Gambler's Ruin Revisited: The Effects of Skew and Large Jackpots

R. Michael Canjar

Department of Mathematics and Computer Science, University of Detroit Mercy

Abstract. The mathematical problem of determining a gambler's risk of ruin has a classical solution, which is the root of a certain equation. However, except for "coin-tossing" games, that solution cannot generally be expressed in closed form. Traditional analysis has estimated the RoR using a formula based upon a normal approximation:

$$\mathrm{RoR} = \exp(-2\mu B/\sigma^2).$$

Here μ and σ^2 are the game's mean and variance and B is the gambler's bankroll. However, this formula performs poorly in games in which there is large skew, as in games in which the player's advantage stems from a large jackpot.

In this paper we discuss the underlying mathematical theory of gambler's ruin with a special emphasis on "machine-like" games. We present the Taylor series that expresses the risk of ruin in terms of the moments of the probability distribution associated with the game. We then show that a variant of the standard formula, obtained by replacing the variance with the second moment, is conservative for this type of game. That is, it overestimates the RoR. Moreover, we show that the approximation is accurate for games with moderate skew, and we present a formula that adjusts the normal approximation for moderate skew.

We conclude by presenting closed-form approximations for estimating the risk of ruin in games in which there is a large jackpot. We illustrate the applicability of these formulas with examples from video poker.

1 Introduction

A gambler's risk of ruin (RoR) is the probability that a gambler will lose his/her bankroll if he/she repeatedly wagers on a specified game. There is a classical solution to this problem, which was developed by Laplace, De Moivre, Lagrange, and Bernoulli. Unfortunately, the key equation does not generally admit a closed-form solution, except for the important special case of a coin-tossing problem. For some time gamblers have been estimating the risk of ruin with this formula from Brownian motion:

$$\rho = \exp\left(\frac{-2\mu B}{\sigma^2}\right). \tag{1}$$

Here μ and σ^2 are the game's mean and variance and B is the gambler's bankroll.

This formula is based upon a continuous model, where the position of the particle follows the normal distribution. The formula was adapted to blackjack by Sileo (1992) where it is justified by modeling the game based upon an "equivalent coin-tossing game." This method goes back to Griffin (1981).

Now I have always had some uneasiness about this method. First, it really does not indicate the circumstances under which a "coin-tossing equivalent" is an adequate approximation for the game. Nor does it indicate how good the approximation is. Finally, it is curious that the equation above, based upon the continuous approximation, is actually "wrong" for coin-tossing games, where $\exp(-2\mu B/\,\mathrm{SS})$ is a better approximation, SS being the sum of squares $\sigma^2 + \mu^2$.

In 1999 on the mathematics of blackjack discussion group, Evgeny Sorokin brought to our attention the fact that this approximation fails very badly in games in which there is high skew. Indeed, let us consider a game in which there is a huge bonus payoff, or jackpot, J. It is clear that if we increase the size of J, our RoR decreases. However $\mu B/\sigma^2$ decreases, so the approximation suggests a higher RoR! This suggests that the gambler adopt the ridiculous strategy of refusing the bonus payoff whenever it is offered. (In Appendix 1, I have provided graphs, Figures A1–A3, that illustrate this phenomenon, which I refer to as the "Sorokin effect.") Sorokin then went on to solve the problem for certain games of this type. For these games he independently derived the classical formula (3) given below. A good description of his work is found in Dunbar and B. (1999).

These insights revived my interest in the RoR problem and led to this study. Currently there seems to be a disconnect between the gambling literature and the mathematical literature. Much of the mathematical work is not discussed in the gambling literature and appears to be unknown to many of the authorities. However some of the results in the gambling literature, such as (1), cannot be easily found in the mathematical literature. Part of the purpose of the paper is to help bridge the gap between these. Of course, that is a secondary purpose; my primary purpose is to present new results concerning skewed games.

The first sections concern the classical mathematical theory of gambler's ruin. In that theory, we approximate the key ruin equation with a Taylor series ((8) below) whose terms consist of the moments of the game. Eq. (1) follows immediately from this approximation. Moreover, from this it is clear that our approximation is not valid when there is high skew. From it, we can develop a skew correction to (1), which is also presented below.

Unfortunately, even the moderate-skew correction breaks down when there is very large skew, as when there is a huge jackpot. Because there are many positive-expectation games of this type, I spent some time trying to develop "closed-form" approximations for these types of games. I have been partly successful in this. This discussion is toward the end of this paper, in Section 7. The methods that I use may be called "infinite-jackpot techniques." That is, we study the simpler equation that we would have for an infinite jackpot and then we correct that equation for large finite jackpots.

I had another interesting goal in this study and that was to show that (1) at least gives a conservative estimate of the risk of ruin, for games that are generally available to the public. By a conservative estimate, I mean that the actual risk of ruin is less than that given by the equation. Now as it stands this cannot be quite true, in that coin tossing games have slightly higher risk of ruin than given by (1). However, if we replace σ^2 by SS in (1), the resulting formula is a conservative estimate for coin-tossing games, and for many other games as well. For want of a better name, I am calling such games "well behaved."

Note that not all games are well behaved. Games with negative skew cannot be well behaved; for these games (1) underestimates the risk of ruin. A positive third moment is a necessary but not sufficient condition for a game to be well behaved. However, I have proven that "machine-like" games (such as video poker, slot machines, and lotteries) are well behaved when they have positive expectation. The proof of this is in Section 5. It is fairly technical.

Since I wrote the first draft of this paper, I became aware of the excellent work of Ethier and Khoshnevisan (2002). Some of their work overlaps and is parallel to mine. They independently discovered the Taylor series that relates risk of ruin to the moments of the underlying probability distribution, and they derived the estimate involving the third moment that is presented below in Section 6. But they went further and obtained an additional result that utilizes the fourth moment. These give tighter bounds on the risk of ruin. They also shed more light on the issue of "well-behaved" games mentioned above.

While parts of the paper are very technical, there are some aspects of our theory that have practical application to games such as video poker. Some of these applications are considered in the final sections.

2 Notation and terminology

We will consider the situation where a gambler makes repeated bets on a certain game G. The game has payoffs $X_0, X_1, X_2, \ldots$ and associated probabilities $p_0, p_1, p_2, \ldots$. We will consider the game to be nontrivial if there is at least one positive and one negative X; otherwise G does not represent gambling. We will write M_k for the kth moment: $\sum pX^k$. The first moment M_1 is called the mean and will be denoted by μ. The second moment M_2 will sometimes be denoted by SS, for sum of squares. We write Var for the variance SS $-\mu^2$ and use σ for its square root, the standard deviation. Note that the skew associated with G is $(M_3 - 3\mu\,\mathrm{SS} + 2\mu^3)/\sigma^3$. Finally we call the game G positive if $\mu > 0$ and negative if $\mu < 0$.

3 The classical theory

The classical theory is best understood by considering the situation where the gambler has a win target T, and will stop upon reaching that target. The risk of ruin in this circumstance is given by

$$\rho = \frac{\lambda^{T+B} - \lambda^B}{\lambda^{T+B} - 1}, \tag{2}$$

where λ is the nontrivial solution of

$$P(\lambda) = \sum p\lambda^X - 1 = 0. \tag{3}$$

Note that $\lambda = 1$ is always a zero of P; the other solution is required for (2). Note that if $\lambda < 1$, which we will see occurs precisely if the game is positive, then ρ will converge to λ^B as $T \to \infty$. This gives us the formula for risk of ruin for indefinite play. It is easy to solve (3) for the special case in which G is a coin-tossing game, where the only payoffs are $+1$ with probability p and -1 with probability q. Here the solution is $\lambda = q/p$.

Many probability books contain a development of this theory for the "coin-tossing" problem and this equation is fairly well known. The more general equation, and its proof, is not as well known. It is due to De Moivre (1712), and it involves a rather neat trick. It is so clever that it bears repeating. We imagine painting the numbers $\lambda, \lambda^2, \lambda^3, \ldots, \lambda^B$ on the gambler's chips, and we paint the house's chips $\lambda^{B+1}, \lambda^{B+2}, \ldots, \lambda^{B+T}$. We have the dealer carefully stack the chips, so that at any given stage the gambler wagers λ^n against the house's λ^{n+1}. Note that if the player were to receive a 3 to 1 win, she would be paid $\lambda^{n+1}, \lambda^{n+2}, \lambda^{n+3}$; if the payoff were -2 she would give the house λ^{n-1}, λ^n. We choose λ so that the game has zero expectation in terms of the "funny money" that is painted on the chips. If we write out this requirement, we see that what is needed is precisely the nontrivial zero of $P(\lambda)$. Now, for a zero-expectation game, the probability of ruin is easy to compute: It is $\rho = V(T)/[V(B) + V(T)]$, where $V(B)$ is the value of the bankroll and $V(T)$ is the value of the win target. Here we compute these values by adding up the values that have been painted on the chips. These form a finite geometric series, which, upon simplification, gives (2).

There is a minor point about (2) that deserves some attention, but unfortunately has not received much in the literature on advantage gambling. I believe this was first raised in Coolidge (1909). That concerns the parameters B and T. These have been described as the bankroll B and win target T. However in this equation they should actually be called the "effective bankroll" and "effective win target." The problem can be best illustrated in the context of an example:

A gambler has a bank of 100 units and plays a game where 2 units are risked for an opportunity to win 5. When does the sequence of plays terminate? Obviously if the player loses exactly 100 units, she is ruined. But what if the loss is 99 units? Now the player cannot afford to wager 2 units. We may consider her ruined at this point. If we adopt that interpretation, we see that the loss when ruined is sometime 99 units, and sometimes 100 units. The effective bank is some value in the interval $[99, 100]$. Of course, we may allow the gambler to play the game with 1 unit, and possibly incur an overdraft. In this case, there is an "effective bank" B' is in the interval $[100, 101]$. We may define the effective bank B' in this way. Average the values of λ^{B^*} over all sequences of wins and losses that end in ruin, where B^* is the total loss in a sequence. Let $\mathrm{E}\lambda^{B^*}$ denote this (conditional) expected value. Then B' satisfies $\lambda^{B'} = \mathrm{E}\lambda^{B^*}$. The effective win target is treated similarly.

