

ESTIMATING THE PROBABILITY THAT THE GAME OF MONOPOLY NEVER ENDS

Eric J. Friedman
Shane G. Henderson
Thomas Byuen
Germán Gutiérrez Gallardo

School of Operations Research and Information Engineering
Cornell University
Ithaca NY 14853, USA

ABSTRACT

We estimate the probability that the game of Monopoly between two players playing very simple strategies never ends. Four different estimators, based respectively on straightforward simulation, a Brownian motion approximation, asymptotics for Markov chains, and importance sampling all yield an estimate of approximately twelve percent.

1 INTRODUCTION

Have you ever played a game of Monopoly and, after a few hours, started to wonder if the game would ever end? We found that non-ending games were a distinct possibility while conducting simulations to try to determine effective strategies for the game. So we naturally became curious: What is the probability that the game of Monopoly never ends? By “never ending” we mean over an infinite horizon, and not just that the game lasts longer than reasonable players might be willing to play. It is not easy to estimate this probability using standard simulation, since one can never be certain that a game that is still going will not eventually end.

The mechanism by which games can go on forever has been understood at least since [Lehman and Walker \(1975\)](#). Essentially, the game reaches a stage where players continue to accumulate cash indefinitely. Although there is a chance, at any point in time, that any player could fall on a run of bad luck and lose their cash, this probability is small enough that the player’s wealths simply grow to infinity. There is no shortage of cash in the game, since the rules of Monopoly ([Fritzlein 2009](#)) state that if the bank ever runs out of money, one can simply “print more” on extra sheets of paper.

The probability α that the game goes on forever depends on the strategies that the players adopt. We detail our assumptions about the strategies the players adopt in Section 2. For now, suffice it to say that we use very simple strategies that involve buying properties whenever possible and building houses or hotels whenever possible, while still maintaining a modest threshold of cash on hand to deal with contingencies.

In this paper we give four different simulation-based estimators of α , assuming two-player games. The estimators are built from a combination of an extremely detailed simulation model coded in Java, and an approximating Markov chain model.

Surprisingly (we think) there have been few serious academic studies of the game of Monopoly. Our survey revealed a few relevant books, including [Lehman and Walker \(1975\)](#) in which two (then) Cornell undergraduates studied strategy using what were, at the time, state of the art simulation techniques – playing many monopoly games. They also performed some analysis based on Markov chains. Somewhat more recently, [Darzinskis \(1987\)](#) simulated many monopoly games to derive good strategies.

More recent work was spurred by Stewart’s analysis in Scientific American ([Stewart 1996a](#), [Stewart 1996b](#)), which describes the use of Markov chains in the analysis. This led to several analyses on websites (although such analyses have no doubt occurred since the game was introduced). For example, one computes accurate site-visiting probabilities for Markov chains ([Collins 2009](#)) and another applies these to develop a finance-inspired analysis of good strategies ([Darling 2009](#)). Lastly, for reference, the rules and other facts about monopoly can be found on [Fritzlein \(2009\)](#).

While many have cursed the length of monopoly games, we have not found any detailed analyses of the game's length, nor any studies that compute the probability that the game goes on forever. But surely, millions of players have thought that the game might never end.

Section 2 describes the simulation model and the player strategies we adopt. Then, in Section 3 we describe a simple estimator of an upper bound on α . Section 4 explains in more detail how games can go forever, describing "wealth plots" as a conceptual basis for understanding how this arises. It also explains how we can view games as consisting of potentially two stages. One can then estimate the probability α that the game goes on forever by conditioning on first reaching Stage 2 and the state of the board upon reaching Stage 2. Section 5 gives an approximation for the conditional probability based on a Brownian approximation of Stage-2 dynamics of player wealths. The conditional probability is instead estimated using asymptotics for Markov Chains using a standard Markov chain model of Stage 2 dynamics in Section 6. In Section 7 we implement an importance sampling estimator using the Markov chain asymptotics. Section 8 rounds out the paper with discussion of the results and an outline of future research efforts.

2 THE SIMULATION MODEL AND PLAYER STRATEGIES

We have developed a very detailed simulation model of the game of Monopoly in Java. The simulation model is object based, and can accommodate more than two players, although we restrict attention to the two-player game in this paper. To verify the model we have performed a large number of detailed game traces, and compared the simulation model's predictions with available results on, e.g., frequencies of landing on particular squares, average net cash influxes from the bank per turn, and so forth. We have also compared the model's results with those of a Markov chain model that is detailed in Section 6 below.

