger values of $\lambda_{0.95}$ that approach 1.

Moreover, we even see λ take on a negative value. This means that the small amount of data suggested the game was unlikely to be favorable to us with probability 0.95. So, given this data, and our chosen level of $c = 0.95$, we would choose not to play.

Varying c Now, let's vary c with the simulated data fixed. We should expect λ_c to increase as c decreases, since we are allowing for more risk.

Additionally, the value c' for which $p_{c'} = p_e$ (i.e., the percentile of the mean of p) will give $\lambda_{c'} = 1$. Any values of c larger than c' will give $\lambda_c < 1$ and any values $c < c'$ will give $\lambda_c > 1$. If the posterior distribution for $p|y \sim Beta$ has no skew, then the median and mean are the same: this will occur at $c' = 0.50$.

Let's simulate 100 data points and test $c = 0.95, 0.75, 0.50, 0.25$.

	c	lambda
1	0.95	0.72
2	0.75	0.88
3	0.5	0.99
4	0.25	1.09

Figure 3: Varying Probability of Underestimation

As expected, λ_c increases as c decreases and at $c = 0.50$, we have λ_c very close to 1. Moreover, if we allow more risk at the $c = 0.25$ level, we have $\lambda_c > 1$.

4 Managing Ruin

To be written up.

5 Conclusions

To be written up.

6 Appendices

6.1 A. Dependence of Parameters a and b

Let $k = \frac{a}{b}$. Then

$$f^* = \frac{1}{b}(p(k+1)-1).$$

Let X_n denote our wealth after the nth play and W_n denote the number of wins after the nth play. Let's examine our relative wealth:

$$\frac{X_n}{X_0} = (1+fb)^{W_n}(1-fa)^{n-W_n}.$$

If we use $f = \lambda f^*$ and replace $a = kb$, we find that this simplifies to

$$\frac{X_n}{X_0} = [1+\lambda(p(k+1)-1)]^{W_n}[1-\lambda k(p(k+1)-1)]^{n-W_n}.$$

The key here is that our relative wealth is only dependent on parameters p, λ, and k and the stochastic process W_n (which only depends on parameter p). So, if we fix $f = \lambda f^*$, then when considering the properties of the game, such as the probability of ruin, we can eliminate a parameter, which will be useful for analysis.

6.2 B. Derivation of Constants

If we let W_n be the number of wins after n plays, then

$$X_n = (1+fb)^{W_n}(1-fa)^{n-W_n}X_0.$$

Then, an alternate characterization of $X_n/X_0 \le r$ is

$$(1+fb)^{W_n}(1-fa)^{n-W_n} \le r.$$

Taking logarithms gives

$$W_n \log(1+fb) + (n-W_n)\log(1-fa) \le log(r)$$

and with some algebraic manipulation,

$$W_n \le \frac{\log(r)}{\log\left(\frac{1+fb}{1-fa}\right)} + \frac{\log\left(\frac{1}{1-fa}\right)}{\log\left(\frac{1+fb}{1-fa}\right)}n.$$

Note that this manipulation requires

$$\log\left(\frac{1+fb}{1-fa}\right) > 0.$$

If $fa < 1$, this reduces to $b > -a$, which is trivially true since both a, b are chosen to be strictly positive. The statement $fa < 1$, on the other hand, is informative. The intuitive meaning is that when we bet we choose a fraction f such that we do not risk losing more than we currently have. This is a very reasonable choice, since if we do not do so, we will encounter ruin with probability 1.

So, if we let

$$C = C_{a,b,r,f} = \frac{\log(r)}{\log\left(\frac{1+fb}{1-fa}\right)} \text{ and } D = D_{a,b,f} = \frac{\log\left(\frac{1}{1-fa}\right)}{\log\left(\frac{1+fb}{1-fa}\right)}$$

we can recharacterize the event of ruin. Ruin upon the n-th play is simply

$$W_n \leq C + Dn.$$

6.3 C. Proof of Probability Mass Preservation

For simplicity, just consider $h(p)$ strictly increasing. Then $\frac{d}{d\lambda}h^{-1}(\lambda) > 0$, so we can leave off the absolute value in the transformation theorem. First, apply the theorem:

$$P(\lambda \leq \lambda_0) = \int_{-\infty}^{\lambda_0} f_\Lambda(\lambda)d\lambda \tag{1}$$

$$= \int_{-\infty}^{\lambda_0} f_\Theta(h^{-1}(\lambda))\frac{d}{d\lambda}h^{-1}(\lambda)d\lambda. \tag{2}$$

Now, perform the substitution $p = h^{-1}(\lambda)$. Then $dp = \frac{d}{d\lambda}h^{-1}(\lambda)d\lambda$.

$$P(\lambda \leq \lambda_0) = \int_{-\infty}^{h^{-1}(\lambda_0)} f_\Theta(p)dp$$

$$= P(p \leq h^{-1}(\lambda_0)).$$

Of course, this is entirely natural. The transformation theorem is constructed precisely so that probability mass is preserved under the transformation.