There are other approaches to developing (2), based upon linear difference equations. These may be found in Feller (1968). My fondness for De Moivre's approach is that it admits the cleanest treatment of the issue raised in the last paragraph. In addition, it is easier to generalize this argument for more complicated types of games, such as games with irrational or continuous payoff structures.

4 The ruin parameter and moments

While the parameter λ plays the crucial role in (2), it will be more convenient to work with its negative log. Hereafter we will write α for $-\ln\lambda$ and call α the *ruin parameter* (rp) for the game G. Note that our risk of ruin, λ^B is now is now given by

$$\rho = \exp(-\alpha B). \tag{4}$$

Our ruin parameter represents what an advantage player may call a "permitted betting fraction." It is inversely proportional to what is considered the "required bank."

Let me elaborate on this. When given an opportunity to make a wager with positive expectation, a player would obviously maximize her expectation by betting an amount as large as possible. However, larger bets obviously increase the risk of ruin, eventually to an undesirable level. Therefore advantage players typically determine their bet sizes so that they maintain some prescribed risk of ruin.

For example, suppose that the player determines that a 5% risk of ruin is acceptable. The player needs $\alpha B \geq -\ln(0.05) \approx 3$. This means that for any wager the player's bank should be at least $3/\alpha$ bets. If the bet sizes were fixed by the casino, this would tell the player how much bank would be required in order to stay within a 5% risk of ruin. On the hand, if the player may select the wager size, this indicates what it should be. The bank B should be at least $3/\alpha$ bets, so each bet should be at most $(\alpha/3)B$. Here $\alpha/3$ represents the fraction of the bank that we are permitted to bet, if we wish to stay within a 5% RoR.

If the player wished to keep the risk of ruin below 1%, then the permitted betting fraction would be $\alpha/[-\ln(0.01)] \approx 0.22\alpha$ and the required bank is the reciprocal of that, or about $4.6/\alpha$.

There are several reasons why I prefer to work with α and not with λ:

1. The ruin parameter has simpler dimensions, namely reciprocal bankroll units. That is, if the bankroll is measured in dollars ($), α has dimensions of $\$^{-1}$.
2. The rp gives a clearer picture of the actual risk of ruin. For example, if game 1 has a λ_1 of 0.999, and game 2 has a λ_2 of 0.9995, it may appear that these games are relatively similar. In reality, the corresponding αs are approximately 0.001 and 0.0005, indicating that twice as much bank is needed for game 2 as for game 1.
3. The domain of α is the entire real line $(-\infty, \infty)$, whereas as the domain of λ is $(0, \infty)$.
4. The rp α appears naturally in the mathematical theory and equations that we will develop below.

Rewriting (3), we note that α must satisfy

$$g(\alpha) = \sum p e^{-X\alpha} - 1 = 0. \tag{5}$$

We will call this the *ruin equation*, and call g the *ruin function*. (Since we think of g as a function of a generic variable α, we denote the ruin parameter by α_0 from now on.) The function g is actually easier to work with than the previous function P. Note that the derivatives of g are given by

$$g^{(n)}(\alpha) = (-1)^n \sum p X^n e^{-X\alpha}. \tag{6}$$

In particular, g'', $g^{(4)}$, and all the even-order derivatives are always positive. If we evaluate the derivatives at 0, we note that $g'(0) = -\mu$, $g''(0) = \mathrm{SS}$, and

$$g^{(n)}(0) = (-1)^n M_n, \tag{7}$$

which tells us that the Taylor series for g is given by

$$g(\alpha) = -M_1\,\alpha + \frac{1}{2} M_2\,\alpha^2 - \frac{1}{6} M_3\,\alpha^3 + \cdots . \tag{8}$$

We are now prepared to discuss some of the key properties of g. These will be summarized in the three theorems of this section. These are not original theorems, and so we give only brief sketches of their proofs. They codify results that are consistent with our intuition, and should enhance our confidence in the validity of the basic theory.

To aid in visualizing these ruin functions, I have included some graphs in Appendix 2; see Figures A4 and A5. Both of these are done for positive games. The first function was done for a coin-tossing game, which is an example of a symmetric ruin function described below. The second graph is for a game with significant positive skew; in this case it was produced by adding a special jackpot to the first game.

Now we proceed to discuss the general properties of a ruin function g. Since we are assuming the payoffs have mixed signs, the limits of g at both $+\infty$ and $-\infty$ is $+\infty$. Since g'' is always positive, g has exactly one relative (hence absolute) minimum and no more than two zeros. Note that $g(0) = 0$. If $\mu > 0$ then $g'(0) = -\mu < 0$, and g must have another zero α_0 that is positive. In this case, $\lambda = \exp(-\alpha_0) < 1$. Conversely, if $\mu < 0$, then the other zero of g must be negative. If $\mu = 0$, then g has its absolute minimum at $\alpha = 0$ and has no other zeros. We summarize this with

Theorem 1. *If $\mu = 0$, then the ruin function g has only one zero, at $\alpha_0 = 0$ and $\lambda = 1$. If $\mu > 0$, then g has a unique nontrivial solution $\alpha_0 > 0$, in which case $\lambda < 1$. If $\mu < 0$, then g has a unique nontrivial solution $\alpha_0 < 0$, in which case $\lambda > 1$.*

Of course, this result is very consistent with our intuition, that the RoR for indefinite play is less than 1 for advantaged players, and only those players. There are other intuitively clear properties of the ruin parameter that we may formally establish with the help of the ruin equation. The first of these concerns compound games (Epstein 1995, p. 67). Suppose that the player plays a variety of games $G_1, G_2, \ldots, G_n$ and scales her bets so as to maintain a certain ruin parameter. We may view these as one large game G, in which $G_1, G_2, G_3, \ldots, G_n$ are sub-games. The ruin parameter for G will be the same as in the sub-games.

Theorem 2. *If a game G consists entirely of sub-games $G_1, G_2, \ldots, G_n$ that have ruin parameter $< (\leq)\ \alpha$, then the ruin parameter for G is $< (\leq)\ \alpha$.*

To prove the theorem, simply note that the ruin function for G is a linear combination of the ruin functions of the sub-games. This can be seen by simply writing out the ruin functions for the compound game, the sub-games, and performing some algebraic simplifications.

This result emphasizes the importance of performing a ruin analysis for all games that are played. If a gambler plays a wide variety of games for very brief periods, it is of course unlikely that he will be ruined in any one of them. However the gambler's overall or global risk of ruin may be controlled by controlling the risk of ruin in each of the sub-games.

Sometimes we have two games G and H, which are essentially the same game, except that the payoff structures are different. Examples of this occur in video poker games. G may be the so-called "full pay" version of Jacks or Better, in which a full house pays 9 for 1 and a flush pays 6 for 1 (referred to as a 9/6 game). H may be the same game with the less favorable payoff structure in which these two payoffs are shorted to 8 and 5 respectively (referred to as an 8/5 game). Or the jackpot (royal flush payoff) for H may be larger than for G. Formally, we will say that H has a *more favorable* payoff structure than G if the events for G and H are the same, and the payoffs satisfy $X_k^H \geq X_k^G$ for each k. We will say it is *strictly more favorable* if, in addition, there is at least one event with nonzero probability $p_k > 0$, for which $X_k^H > X_k^G$. Intuitively our ruin prospects should be better for H than for G, and that is the content of the next theorem. As noted in the preface, a serious flaw in the approximation in (1) is that this result does not hold for the approximate values it suggests; as the payoffs improve the increase in σ^2 overcomes the increase in μ.

Theorem 3. *If H has a more favorable payoff structure than G, then their ruin parameters satisfy $\alpha_H \geq \alpha_G$. Moreover, if H is strictly more favorable, then $\alpha_H > \alpha_G$.*

For positive games, this theorem follows easily from the fact that the respective ruin functions satisfy $h(\alpha) \leq g(\alpha)$ for $\alpha > 0$, where h is the ruin function for H and g is the ruin function for G. Moreover we have $h(\alpha) < g(\alpha)$ when H has a strictly more favorable payoff structure. From this we conclude that $h(\alpha_G) \leq (<)\ 0$. Since h is negative only on the interval $(0, \alpha_H)$, we have $\alpha_G < \alpha_H$. While the case where both games are negative may be of less interest, it may be treated by a similar argument. The case where H is positive and G is negative is trivial.

It is also easy to generalize our discussion to the case where we have a continuous game in which the payoffs X are continuous. Here we may replace the g of (5) with the more general

$$g(\alpha) = \int e^{-X\alpha}\, dP - 1. \tag{9}$$

It is interesting to apply (8) to the case where the payoff function X has the Gaussian or normal distribution. We will call such games *normal*. By direct computation we may obtain for normal games:

$$g(\alpha) = \exp(\alpha^2\sigma^2/2 - \alpha\mu) - 1, \tag{10}$$

which admits the solution $\alpha_0 = 2\mu/\sigma^2$ mentioned in the introduction (eq. (1)). Note that this is not the formula for a discrete coin-tossing game, but rather is one where the participants can wager on the value of a continuous variable.

Eq. (10) reveals one additional property about normal games. If we let α_0 be the ruin parameter $(2\mu/\sigma^2)$, we see that the function g is symmetric about $\alpha = \alpha_0/2$. That is, $g(\alpha) = g(\alpha_0 - \alpha)$. We will call a ruin function with this property *symmetric* and we will call a game *symmetric* if its ruin function is symmetric.

Any ruin function will achieve an absolute minimum at a vertex, whose α-coordinate we will call α_v. α_v is the unique zero of g'. A symmetric ruin function will be symmetric about α_v which will necessarily equal $\alpha_0/2$.