The Markov chain model needs to have a small state space to allow computations to be performed quickly. Accordingly, throughout the remainder of this paper, both the simulation and the Markov chain operate under the following assumption. In the real game of Monopoly, a player's turn can involve rolling the dice up to three times, depending on the number of doubles rolled. In our version of the game, each player only rolls the dice once, and doubles are treated exactly the same as any non-double roll. Players must pay rent and so forth when rolling doubles, and this is in accordance with the official rules of the game. Each step of the Markov chain corresponds to one *round*, i.e., one roll of the dice for each player. We will measure time in rounds rather than turns in the remainder of this paper.

Under this assumption, the "3 doubles in a row leads to jail" rule is ignored. The impact of this assumption is that players go to jail less often than in the real game. However, the impact seems slight, since players roll 3 doubles in a row only once every 216 turns, or once every 258 rolls.

Our players play a common, simple strategy. That strategy can be approximately characterized as follows.

1. Players maintain a reserve threshold of cash that is easily varied but we set to the maximum of \$200 and the maximum rent on a property owned by another player.
2. Players buy every property they land on that is for sale, unless doing so would reduce their cash below the reserve threshold.
3. Players never bid on properties that are up for auction.
4. Players build on streets as long as they have the cash on hand to do so, building on the more expensive properties when they have a choice. Building is always done evenly, so that for example, if 3 properties form a Monopoly, then either 3 houses are purchased or none are purchased.
5. When a double is thrown while in jail, the player comes out and advances their token by the double that was rolled, and this completes their turn. (They do not roll again.)
6. A player does not pay to get out of jail, unless he/she has rolled non-doubles on three successive attempts. After paying to get out of jail, the last of the three rolls is taken as the player's roll.
7. When a player receives a "Get out of Jail Free" card, the card is immediately sold back to the bank for \$50.
8. Players do not trade properties.
9. When a player receives an "Advance token to your choice of railroad" card, the player moves to the railroad they are currently on or, if not currently on a railroad, moves to the one immediately preceding them. This will result in them receiving \$200 for passing Go unless they are on the first Community Chest square after Go.

3 A SIMPLE ESTIMATOR

Our first estimator of the probability α that a game goes on forever is very natural and simple. Simulate a large number of games, keeping track of the fraction of games that are still going after n rounds have been completed, for various values of




Figure 1: Estimated fraction of games still going after n rounds have been completed. The solid line gives the estimate and the dashed lines represent (pointwise) 95% confidence intervals.

n . Then look for a horizontal asymptote on the graph. We took exactly this approach with our simulation model, simulating 3100 games. Abbreviated results are given in Figure 1.

The curve in Figure 1 appears to asymptote to a value somewhere near 10%. A 95% confidence interval for the fraction of games still going after $n = 10000$ turns is 0.12 ± 0.01 . This represents an upper bound on α . Unfortunately, we have no way of knowing whether the curve would drop to an even lower level were we to simulate for longer. To obtain a better sense of the probability that games go on forever we need a different approach. To derive such an approach, it is helpful to try to understand exactly how games can go on forever.

4 HOW GAMES GO ON FOREVER

To gain a conceptual understanding of how the game can go on forever, let $W_n = (W_n(1), W_n(2))$ be the wealth of the two players after n rounds, i.e., after each player has rolled the dice n times. Here, we measure wealth as the cash on hand plus the value of all other properties, houses, and hotels if they were to be sold back to the bank. (These are worth one half of the face value in this case.) If either player's wealth ever becomes negative then they are bankrupt and the game finishes. If we plot the sequence of points W_0, W_1, W_2, \dots (see Figure 2) then we see that the game goes on forever as long as the plotted points stay within the nonnegative quadrant.

At first, the game configuration changes rapidly, with properties being bought and houses built. Call this period "Stage 1." Some games are completed within Stage 1, but others continue beyond this transient stage, reaching a point where the board configuration has essentially solidified, with no further property exchanges or house building unless one player or other approaches bankruptcy. Let us call this second period "Stage 2." For our simulations, we defined Stage 2 as being the



Figure 2: Conceptual plot of the wealth process.

first time that all properties are purchased and both players have at least \$5000 in cash. For games that reach Stage 2, one can compute the “drift” of the player wealths $\beta = (\beta(1), \beta(2))$, where the term “drift” means the average change in player wealth over a large number of turns. If either drift is negative, then one is essentially certain that one of the players will go bankrupt eventually and the game will end. But if both drifts are positive, then there is a chance that the wealth process will escape to infinity before hitting either axis, and the game goes on forever. Figure 3 gives some examples of such plots for simulated games. Stage 1 is indicated by a cluster of points near the origin and then, after reaching Stage 2, the wealth of the two players accumulates and we see rays heading up and to the right.