Obviously, normal games are symmetric. Coin-tossing games are also symmetric. In such a game, we win 1 unit with probability p and lose 1 unit with probability q. The ruin function for this game is given by

$$g(\alpha) = pe^{-\alpha} + qe^{\alpha} - 1. \tag{11}$$

The ruin equation has solution $\alpha_0 = \ln(p/q)$. The symmetry of g may be easily seen if one replaces p with qe^{α_0} in g.

Note that, for symmetric ruin functions,

$$g^{(n)}(\alpha_0) = (-1)^n g^{(n)}(0) = M_n. \tag{12}$$

Of course, not all games are symmetric. For example, for a game to be symmetric it is necessary that all its odd moments have the same algebraic sign. To see this, suppose that n is odd. Since $g^{(n+1)}$ is positive, $g^{(n)}$ is increasing, hence when $\alpha_0 > 0$, we have $g^{(n)}(0) < g^{(n)}(\alpha_0) = -g^{(n)}(0)$ by (12), and therefore $M_n = -g^{(n)}(0) > 0$. Conversely, if $\alpha_0 < 0$, all the odd moments are negative. If $\alpha_0 = 0$, then they are all 0.

However a compound game will be symmetric if all of its sub-games are symmetric and have the same ruin parameter. For an example of such a game, consider a game in which we make different sized bets on the outcome of a biased coin toss. If the same coin, or rather the same bias, is used for all the bets, then the resulting compound game will be symmetric.

The vertex plays an important role in Epstein's (1995, p. 67) "criterion of survival." This was proposed as a betting strategy for a compound game. Here the game consists of several sub-games $G_1, G_2, \ldots, G_n$, at least one of which is positive. In each sub-game, the player chooses an amount to bet, which we will designate as $X_1, X_2, \ldots, X_n$. The goal is to minimize the overall risk of ruin. If there are no constraints on the bet sizes $X_1, X_2, \ldots, X_n$, then this is an indeterminate problem. One constraint considered by Epstein was that the player was required to wager at least one unit in each sub-game, and that at least one of the sub-games has negative expectation.

Although he was motivated by the game of blackjack, Epstein actually considered the problem where each of the sub-games involves coin tossing. Epstein showed that, in the optimal solution, the sizes of the unconstrained bets would be proportional to $\ln(p_k/q_k)$, which is, of course, the ruin parameter for the game G_k. One might guess that this would be the general result, but that turns out to be slightly different. In the general case, the unconstrained bets are proportional to $[\alpha_v]_k$, the vertex of the ruin function g_k for G_k.

To see this, let $f_0, f_1, \ldots, f_n$ be the frequencies for each of the sub-games. Let g be the ruin function for the compound game G. The ruin equation may be expressed as

$$g(\alpha) = \sum f_k g_k(X_k \alpha) = 0. \tag{13}$$

We now view α as an implicitly defined function of the $X_1, X_2, \ldots, X_n$. We wish to maximize α. We may compute the derivatives of α by implicit differentiation; for the unconstrained bets, these must be 0 when we have an optimal solution. Toward this end, let $S(X_1, X_2, \ldots, X_n, \alpha)$ be the expression in (13), so that $g(\alpha) = S(X_1, X_2, \ldots, X_n, \alpha)$. Then $\partial\alpha/\partial X_k = -(\partial S/\partial X_k)/(\partial S/\partial\alpha)$. For this to be 0, the numerator $\partial S/\partial X_k$ must be 0, so we obtain

$$\frac{\partial S}{\partial X_k} = f_k g_k'(X_k \alpha)\alpha = 0. \tag{14}$$

Since $\alpha > 0$ for the optimal solution, we must have $g_k'(X_k\alpha) = 0$. But then $X_k\alpha$ must be the vertex of the function g_k. Hence for each unconstrained bet X_k,

$$X_k = [\alpha_v]_k/\alpha, \tag{15}$$

where the $[\alpha_v]_k$ is the vertex for g_k. Here α is the solution of (13). The key point is that it is a constant relative to k so that the optimal value of the unconstrained bets will be proportional their corresponding vertices.

As noted above, often the approximation $\alpha = 2\mu/\operatorname{Var}$ is used to estimate the ruin parameter. This is the exact answer for a normal game, but it slightly overestimates α for coin-tossing games. A better estimate for those games is $\alpha_1 = 2\mu/\mathrm{SS}$. For these games, this estimate is conservative in that it underestimates α, which means that it overstates the risk of ruin. We will show this is true for a large class of games. We will call such games *well behaved*. Hereafter we will use α_0 to represent the exact zero of the ruin equation, and α_1 for the approximation $\alpha_1 = 2\mu/\mathrm{SS}$. Games are well behaved if $\alpha_1 \le \alpha_0$.

We obtain the approximation α_1 by solving the quadratic approximation for the ruin function, effectively ignoring the higher moments. Note however that if the third moment is negative, our quadratic will overestimate α. Because the third derivative g''' is increasing and $g'''(0) = -M_3$, $g'''(\alpha) > 0$ for all positive values of α. Taylor's theorem tells us that the quadratic approximation to g will be less than g, from which is concluded that $\alpha_1 > \alpha_0$. That is, *a game with negative third moment can never be well behaved*. Indeed, we will see that with highly negative skew, the actual ruin parameter will be lower than α_1.

We might hope that a positive third moment would be sufficient to give us a well-behaved game. However that is not the case. It is possible to construct payoff schedules with positive M_3, but for which M_4 is very large as well. The high M_4 will lower g and push the ruin parameter to the left. However these are somewhat artificial examples. Games that are offered to the public generally do not have negative skew or other unusual payoff structures. In the next section we will see that most casino games are well behaved.

5 Well-behaved games

In this section we will identify two classes of games and show that they are well behaved. One such class has already been mentioned: symmetric games. Another such class are games that are like slot or video machines. Consider a game with the following payoffs: $\{-1, 0, 1, 10, 20\}$. Here players enter bets of 1 unit. They either lose that unit, or they push, or they win an amount greater than or equal to 1 unit. This is an example of a type of game that we will call machine-like. Formally, we will say that a game is *machine-like* if the payoffs have the following structure: There is a numerical value S such that all the positive payoffs X have magnitude $|X| \geq S$, and the negative payoffs L have magnitude $|L| \leq S$. S may be thought of as the size of the bet. In the earlier example, S was 1.

Note that, by our definition, a game with the following payoff set is also machine-like: $\{-2, -1, 0, 2, 2.5, 3\}$. Here $S = 2$.

Theorem 4. *Positive symmetric games and positive machine-like games are well behaved (i.e.,* $\alpha_1 \leq \alpha_0$*).*

There are other casino games in which the player may have positive expectation, such as blackjack, when a counting strategy is used. Indeed blackjack players have been using (1) as an approximation for some time (Sileo 1992, Schlesinger 2005). Is this a conservative estimate? Is blackjack well behaved?

It seems intuitively clear that if the ruin parameter α_0 is "small," then the contributions of the fourth and higher moments will be negligible. Then the game would be well behaved iff $M_3 > 0$. Recently, S. Ethier informed me that the inequalities established in Ethier and Khoshnevisan (2002) may be used to confirm this. Let ν be the maximum possible loss that the player may suffer. Then a *sufficient* condition for a game to be well behaved is that $\mu/\,\mathrm{SS} \leq 1/\nu$ and $M_3 > 0$.

To illustrate, in blackjack with standard rules a player may sometimes increase her bet 8 times, by splitting to 4 hands and doubling on each. She may also take, and lose, insurance, which is another half bet. So the maximum possible loss is 8.5 times MaxB, the maximum bet. Now among blackjack players, the expression $\mathrm{SS}\,/\mu$ is sometimes called the "Kelly-equivalent bank," which I will abbreviate below as KEB. Since blackjack is positively skewed, we see that a sufficient condition for it to be well behaved is that $\mathrm{MaxB} \leq \mathrm{KEB}\,/8.5$. Blackjack players almost never bet that large, and so in practice, we can be assured that BJ is well behaved. Indeed, a loss of 8.5 bets is an extremely rare occurrence (it would require losing an insurance bet to a dealer's ace, splitting 8s four times, drawing a 2 or 3 to each, doubling down, and losing all four doubled hands) and so it is possible to justify even higher estimates.

Typical casino games will have $M_3 > 0$ and will also have $\mu/\,\mathrm{SS}$ "small" in the sense described above. Hence these games are almost invariably well behaved.

5.1 Technical proofs

We will devote the remainder of this section to the proof of Theorem 4. While the proof requires only knowledge of elementary calculus, it is rather technical. However, readers who are not interested may skip the rest of this section; none of the material herein will be used in the later sections.

Let α_v be the vertex on the graph of the ruin function g, so that $g'(\alpha_v) = 0$. Form the function Asym defined by

$$\operatorname{Asym}(h) = g(\alpha_v - h) - g(\alpha_v + h). \tag{16}$$

Note that for symmetric games $\operatorname{Asym}(h) = 0$ for all h. Think of Asym as measuring the asymmetry of the game. We will see that for all of the games under discussion, $\operatorname{Asym}(h) \geq 0$ whenever $h \geq 0$. Indeed, they actually enjoy the stronger property:

$$\operatorname{Asym}'(h) = -[g'(\alpha_v - h) + g'(\alpha_v + h)] \geq 0 \quad \text{for} \quad h \geq 0. \tag{17}$$

We will call this *nonnegative asymmetry* (*nna*). To understand the meaning of nna, note that $g'(\alpha_v - h)$ is negative and is the slope of the graph of our ruin function on the left of the vertex, while $g'(\alpha_v + h)$ is the slope at a corresponding point on the right side of the vertex. A function has nonnegative asymmetry if the magnitude of the former exceeds that of the latter; that is, if its rate of descent into the vertex is greater than its rate of climb out of the vertex, measured at points that are equidistant from the vertex.

Nonnegative asymmetry is a fairly strong property, and it has the following consequences. If a game has nna, then

1. $g'''(\alpha_v) \leq 0$.
2. $M_3 \geq 0$.
3. $\alpha_v \geq \alpha_1/2 = \mu/\operatorname{SS}$.
4. $2\alpha_v \leq \alpha_0$.
5. $\alpha_1 \leq \alpha_0$.

To establish these, we first compute the derivatives of Asym and we see that $\operatorname{Asym}^{(n)}(0) = 0$ when n is even and $\operatorname{Asym}^{(n)}(0) = -2g^{(n)}(\alpha_v)$ when n is odd. Since α_v is the vertex, we have that $\operatorname{Asym}'(0) = -2g'(\alpha_v) = 0$, and $\operatorname{Asym}''(0) = 0$. Thus the third derivative $\operatorname{Asym}'''(0)$ must be nonnegative. If it were negative, then Asym would be negative on some interval $(0, \delta)$, contradicting the definition of nna. But $g'''(\alpha_v) = -\operatorname{Asym}'''(0)/2$, so condition 1 must hold. Condition 2 follows immediately, inasmuch as g''' is increasing. Since it is nonpositive at α_v, it must be nonpositive at 0. But $g'''(0) = -M_3$.

Now consider the first-degree Taylor polynomial for $g'(\alpha)$. We have

$$g'(\alpha) = -\mu + \operatorname{SS}\alpha + g'''(\zeta)\alpha^2/2. \tag{18}$$

The last term involves the value of the third derivative of g at an intermediate point between 0 and α. But throughout the interval $(0, \alpha_v)$, $g''' \leq 0$. Hence the remainder is negative. But $g'(\alpha_v) = 0$, from which we obtain

$$0 = g'(\alpha_v) = -\mu + \mathrm{SS}\,\alpha_v + R(\alpha_v) \leq -\mu + \mathrm{SS}\,\alpha_v, \tag{19}$$

where $R(\alpha_v) \leq 0$. Solving we obtain $\alpha_v \geq \mu/\,\mathrm{SS} = \alpha_1/2$, establishing condition 3.

Moreover, because $\mathrm{Asym}(0) = 0$ and $\mathrm{Asym}'(h) \geq 0$ for $h \geq 0$, we may conclude that $\mathrm{Asym}(h) \geq 0$ whenever $h \geq 0$. In particular $\mathrm{Asym}(\alpha_v) = -g(2\alpha_v) + g(0) \geq 0$. But $g(0) = 0$, so $g(2\alpha_v) \leq 0$. From this we conclude that $2\alpha_v$ lies to the left of α_0, which gives us 4. But we know from 3 that $2\alpha_v \geq \alpha_1$, which gives us 5 and tells us that G is well behaved.

Thus we see that nna guarantees that a game is well behaved. We will now show that the games described above satisfy nna. First, we rewrite the expression $\mathrm{Asym}'(h)$. We take the derivative of g (eq. (6)) and substitute in $\alpha_v + h$ and $\alpha_v - h$. We simplify by removing a common factor of $\exp(\alpha_v)$. Finally, we will write $T(z)$ for $[\exp(z)+\exp(-z)]$, which is also $2\cosh(z)$. With this we obtain:

$$\begin{aligned}\mathrm{Asym}'(h) &= -\sum p(-X)\exp(-X[\alpha_v + h]) - \sum p(-X)\exp(-X[\alpha_v - h])\\ &= \sum pX\exp(-X\alpha_v)[\exp(Xh) + \exp(-Xh)]\\ &= \sum pX\exp(-X\alpha_v)[T(Xh)]. \end{aligned}\tag{20}$$

Note that T is increasing on the positive real axis: if $0 < z_1 < z_2$ then $T(z_1) < T(z_2)$.

Now for a machine-like game, we have negative payoffs $-Y_1, -Y_2, \ldots$ and positive payoffs $X_1, X_2, \ldots$. Recall that we also have a number S so that $Y_k \leq S$ and $X_k \geq S$. For convenience, we use $q_1, q_2, \ldots$ for the probabilities of obtaining $-Y_1, -Y_2, \ldots$ and $p_1, p_2, \ldots$ for the probabilities associated with the $X_1, X_2, \ldots$. Our equation now becomes

$$\mathrm{Asym}'(h) = \sum pX\exp(-X\alpha_v)[T(Xh)] - \sum qY\exp(Y\alpha_v)[T(Yh)]. \tag{21}$$

We will assume that $h \geq 0$. Then for each positive payoff X, $T(Xh) \geq T(Sh)$. Similarly $T(Yh) \leq T(Sh)$, or $-T(Yh) \geq -T(Sh)$. Substituting these into the last equation yields

$$\mathrm{Asym}'(h) \geq \left[\sum pX\exp(-X\alpha_v) - \sum qY\exp(Y\alpha_v)\right] T(Sh). \tag{22}$$

Now the expression within the brackets is just $-g'(\alpha_v) = 0$. Thus we conclude that $\mathrm{Asym}'(h) \geq 0$ whenever $h \geq 0$, or that our game has nonnegative asymmetry.

It is clear that symmetric games exhibit nna.

6 Moderate skew

We will consider positive games, and seek a simple formula to estimate the ruin parameter α_0 and the effect of the third moment on it. Let $C(\alpha)$ be the cubic Taylor polynomial for $g(\alpha)$:

$$C(\alpha) = -M_1\alpha + \frac{1}{2}\,M_2\alpha^2 - \frac{1}{6}\,M_3\alpha^3. \tag{23}$$

Because $g^{(4)}(\alpha) > 0$ for all α, we know that the remainder term is positive, hence $g(\alpha) \geq C(\alpha)$. Now C obviously has a zero at the origin, and may have two other zeros. If $M_3 > 0$ the other zeros, if they exist, would have to be positive. This can be seen geometrically from the graph of g, which is above the horizontal axis when $\alpha < 0$ or by considering the structure of the quadratic formula, discussed below. However, if $M_3 < 0$ there will be two solutions with mixed signs. In either case, let α_c be the smallest positive zero of C, provided one exists. Then $g(\alpha_c) \geq 0$ so $\alpha_0 < \alpha_c$, since g is negative only on the interval $(0, \alpha_c)$.

The zeros of C may be found easily, since we simply solve a quadratic, after we factor out α. We employ the quadratic formula, with $a = M_3/6$, $b = -\,\mathrm{SS}/2$, and $c = \mu$, taking the smallest of the two roots. Our equation will have a solution provided $b^2 \geq 4ac$ or $\mu M_3 \leq (3/8)(\mathrm{SS})^2$. We will say that a game has *moderate skew* if this inequality is satisfied. Note that any positive game with negative M_3 will meet this criterion.

We will simplify the answer we obtain from the quadratic formula. First, we introduce a temporary function T so that

$$\alpha_c = \frac{|b|}{2a}\left[1 - \sqrt{1-(4ac)/b^2}\right] = \frac{|b|}{2a}T(4ac/b^2). \tag{24}$$

Now $4ac/b^2$ is $4\mu(M_3/6)/(\mathrm{SS}/2)^2 = (8/3)\mu M_3/(\mathrm{SS})^2$. Let us write γ for the dimensionless quantity $\mu M_3/(\mathrm{SS})^2$. This parameter γ gives us a measure of the relative skew for our game. Substituting $|b| = \mathrm{SS}/2$ and $a = M_3/6$ into the last equation gives us our estimate

$$\alpha_c = \frac{3\,\mathrm{SS}}{2M_3}T(8\gamma/3). \tag{25}$$

Note that for s close to 0, $T(s) \approx s/2 + s^2/8 + s^3/16$. This motivates the introduction of a new function, which we call $\mathrm{Err}(s)$, since it will ultimately tell us the error in the approximation we are developing. This is defined so that $T(s) = (s/2)[1+\mathrm{Err}(s)]$. That is, $\mathrm{Err}(s) = 2T(s)/s - 1$. We define $\mathrm{Err}(0)$ to be 0, so that Err is a continuous function. We now replace $T(s)$ with $(s/2)[1+\mathrm{Err}(s)]$, and for $s = 8\mu M_3/(3(\mathrm{SS})^2)$ we obtain

$$\alpha_c = \frac{2\mu}{\mathrm{SS}}[1+\mathrm{Err}(8\gamma/3)] = \alpha_1\left[1+\mathrm{Err}\left(\frac{8M_3\mu}{3(\mathrm{SS})^2}\right)\right]. \tag{26}$$

Again, this gives us an upper bound on the exact value of α_0. We have

Theorem 5. *If $\mu > 0$ and $M_3 \leq (3/8)(\mathrm{SS})^2/\mu$, and if α_c is defined via (25), then*

1. $\alpha_0 \leq \alpha_c$

2. For well-behaved games, $\alpha_1 \leq \alpha_0 \leq \alpha_c$.

Remark. As mentioned in the preface, this result was independently discovered by Ethier and Khoshnevisan (2002). They have actually proven a stronger result. We have used the third moment to bound α_0 above, which they also established. However, they also use the fourth moment to obtain a lower bound for α_0.

The restriction on M_3 is necessary in order that our quadratic may be solved. The second part of the theorem tells us that our approximation $\alpha_0 \approx \alpha_1$ is valid for moderate skew. Our Err function quantifies the error in using the approximation. Note that if Err $= 0.1$, then the error in our approximation is less than 10%. Now for small s, $\mathrm{Err}(s) \approx s/4$, so $\mathrm{Err}(8\gamma/3) \approx 2\gamma/3$. Thus we obtain the following estimate, which is valid when M_3 is "relatively small":

$$\alpha_0 \approx \frac{2\mu}{\mathrm{SS}}\left[1 + \frac{2M_3\mu}{3(\mathrm{SS})^2}\right]. \tag{27}$$

To give the reader a feel for the size of this error under other circumstances, we provide a table of representative values:

Table 1. Errors in skew adjustment.