Consider how to estimate $1 - \alpha$, the probability that the game ends in finite time. This probability is simply the probability that the game ends in Stage 1, plus the probability that the game ends in finite time after reaching Stage 2. The first probability is easily estimated, since there is a clear time at which to stop the simulations – the minimum of the times when one player goes bankrupt and the time at which Stage 2 is reached. It is the second probability that is challenging to estimate.

Suppose we have reached Stage 2, so that the board configuration is fixed, and let us think of this point as time zero. Let $A(i)$ be the event that Player i goes bankrupt, $i = 1, 2$, i.e., the wealth of Player i hits 0. Here we assume that each player continues to play over an infinite horizon, even if the opposing player goes bankrupt. Conceptually, play continues after such a point, with payments going to and from a *phantom* player.

We are interested in $P(A(1) \cup A(2))$, the probability that one of the players goes bankrupt and the game ends. This probability seems complicated to estimate, since it involves a hitting problem in two dimensions. Boole’s inequality bounds this probability by $P(A(1)) + P(A(2))$, and so we now have two hitting problems each in one dimension, which seems much more tractable. We expect that these bankruptcy probabilities are very small for the cases of interest, and that one will tend to dominate the other, so the overestimation implied through the use of Boole’s inequality should be slight. We have suppressed the dependence of these probabilities on the initial wealth $w = (w(1), w(2))$ of the two players, and the initial layout of the board (player locations, property ownership, house and hotel development etc.). We will later want to vary these parameters, so $P(A(1)) \neq P(A(2))$ in general.




Figure 3: Plots of the wealth process for two simulated games. In both cases we see the linear trend of the curve that corresponds to the drift β .

5 BROWNIAN MOTION

The next question we must address is how to estimate $P(A(1))$, the probability that Player 1 goes bankrupt, where we again suppress the initial conditions corresponding to the board configuration and initial player wealths upon reaching Stage 2. (The calculations for Player 2 are identical.) Let \tilde{w} be Player 1's initial wealth, and \tilde{W}_n be Player 1's wealth after n rounds, with $\tilde{W}_0 = \tilde{w}$. The player goes bankrupt if there exists an $n \geq 0$ such that $\tilde{W}_n \leq 0$.

Looking at Figure 3, it is natural to approximate $(\tilde{W}_n : n \geq 0)$ by a Brownian motion starting at \tilde{w} in one dimension. Such a process is characterized by a drift γ , and a variance parameter η^2 . If $\gamma < 0$, then the Brownian motion drifts to $-\infty$, so is certain to hit 0, and so $P(A(1)) = 1$. Similarly, if the drift $\gamma = 0$ then the Brownian motion hits any point on the real line in finite time, and so $P(A(1)) = 1$. When $\gamma > 0$, the Brownian motion drifts to $+\infty$ so is not certain to hit 0, but instead hits 0 with probability $\exp(-2\gamma/\eta^2)$. (This is a standard calculation that can be found, e.g., in [Ross 1996](#), Section 8.4.) We thus estimate γ and η^2 from the path of $(\tilde{W}_n : n \geq 0)$ and then estimate $P(A(1))$ by $\min\{\exp(-2\gamma/\eta^2), 1\}$.

The drift γ is easily estimated as $(\tilde{W}_n - \tilde{W}_0)/n$ for some large n . To estimate η we first remove the drift to get $Y_i = \tilde{W}_i - \tilde{W}_0 - i\gamma$ for all i , and then use non-overlapping batch means with batches of size b : We first compute the k batch means M_j , $j = 1, \dots, k$, where M_j is the average of the observations $Y_{(j-1)b+1}, \dots, Y_{jb}$, $j = 1, \dots, k$ (assuming $kb = n$). We can then estimate η^2 by b multiplied by the sample variance of M_1, M_2, \dots, M_k . In our code we took $b = 60$ rounds and $k = 170$ for a total runlength of $n = 10200$ after reaching Stage 2.

This estimator of second-stage bankruptcy probabilities is compared with two others in Section 8, once we define those two other estimators.