$\gamma = \mu M_3/(\mathrm{SS})^2$	$s = 8\gamma/3$	$\mathrm{Err}(s)$	$\gamma = \mu M_3/(\mathrm{SS})^2$	$s = 8\gamma/3$	$\mathrm{Err}(s)$
0.000	0.000	0.000	0.250	0.667	0.268
0.050	0.133	0.036	0.300	0.800	0.382
0.100	0.267	0.077	0.350	0.933	0.590
0.150	0.400	0.127	0.375	1.000	1.000
0.200	0.533	0.188	> 0.375	> 1.000	n/a

We should also note that these equations indicate the devastating effects that negative skew can have. As noted at the beginning of this section, we always have $\alpha_0 \leq \alpha_c$. But when $M_3 < 0$, $\alpha_c < \alpha_1$. This follows since the cubic function (23) is greater than its quadratic part when $\alpha < 0$, and so the cubic makes it back up to the horizontal axis before the quadratic does. Thus $\alpha_0 \leq \alpha_c \leq \alpha_1$. Our quadratic formula shows us that as $M_3 \to -\infty$, α_c is approximately $\sqrt{-6M_1/M_3}$. It is fairly easy to construct games where this ratio can be made arbitrarily small.

Example 1. We illustrate this formula with an existing game: Deuces Wild video poker. This example was discussed in Dunbar and B. (1999) and Ethier and Khoshnevisan (2002). The data come from Jensen (2001) and are based on a particular optimal drawing strategy; the optimal strategy is not unique in this game.

Table 2. Deuces Wild payoff schedule.

hand	payback	probability $\times\binom{52}{5}\binom{47}{5}5$	payoff
nonwinner	0	10,901,684,089,380	-1
trips	1	5,675,175,572,112	0
straight	2	335,108,390,364	1
flush	2	1,116,335,570,916	1
full house	3	423,165,297,240	2
four-of-a-kind	5	1,294,427,430,576	4
straight flush	9	83,218,233,480	8
five-of-a-kind	15	63,818,309,856	14
royal flush (deuces)	25	35,796,957,696	24
four deuces	200	4,060,462,824	199
royal flush (natural)	800	440,202,756	799
sum		19,933,230,517,200	

We may compute the probability of each outcome by dividing the number of occurrences by the total. Since we are using exact data, our probabilities will sum to exactly 1. However, we must caution the reader on the necessity of this. If approximate data are used, and if the sum of the probabilities does not equal 1, then our equations will associate the residual probability to an infinite jackpot, in which case our game will have infinite expectation. One could still obtain an approximation by adjusting the probabilities so that they sum to exactly 1.

Using these data, we may compute the first three moments and obtain $\mu = 0.007620$, SS $= 25.84$, and $M_3 = 12{,}909$. Now $\alpha_1 = 2(0.007620)/25.84 = 0.000590$. This would be the ruin parameter if we used the conventional coin-tossing equation. This underestimates our ruin parameter, because it ignores skew. Here $\gamma = (0.007620)(12{,}909)/(25.84)^2 = 0.147$. A simple estimate is that Err is $(2/3)(0.147) = 0.098$.

Using (27), we estimate that the ruin parameter is $(0.000590)(1.098) \approx 0.000648$. Using our table, we can note that $\mathrm{Err}(s) \leq 0.127$. From this we establish an upper bound on the rp of $(0.000590)(0.127) \approx 0.000665$. In other words, we know that the exact answer lies in the interval $(0.000590, 0.000665)$.

To see how accurate these methods are, we could obtain the exact solution with a solving utility. Using Excel, I came up with 0.000653.

What do these number mean? The fact that the traditional equation underestimates α_0 by about 10% means that we can actually use a 10% smaller bank than it predicts. For example, the bank required to maintain a 5% risk of ruin is $-\ln(0.5)/(0.000653) = 4586$ units. If the machine requires wagers of \$1.25, this would be \$5,732. The player using the conventional method would think that a larger bank of \$6,347 were required.

The methods of this section are not valid when M_3 is relatively large. This can occur when we have a large jackpot, which is the subject of the next section.

7 Large jackpots

We now consider a special type of positive game H, in which there is a reasonably "normal" game G combined with a special large jackpot J. The probability of getting the jackpot is a very small number, hereafter called ε. The rules of H dictate that with probability $1-\varepsilon$, we obtain the results of G. G will be called the *base game*. We may always find such a base game G for any given H by doing the following: The payoffs for G are the payoffs of H, except for the jackpot J, with their probabilities inflated by a factor of $(1-\varepsilon)^{-1}$. Let g be the ruin function for the game G and let μ_G and SS_G be the parameters of G. Note that the parameters of H are given by $\mu_H = (1-\varepsilon)\mu_G + \varepsilon J$ and $\mathrm{SS}_H = (1-\varepsilon)\,\mathrm{SS}_G + \varepsilon J^2$. We are interested only in the case that H is positive. We allow the possibility that the base game is negative, in which case we assume that J is sufficiently large to overcome the negative expectation of G. Thus, the ruin parameter α_H for H will be positive. The ruin equation for H is given by

$$(1-\varepsilon)g(\alpha) = \varepsilon(1-e^{-J\alpha}). \tag{28}$$

Let $R(\alpha) = 1 - e^{-J\alpha}$, so that the right side of this equation $\varepsilon R(\alpha)$. Note that this is positive for $\alpha > 0$ and will usually be very small. Note that the left side of (28) is zero at both $\alpha = 0$ and at $\alpha = \alpha_G$, the ruin parameter of G. Let α_M be the maximum or right-most zero of g. This is equal to the ruin parameter α_G when G is positive, and equal to 0 otherwise. Since g is positive to the right of α_M, the nontrivial solution to (28) will be found there. Let h_J be the improvement in α produced by J: $h_J = \alpha_H - \alpha_M$. We may use the mean-value theorem and obtain

$$h_J = \frac{\varepsilon R(\alpha_H)}{(1-\varepsilon)g'(\zeta)}, \tag{29}$$

where ζ is some number in (α_M, α_H). But g' is increasing, so we have that

$$h_J \leq \frac{\varepsilon R(\alpha_H)}{(1-\varepsilon)g'(\alpha_M)}. \tag{30}$$

Note that this is the estimate we would obtain by carrying out one step of Newton's method for solving (28). It gives us an upper bound on how much improvement can be obtained. Now when G is negative, $g'(\alpha_M)$ will be $|\mu_G|$; when G is positive, its value will depend on the specific game. For symmetric games, it will also be $|\mu_G|$, but this will not be true for an arbitrary G. We will shortly consider both of these important special cases separately. However the reader should note that as long as ε is small in comparison to $|\mu_G|$, the improvement h_J will be relatively small.

Before proceeding to our special cases, there is one other important general observation to be made. Consider the limiting case as $J \to \infty$. Let α_∞ be the limiting value of the α_H. It is easy to check that α_∞ satisfies the following equation, which comes from replacing $e^{-J\alpha}$ with 0 in (28):

$$g(\alpha_\infty) = \frac{\varepsilon}{1-\varepsilon}. \tag{31}$$

This gives us the ruin equation for an "infinite jackpot." An infinite jackpot should be thought of as being so large that once it is won by a gambler, there is no possibility that the gambler will be subsequently ruined. Thus, an infinite jackpot is equivalent to the gambler's ruin problem for the following scenario: There is a (rare) event of probability ε whose occurrence terminates the sequence of plays, in a state in which the gambler is considered as "nonruined." Computations based on an infinite jackpot will give our probability of ruin in the following real situation: *We have an ordinary game, with a finite jackpot, and we wish to know the probability that we will be ruined before hitting the jackpot once.*

Note that even with an infinite jackpot, we still have a positive risk of ruin, and a finite ruin parameter. In fact we always have that

$$\alpha_\infty - \alpha_M \leq \frac{\varepsilon}{(1-\varepsilon)g'(\alpha_M)}, \tag{32}$$

which gives us an upper bound on the improvement available from J.

7.1 Normally based games: positive-based case

Let us consider the case where the base game is normally distributed, so that its ruin function is given by (10). These games are a mathematical abstraction, as casino games usually have discrete payoffs which cannot have the normal distribution. However, we use this model as an approximation to many casino games. These are games where the traditional mean-variance approach gives us a good approximation of the game without the jackpot.

We may rewrite (10) as

$$g(\alpha) = \exp\{Q(\alpha)\} - 1 \quad \text{where} \quad Q(\alpha) = \alpha(\alpha \operatorname{Var}_G/2 - \mu). \tag{33}$$

We now use the approximation $g(\alpha) \approx Q(\alpha)$. In our situation where the base game is normal, this can be justified by the approximation $\exp(x) - 1 \approx x$. More generally, we are approximating $g(\alpha)$ by the first two terms of its Taylor series, and treating Var as approximately M_2. (Alternatively, we can expand $\ln(1 + g(\alpha))$ in a Taylor series.) Our ruin equation for the compound game H may now be approximated by

$$\pm\mu_0\alpha + \frac{v_0}{2}\,\alpha^2 = \varepsilon(1 - e^{-J\alpha}). \tag{34}$$

Here we let $\mu_0 = (1-\varepsilon)|\mu_G|$ and $v_0 = (1-\varepsilon)\operatorname{Var}_G$. These represent the contribution that the sub-game G makes to the overall mean and to the overall sum of squares for H; note our convention that $\mu_0 > 0$. The cases where the base game is positive and where the base game is negative are sufficiently different as to merit separate analysis. The sign of μ_0 will be negative when the base game is positive, and vice-versa.

We will first consider the case where the base game is positive. For the interesting case where the jackpot is infinite, our equation becomes a quadratic, which we may solve explicitly. Only one of the roots is positive:

$$\alpha_\infty = \frac{\mu_0 + \sqrt{\mu_0^2 + 2v_0\varepsilon}}{v_0}. \tag{35}$$

Rewriting this we obtain

$$\alpha_\infty = \frac{2\mu_0}{v_0}\left[\frac{1+\sqrt{1+2v_0\varepsilon/\mu_0^2}}{2}\right]. \tag{36}$$

Among blackjack players, the term N_0 has been used to represent v/μ^2. This term was coined by Brett Harris (1997) as a measure of how many plays it would take a game to "get into the long run." It is the square of σ/μ, which is known as the coefficient of variation. Using this notation, we obtain

$$\alpha_\infty = \frac{2\mu_0}{v_0}\left[\frac{1+\sqrt{1+2\varepsilon N_0}}{2}\right]. \tag{37}$$

The expression εN_0 measures how frequently the jackpot occurs relative to the time it takes for the base game to get into the long run. The factor $2\mu_0/v_0$ is the ruin parameter we would have if there were no jackpot. That value is being increased by a factor that depends only on the frequency of the jackpot before getting into the long run. As we would expect, when εN_0 is very small, there is almost no increase in the ruin parameter.

We clarify this equation with an example:

Example 2. A blackjack player is going to play in a promotion where there is a special jackpot associated with a rare combination of cards. This jackpot occurs once every 5000 rounds. Using her count strategy, the player estimates that her expected win is 0.01 units per round, with a variance of 2 units each round. She has brought only 50 units with her; she wants to know the risk that she will be ruined before hitting this jackpot.

Without the jackpot, her ruin parameter would be $2(0.01)/2 = 0.01$. $N_0 = 2/(0.01)^2 = 2\times 10^4$. $\varepsilon = 2\times 10^{-4}$ so $2\varepsilon N_0 = 8$. Thus, her ruin parameter for an infinite jackpot would be $(0.01)[1+\sqrt{1+8}] = 0.02$. With her 50-unit bankroll, the risk that she will be ruined before hitting the jackpot is $\exp(-0.02\cdot 50) = 37.8\%$, if she plays indefinitely.

Our player can calculate that, if she wishes to lower this risk to 5%, she will need to bring $-\ln(0.05)/(0.02)$ or about 150 units.

7.2 Negative-based case: applications to video poker

We now consider the interesting special case when the sub-game G is negative. Of course, we are assuming that the jackpot is large enough that H is a positive game; this means that $J > (1-\varepsilon)|\mu_G|/\varepsilon$.

Video poker is a good example of the type of game we are considering. Certain varieties of VP do offer the knowledgeable player a positive expectation. However, this is always due to the large payoff for a rare event, such as a royal flush. Without this large jackpot, the game would be negative. Even without the jackpot, VP is not normally distributed, but has some positive skew. This means that our model should actually be conservative in estimating the RoR for video poker.

Recall our convention that $\mu_0 = (1-\varepsilon)|\mu_G| > 0$. That is, here μ_0 represents the advantage the house would have, if there were no jackpot. Thus our equation becomes

$$+\mu_0\alpha + \frac{v_0}{2}\,\alpha^2 = \varepsilon(1 - e^{-J\alpha}). \tag{38}$$

For the limiting case where J is ∞, we again obtain a quadratic, which we may solve explicitly. Again we consider only the positive root. After simplifying, by rationalizing the numerator, we obtain

$$\alpha_\infty = \frac{2\varepsilon}{\mu_0 + \sqrt{\mu_0^2 + 2\varepsilon v_0}}. \tag{39}$$

Note that in an extreme situation, where εv is relatively small (compared to μ^2) we have the formula $\mathrm{RoR} = \exp(-\varepsilon B/\mu_0)$, which may be derived directly through simple probabilistic considerations. The formula shows how much the variance of the base game increases the RoR.

We might wish to be able to tabulate solutions to (38). That equation is quite formidable in the terms of the number of variables it involves, but we may simplify it by introducing some new, dimensionless variables. Let $J_0 = \mu_0/\varepsilon$, which represents the break-even jackpot; $J = J_0$ means that the overall advantage is precisely 0. We introduce the dimensionless parameter r:

$$r = \frac{J - J_0}{J_0} = \frac{\varepsilon J}{\mu_0} - 1. \tag{40}$$

This parameter r is proportional to the overall expected value. In the following equations, we will write r' for $r + 1 = \varepsilon J/\mu_0$. It is easier to write the equation in terms of r', but we later use r because then our solutions will contain the origin $\alpha = 0$ and $r = 0$. Rewriting (38),

$$+\mu_0\alpha + \frac{v_0}{2}\,\alpha^2 = \varepsilon\left[1 - \exp\left(-\frac{\mu_0 r'\alpha}{\varepsilon}\right)\right]. \tag{41}$$

This suggests a further substitution of

$$\alpha' = \frac{\mu_0 \alpha}{\varepsilon}, \tag{42}$$

giving us

$$\varepsilon\alpha' + \varepsilon^2 \frac{v_0}{2\mu_0^2}(\alpha')^2 = \varepsilon(1 - e^{-r'\alpha'}). \tag{43}$$

We divide by ε and then introduce a new parameter

$$\eta = \frac{\varepsilon v_0}{\mu_0^2} = \varepsilon N_0, \tag{44}$$

so that our equation becomes

$$\alpha' + \frac{\eta}{2}(\alpha')^2 = 1 - e^{-r'\alpha'}. \tag{45}$$

We are thus left with a family of solutions, indexed by the parameter η. Recall that the variables r' and α' are proportional to our actual variables, J and α. The parameter η depends on the characteristics of the base-game and represents the relative frequency of the jackpot. Recall that N_0 represents how long we must play for the base-game to get into the "long-run"; η tells us how many jackpots we can expect to have during that time. $\eta = 0$ represents the limiting case where the base game has no variance, and $\eta = \infty$ represents the case where the base-game has no advantage.

The reader may wish to note that for the infinite jackpot,

$$\alpha'_\infty = \frac{-1 + \sqrt{1 + 2\eta}}{\eta} = \frac{2}{1 + \sqrt{1 + 2\eta}}. \tag{46}$$

If we plot solutions of (45) we will notice that the equations all have a similar shape, and that the parameter has little effect of the shape. To illustrate this, I have included in Appendix 3 graphs of the two limiting cases $\eta = 0$ and $\eta = \infty$. I have actually changed the scale so that our x and y co-ordinates are of the form $y = \alpha/\alpha_\infty$ and $x = Kr$. The constant K is chosen so that all the curves have the same slope, 2, at the origin. To compute this constant, we may take the derivative of (45), after solving it for r and then equating. Out of mercy for the reader, we will omit the details of this computation. Expressing these in terms of the original variables gives us

$$y = \frac{\alpha}{\alpha_\infty}, \tag{47}$$

$$x = \frac{\varepsilon(J - J_0)}{\alpha_\infty(v_0 + \mu_0^2/\varepsilon)} = \frac{\text{Mean}}{\alpha_\infty \, \text{Var}_0}. \tag{48}$$

In the last equation, Mean is the overall expectation, but Var_0 is the variance that would occur if $J = J_0$. That is, it is the variance associated with the break-even jackpot.

Figure A6 in Appendix 3 illustrates the two limiting cases, where $\mu_0 = 0$ ($\eta = 0$) and $v_0 = 0$ ($\eta = \infty$). As can be seen, these graphs are quite similar in shape, indicating that the parameter η has a secondary effect on the solutions.

Our graphs look very much like an exponential curve. Motivated by this, I have drawn the graph of

$$y = 1 - e^{-2x} \tag{49}$$

alongside our first two curves. This is also shown in Figure A6. Remarkably, this graph lies right between the two limiting cases. This indicates that the combination of these last two equations give us an approximate solution:

$$\frac{\alpha}{\alpha_\infty} = 1 - \exp\left[-\frac{2\,\mathrm{Mean}}{\alpha_\infty\,\mathrm{Var}_0}\right]. \tag{50}$$

Of course, this would certainly be considered a closed-form equation by any standard. Please note carefully the differences between Var_0 and Mean. The former is a constant, and is the variance associated with the "break-even" jackpot. The latter is the overall mean associated with any jackpot. The author regrets that he does not currently have a good justification for this approximation, other than the fact that "it works." To employ this approximation in a specific problem, proceed as follows:

Step 1. Compute α_∞ using (39). (Note that this does not involve the size of the jackpot).
Step 2. Compute the break-even jackpot $J_0 = \mu_0/\varepsilon$.
Step 3. Compute the variance associated with the break-even jackpot, $V = v_0 + \varepsilon J_0^2 = v_0 + \mu_0^2/\varepsilon$.
Step 4. Compute what we will call the variance term, $\mathrm{VT} = \alpha_\infty V/\varepsilon$.
Step 5. Compute $x = (J - J_0)/\mathrm{VT}$.
Step 6. Compute the adjustment factor of $\mathrm{Adj}(x) = 1 - e^{-2x}$.
Step 7. Estimate $\alpha = \mathrm{Adj}(x)\alpha_\infty$

I call the function Adj the "adjustment" factor, because it allows us to take the known values for an infinite jackpot and adjust them to come up with the values for an actual, finite jackpot.

This is a nice simple formula. Note that all of these calculations can be done on a scientific calculator or on a simple spreadsheet. However, it is only an approximation. Table 3 shows how it stacks up against the exact values. The columns correspond to values of the parameter η described above.

A plus entry means our formula has underestimated the value of α and it should be increased by that percentage. Equivalently, the required bankroll may be decreased by that percentage. A negative entry has the opposite meaning. Another way to say this is that for rare jackpots (small η) our formula may produce bankroll estimates that are too small. The conservative bettor can simply multiply the computed bankroll by 10%; this will guarantee that he/she achieves the required RoR. For those desiring more accuracy, we will provide another formula for this special case.