The above steps are completed for each game that reaches Stage 2 for each player. Let us now reuse some counting indices to describe the overall estimator that Monopoly goes on forever. Define $U_k = 1$ if the k th game reaches Stage 2, and $U_k = 0$ otherwise. Also, let V_k be zero if the k th game does not reach Stage 2, and otherwise let it be the minimum of 1 and the sum of the Brownian estimates of the game ending as outlined above. Our overall estimator that the game goes on forever is then

$$1 - \frac{1}{n} \sum_{k=1}^n [1 - U_k + U_k V_k] = \frac{1}{n} \sum_{k=1}^n U_k (1 - V_k).$$

This estimator is a sample average of i.i.d. observations, so can be analyzed using standard confidence-interval methodology. A 95% confidence interval for the probability that games will go on forever, based on $n = 3100$ replications is 0.12 ± 0.01 . This is the same (to two decimal places) as the confidence interval obtained using the “brute-force” strategy of simulating for a very long time horizon, thereby reinforcing that result.

This estimator is appealing in the sense that it attempts to model the infinite-time behaviour of the wealth process. However, it suffers from the fact that each replication that reaches Stage 2 requires further estimation of the drift and volatility parameters of the approximating Brownian motion, and this is computationally expensive. Furthermore, the estimated

parameters are subject to the usual noise in simulation estimates and are combined in a nonlinear function, and one might expect that this could lead to errors in the overall estimate of the probability that the game goes on forever. We now turn to another approach to estimating this probability that addresses, to some extent, each of these difficulties.

6 MARKOV CHAINS

It has been noted many times that the game of Monopoly can be modeled, at various levels of detail, by a Markov chain. See, for example, [Stewart \(1996a\)](#), [Stewart \(1996b\)](#) where the dice rolls are modeled as a Markov chain, as is the board motion of a single player, resulting in the steady-state probabilities of landing on various squares.

In this section we present a simple Markov model of a 2-player Monopoly game. Our purpose is to apply Markov chain theory to obtain two further estimates of the probability that the game goes on forever, presented in this and the following section. We rely heavily on the theory presented in [Lehtonen and Nyrhinen \(1992\)](#) for ruin probabilities in Markov processes.

We use a Markov chain $X = (X_n : n \geq 0)$, where each state transition represents a roll of the dice by *both* players. Here $X_n = (X_n(1), X_n(2))$, where $X_n(i)$ is the current location of Player i on the board. As is standard, $X_n(i)$ can take one of 3 values for “in jail,” which distinguish how many turns the player has been present in jail. So there are 4 states for jail: “just visiting”, and “in jail for j turns so far”, $j = 0, 1, 2$. Let $d = 42^2 = 1764$ be the number of states in the chain. (There is no state for the “Go to Jail” square.)

Let $B(x, i)$ denote a generic random payment paid to Player i from the bank during one transition of the Markov chain starting from the state x . This quantity could be negative if the player makes a payment to the bank, e.g., when assessed for street repairs. Similarly, let $R(x, i)$ denote any rent or other payment received by Player i from the other player in one transition of the Markov chain starting from the state x . This quantity could also be negative.

The amount of money received by Player i ($i = 1, 2$) over m steps of the Markov chain is then

$$\tilde{M}_n(i) = \sum_{n=0}^{m-1} \tilde{C}_n(i),$$

where

$$\tilde{C}_n(i) = B_n(X_n, i) + R_n(X_n, i).$$

If Player i 's initial wealth is $w(i)$, then the probability that Player i eventually goes bankrupt is

$$P(A(i)) = P(\exists n \geq 0 : w(i) + \tilde{M}_n(i) \leq 0),$$

i.e., Player i goes bankrupt if his net worth hits 0.

Suppose we fix the layout of the board in the sense that all properties are owned and no further building will occur. The conditional distribution of $(\tilde{C}_n(1), \tilde{C}_n(2))$, conditional on $(X_n : n \geq 0)$, then depends only on X_n and X_{n+1} , and so $((X_n, \tilde{M}_n) : n \geq 0)$ is a Markov-additive process. We can therefore apply the theory of ruin probabilities of Markov-additive processes in our context. To apply that theory it is convenient to work with $C_n(i) = -\tilde{C}_n(i)$ and $M_n(i) = -\tilde{M}_n(i)$, and define the probability of bankruptcy as

$$P(A(i)) = P(\exists n \geq 0 : M_n(i) \geq w(i)),$$

i.e., the chance that net outgoings will equal or exceed initial wealth.