Table 3. Exact vs. approximate calculations.

x	Adj(x)	$\eta = 0$	$\eta = 1$	$\eta = 10$	$\eta = 100$	$\eta = \infty$
0.00	0.00	0	0	0	0	0
0.10	0.02	−0.3%	−0.2%	+0.3%	+0.7%	+1.0%
0.25	0.39	−6.0%	−4.3%	+2.8%	+8.6%	+12.1%
0.50	0.63	−8.5%	−6.5%	+1.7%	+8.0%	+11.5%
0.75	0.78	−9.0%	−7.2%	+0.3%	+5.5%	+8.4%
1.00	0.86	−8.5%	−6.9%	−0.6%	+3.5%	+5.6%
1.25	0.92	−7.5%	−6.2%	−1.0%	+2.1%	+3.6%
1.50	0.95	−6.5%	−5.2%	−1.2%	+0.2%	+2.3%
2.00	0.98	−4.4%	−3.6%	−1.0%	+0.3%	+0.9%

The approximation produces the exact value when x is very small or very big. For intermediate values, there is some error. Note that it underestimates the RoR for some games (with small η) and overestimates it for others. In other words, it is attempting to be a "happy medium," which necessarily involves some compromise in accuracy.

Example 3. Let us revisit the game of Deuces Wild video poker, using the pay table and probabilities in Table 2. For this pay table, we may compute that the mean and variance of the base game are −1.005% and 11.7. We compute this by replacing the jackpot of 800 with 0 and computing the mean and variance that result. Note that the frequency of the jackpot is 2.21×10^{-5}. Using this data and the formula above, I estimated the ruin parameter for this game. The details of the computation are shown in the following table.

Table 4. Deuces Wild with cash-back.

cash-back	μ_0	η	α_∞	x	Adj(x)
0	1.005%	2.6	0.126%	0.37	0.52
0.5%	0.505%	10.2	0.156%	0.63	0.72
1.0%	0.005%	102991.7	0.194%	0.78	0.79

approx. α	exact α	error	2μ/SS	error
0.066%	0.065%	−2.1%	0.059%	8.9%
0.111%	0.116%	3.9%	0.098%	15.9%
0.153%	0.175%	12.6%	0.136%	21.9%

As you can see, I also consider various levels of cash-back. Cash-back is easy to accommodate with our formulas; we simply subtract it from the house advantage. I also computed the exact value of α using a solving utility (the solver in Excel) and I have compared the results in the following table. The approximation produced by the formula above is listed under "approx. α" I also included, in the final columns, the computation that would result if we incorrectly applied the traditional mean-variance formula to this game.

Notice that the magnitude of the errors are consistent with Table 3. Note that 1% cash-back with full-pay Deuces Wild gives a game that is almost even, even without the royal flush payoff. This drives η toward ∞. Our approximation did not perform well here, but it is a significant improvement over the traditional methods. In addition, our chart show us that for small x and large η, the error can be over 10%, and we may adjust our computed values appropriately.

7.3 Special case for "rare jackpots"

Our approximation does not perform well for very small and very large values of η. Moreover, when η is small, we are underestimating the risk of ruin, which is generally not desirable.

Small value of η will occur when the jackpot is "rare." Keep in mind that this means rare relative to N_0 of the base game. Thus, large values of μ_0, small values of v_0, and relatively small ε contribute to making η small. Yet these conditions can occur quite easily in games with progressive jackpots. Such games will typically have rare jackpots (small ε) and a large nonjackpot house advantage.

Now the case where $\eta = 0$ cannot actually occur, since this would indicate that there was zero variance in the base game. However, it does represent a useful approximation. In addition, the equations are somewhat easier here. Fortunately, it has a type of closed-form expression for this case. I asked a computer algebraic system (Maple V) to solve this equation, and it produced the following:

$$\mathrm{Adj}(x) = 1 + \frac{\mathrm{LW}(-re^{-r})}{r} \quad \text{with} \quad r = x + 1, \tag{51}$$

where LW is the Lambert W function. This obscure function is the inverse of the function $F(x) = xe^x$. Careful algebraic manipulation will show that this is in fact the solution.

For this case, $\alpha_\infty = \varepsilon/\mu_0$, and our solution simplifies to

$$\alpha = \frac{\varepsilon}{\mu_0}\left(1 + \frac{\mathrm{LW}(-re^{-r})}{r}\right) \quad \text{where} \quad r = \varepsilon J/\mu_0 \tag{52}$$

for the special case where $\eta = 0$. A table of values for this "adjustment" function is given in Table 5.

The following simpler procedure may be used for a more accurate estimate when η is small.

Table 5. Adjustment factor for special case $\eta = 0$.

x	r	Adj	x	r	Adj
0.0	1.0	0.000	1.0	2.0	0.797
0.1	1.1	0.176	1.2	2.2	0.844
0.2	1.2	0.314	1.4	2.4	0.879
0.3	1.3	0.423	1.6	2.6	0.905
0.4	1.4	0.511	1.8	2.8	0.925
0.5	1.5	0.583	2.0	3.0	0.940
0.6	1.6	0.642	3.0	4.0	0.980
0.7	1.7	0.691	3.5	4.5	0.988
0.8	1.8	0.732	4.0	5.0	0.993
0.9	1.9	0.767	∞	∞	1.000

Step 1. Compute α_∞. If v_0 is negligible, this is just ε/μ_0.
Step 2. Compute x as above, and form $r = x + 1$. If v_0 is negligible, this is just $\varepsilon J/\mu_0$.
Step 3. Look up Adj(r) in Table 5.
Step 4. Compute the estimate $\alpha_A = \text{Adj}(r)\alpha_\infty$.

Example 4. We conclude this discussion with an example from video poker: an 8/5 Jacks or Better progressive game. Progressive means that the house makes a contribution to the jackpot with each play. This requires that the initial house advantage be relatively large, and so η will be small.

The following example is motivated by Table 2 on page 20 of Wong (1993). Here we have a progressive 8/5 Jacks or Better video poker machine, with a progressive payoff for a royal flush. Here the progressive jackpot is 3200 times the wager. However, we have to be careful in simply using tabulated data such as this. Because of very slight round-off error, the probabilities will not sum to exactly 1. For some purposes, this small round-off error would not matter. However, here it is serious, because our equations associate the residual probability with an infinite jackpot. We could make an adjustment. One possibility would be to add the residual probability to that of the nonwinner, which would give us a slightly conservative estimate of our RoR. Another possibility would be to simply divide each probability by the sum of the raw data, forcing the new sum to be 1.

An even better alternative would be simply to use exact data. I am grateful to Stewart Ethier for providing me with the exact data for this game, which is found in Table 6. This data assumes that the player is using the optimal drawing strategy for this particular pay table.

Note that $\varepsilon = 0.00317\%$. We may compute μ_0 by setting the royal flush payoff to 0 and taking the mean and we obtain $\mu_0 = 5.547\%$. Similarly v_0 is 3.40. We see that $\eta = \varepsilon v/\mu_0^2 = 3.5\%$, which indicates that the jackpot is

Table 6. Jacks or Better payoff schedule.

hand	payback	probability $\times\binom{52}{5}\binom{47}{5}5$	payoff
nonwinner	0	10,988,952,720,912	−1
jacks or better	1	4,180,806,678,816	0
two pair	2	2,566,397,076,660	1
three-of-a-kind	3	1,479,679,728,204	2
straight	4	217,803,391,608	3
flush	5	221,795,807,184	4
full house	8	228,333,149,928	7
four-of-a-kind	25	46,941,780,072	24
straight flush	50	1,888,070,856	49
royal flush	3200	632,112,960	3199
sum		19,933,230,517,200	

relatively rare. Thus we may neglect the variance and use the simple forms above, with very little error.

Step 1. Compute $\alpha_\infty = 0.00317\%/5.547\% \approx 0.057$.
Step 2. Compute $r = \varepsilon J/\mu_0 = (0.00317)(3200)/(5.547) \approx 1.83$.
Step 3. From Table 5 estimate $\mathrm{Adj}(1.83) = 0.74$.
Step 4. Estimate the rp as $(0.74)(0.057) = 0.042$.

We could use this ruin parameter to estimate either risk of ruin or required bank size. For example, if these were on a \$0.25 machine, where the bet sizes were \$1.25 and we wished to keep our risk of ruin to under 5%, we would need an initial bank of $(\$1.25)(-\ln(0.05)/0.042) = \$8{,}900$.

If we incorrectly used the traditional μ/Var method of (1), we would compute the variance with the royal flush to be 328, and we would estimate our rp to be 0.028, or about 2/3 of its actual value. This means we overstate our required bank by a factor of 50% and obtained \$13,300.

Again, with solving utility I computed the exact ruin parameter by solving the ruin equation and obtained 0.0415, in good agreement with our approximate methods.

If cash-back were available with this game, our values of η would increase, so that we may wish to use the more complicated equations for α_∞ and x from before. I summarize these calculations in Table 7.

Appendices

I include several graphs to illustrate some of the concepts expressed above:

Table 7. Progressive Jacks or Better with cash-back.

cash-back	μ_0	η	α_∞	x	Adj(x)
0%	−5.547%	0.035	0.056%	0.81	0.73
1%	−4.547%	0.052	0.068%	1.20	0.84
2%	−3.547%	0.086	0.086%	1.78	0.92

approx. α	exact α	error	$2\mu/$ SS	error
0.041%	0.041%	0.9%	0.028%	32.4%
0.057%	0.057%	0.1%	0.034%	40.5%
0.079%	0.079%	0.5%	0.040%	49.3%

I. The Sorokin effect. These graphs show what happens in a game with a jackpot as the jackpot increases and they illustrate the failure of the mean-variance approximation.
II. Ruin functions. I include graphs that are typical of ruin functions. One is for a game that has a symmetric ruin function, such as a coin-tossing or normal game. The second is for a game that is skewed.
III. A graphs that demonstrates our exponential approximation for the "large jackpot" problem, discussed in the final section of the text.

All of these graphs were done with the computer algebraic system Maple V.

A1 Sorokin effect

Figures A1–A3 illustrate the Sorokin effect: the breakdown of the mean-variance approximation in the presence of a large jackpot. These graphs illustrate ruin calculations for a game with various size jackpots. The parameter r is the relative size of the jackpot; $r = 1$ gives us a break-even game. We plot values of λ, which is the RoR for a one-unit bank, the ruin parameter α, and its reciprocal. The latter is proportional to the bank required for any prescribed RoR. We plot the actual values as well as the mean-variance approximation. Notice that the latter are going in the wrong direction once r is above 2.

A2 Graphs of ruin functions

Figures A4 and A5 illustrate ruin functions. The first was done for a coin-tossing game, and shows a symmetric function. We then added a large jackpot to the game; the result is plotted in Figure A5. The dashed line shows the position of the vertex; the first function is symmetric about it.

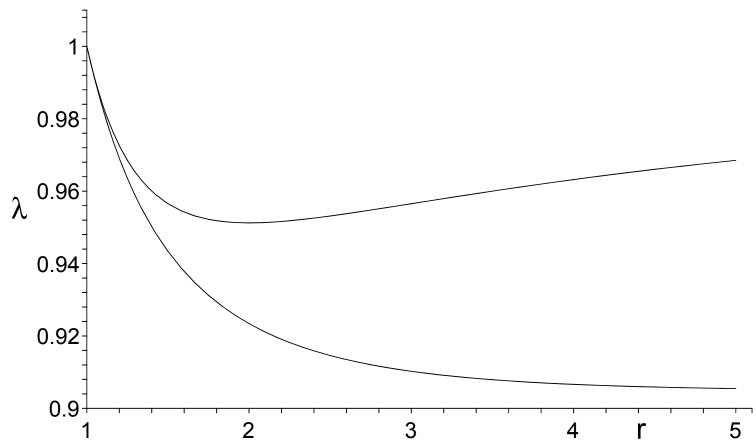

Fig. A1. Risk of ruin (λ) as a function of r. The lower curve represents the actual values; the upper curve gives the mean-variance approximation.

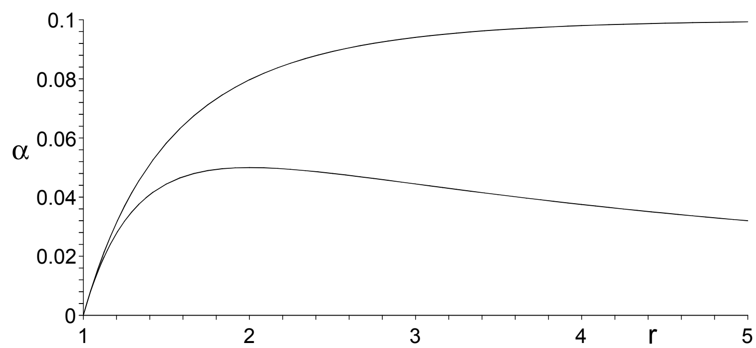

Fig. A2. Required bank as a function of r. The lower curve represents the actual values; the upper curve gives the mean-variance approximation.

A3 Infinite-jackpot approximations

Figure A6 contains graphs of the rare-large jackpot model. On the vertical axis, we plot the ruin parameter, scaled as α/α_∞. On the horizontal axis we plot the expected value for the game, which is a linear function of the relative jackpot size. The scale factor K is chosen so that the slopes of the curves are equal at the origin. The upper curve is the case where base game has zero mean; lower curve is the case where base game has zero variance.

We also impose an exponential function of the form $\alpha = A(1 - e^{x/T})$, which goes right through the middle.

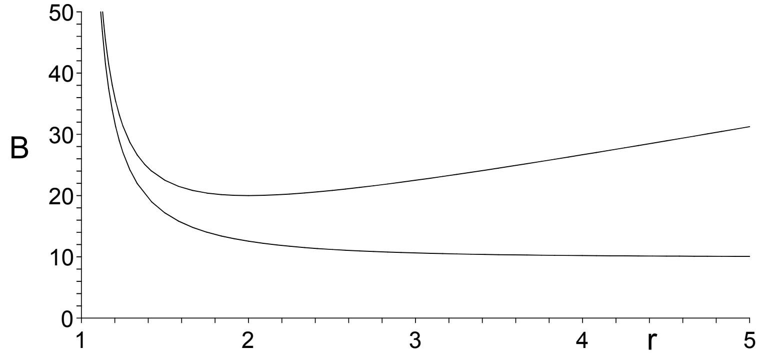

Fig. A3. Ruin parameter $(-\ln \lambda)$ as a function of r. The upper curve represents the actual values; the lower curve gives the mean-variance approximation.

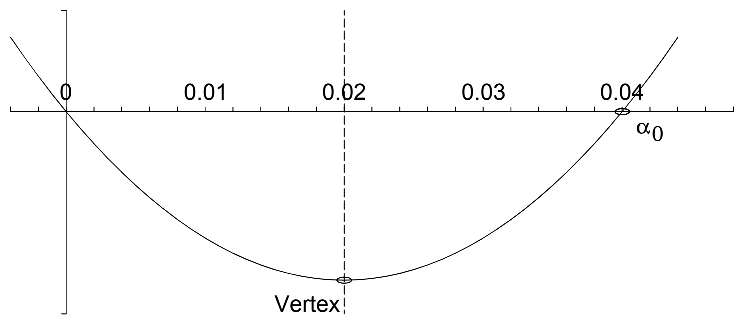

Fig. A4. Symmetric ruin function.

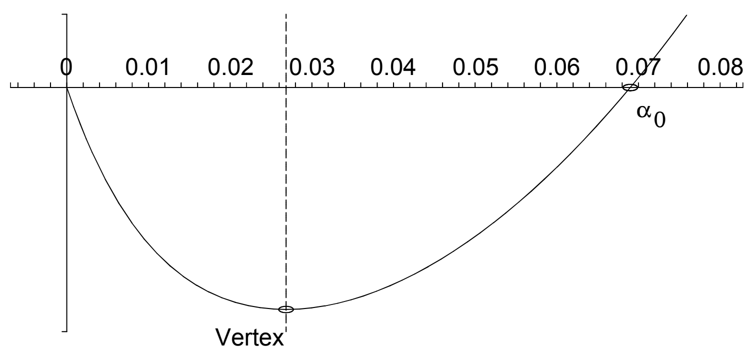

Fig. A5. Addition of large jackpot.

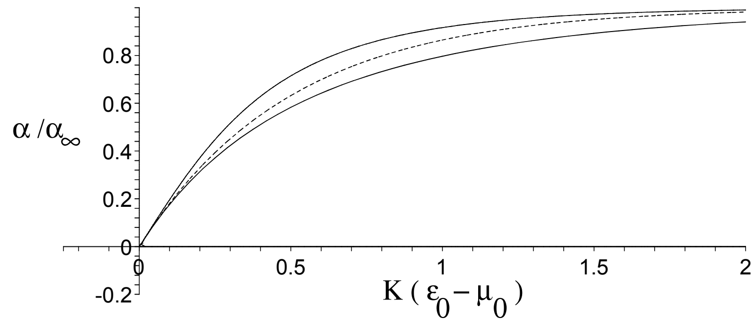

Fig. A6. Exponential approximation for the rare-large jackpot model. The upper curve is the case where base game has zero mean; lower curve is the case where base game has zero variance. The middle curve is an exponential curve.

References

1. Coolidge, J. L. (1909) The gambler's ruin. *Annals of Mathematics, Second Series* **10** 181–192.
2. De Moivre, A. (1712) De mensura sortis, seu, de probabilitate eventuum in ludis a casu fortuito pendentibus, *Philosophical Transactions of the Royal Society of London* **27**, 213–264. English translation by Bruce McClintock, *International Statistical Review/Revue Internationale de Statistique* **52** 229–262, 1984.
3. Dunbar, H. and B., J. (Math Boy) (1999) Risk of ruin for video poker and other skewed-up games. *Blackjack Forum* **19** (Fall) 21–27.
4. Epstein, Richard A. (1995) *The Theory of Gambling and Statistical Logic*, revised edition. Academic Press, New York.
5. Ethier, S. N. and Khoshnevisan, Davar (2002) Bounds on gambler's ruin probabilities in terms of moments. *Methodology and Computing in Applied Probability* **4** 55–68.
6. Feller, William (1968) *An Introduction to Probability Theory and Its Applications*, 3rd edition. John Wiley & Sons, New York.
7. Fisz, Marek (1963) *Probability Theory and Mathematical Statistics*, 3rd edition. John Wiley & Sons, New York.
8. Griffin, Peter A. (1981) *The Theory of Blackjack: The Compleat Card Counter's Guide to the Casino Game of 21*, second edition. GBC Press, Las Vegas.
9. Griffin, Peter A. (1999) *The Theory of Blackjack: The Compleat Card Counter's Guide to the Casino Game of 21*, 6th edition. Huntington Press, Las Vegas.
10. Harris, Brett R. (1997) The theory of optimal betting spreads. `http://www.bjmath.com/bjmath/Betsize/theory.htm`, posted Nov. 13.
11. Jensen, Stephen (2001) Optimal drawing strategy for Deuces Wild video poker. Unpublished manuscript, University of Utah.

12. Sileo, Patrick (1992) The evaluation of blackjack games using a combined expectation and risk measure. In: Eadington, William R. and Cornelius, Judy A. (eds.) *Gambling and Commercial Gaming: Essays in Business, Economics, Philosophy and Science*. Institute for the Study of Gambling and Commercial Gaming, University of Nevada, Reno, 551–563.

13. Schlesinger, Don (2005) *Blackjack Attack: Playing the Pros' Way*, 3rd edition. RGE Publishing, Las Vegas.

14. Thorp, Edward O. (2000) The Kelly criterion in blackjack, sports betting, and the stock market. In: Vancura, Olaf, Cornelius, Judy A., and Eadington, William R. (eds.) *Finding the Edge: Mathematical Analysis of Casino Games*. Institute for the Study of Gambling and Commercial Gaming, University of Nevada, Reno, 163–213.

15. Wong, Stanford (1993) *Professional Video Poker*. Pi Yee Press, La Jolla, CA.