This bankruptcy probability depends on the long-run “drift” of $M_n(i)$ for each i . To compute this drift, first let $c(x, i)$ be the conditional mean of $C_n(i)$ conditional on $X_n = x$. Next, note that the chain X is irreducible and aperiodic on a finite state space, and is therefore positive recurrent. Let π denote its stationary distribution. The strong law of large numbers for Markov chains, e.g., [Meyn and Tweedie \(1993\)](#), Theorem 17.0.1, then ensures that

$$\frac{M_n(i)}{n} \rightarrow \sum_x \pi(x) c(x, i) = \mu(i)$$

as $n \rightarrow \infty$ with probability 1. We refer to $\mu(i)$ as the *drift* of Player i 's outgoings. It corresponds to the negative of the previously computed drift $\beta(i)$, which is the drift of the wealth process.




Figure 4: Scatter plots of the estimated drift for a single player. The solid line is the $y = x$ line. The expected drift on the horizontal axis is computed using steady-state calculations from the Markov chain, whereas the vertical axis values are simulation results. The left (right) plot gives results before (after) debugging, where a bug in the simulation and a difference in modeling assumptions between the simulation and the Markov chain model generated a couple of outlier points and a skewed slope of the points, respectively.

We have used these calculations to help verify our simulation code. In particular, we can compute drifts of games that reach Stage 2 using both the Markov model as above, and the simulation code, and compare those drifts. See Figure 4.

If $\mu(i) > 0$ then Player i will eventually run out of money with probability 1, i.e., $P(A(i)) = 1$, so that the game will certainly end. (Note that this probability is, as mentioned earlier, conditional on the initial conditions of the chain, which we specify below.) If $\mu(i) = 0$ then again $P(A(i)) = 1$, i.e., when the drift is 0, Player i is again certain to go bankrupt. This follows from the functional central limit theorem [Meyn and Tweedie \(1993\)](#), Theorem 17.4.4.

If $\mu(i) < 0$ then the situation is more complicated. Here $M(i) = (M_n(i) : n \geq 0)$ has negative drift, so there is a probability, depending on the initial wealth of the players and the initial state of the chain, that Player i will go bankrupt. We need to compute this probability. We apply a result due to [Lehtonen and Nyrhinen \(1992\)](#) for Markov-additive processes on discrete Markov chains. Lehtonen and Nyrhinen specialized existing large-deviations theory for Markov-additive processes to develop asymptotics and a simulation estimator for the ruin probability for Markov chains. The simulation estimator has provably good performance (asymptotic optimality in the language of rare-event simulation) as the probability of ruin becomes smaller. We now specialize their development to the Monopoly context. In order to do so, we need to compute the root of a certain equation involving Laplace transforms.

Define

$$\hat{p}_{xy}(\theta, i) = E[e^{\theta C_n(i)} I(X_{n+1} = y) | X_n = x]$$

to be the Laplace transform (LT) of the one-step outgoings of Player i on the event that the chain transitions from x to y . This quantity does not depend on n , and here x and y each denote a pair of locations on the Monopoly board. This quantity is 0 if transitions from x to y have zero probability. The transform is defined for all $\theta \in (-\infty, \infty)$, because C_n is bounded.

Let $\hat{P}(\theta, i)$ be the $d \times d$ matrix of LTs defined above, which is a non-negative matrix. The theory of such matrices implies that it has an eigenvalue $\lambda(\theta, i)$ such that

1. $\lambda(\theta, i)$ is simple, real and strictly positive, and $\lambda(\theta, i) > |\lambda|$ for any other eigenvalue λ , and
2. $\lambda(\theta, i)$ has a strictly positive right eigenvector $h(\theta, i)$.

Define $r(\theta, i) = \log \lambda(\theta, i)$ (here and elsewhere we use the natural logarithm unless otherwise specified). It is known that $r(\cdot, i)$ is strictly convex in θ , that $r(0, i) = 0$, and that $r'(0, i) = \mu_i < 0$. So $r(\cdot, i)$ initially dips below 0, and can have at most one positive root $\theta^*(i)$ (so that $r(\theta^*(i), i) = 0$ and $r'(\theta^*(i), i) > 0$). Then it is known that

$$\lim_{w(i) \rightarrow \infty} \frac{1}{w(i)} \log P(A(i)) = -\theta^*(i). \tag{1}$$

The logarithmic asymptotic result (1) suggests the approximation

$$P(A(i)) \approx e^{-\theta^*(i)w(i)}, \quad (2)$$

although there are many other possibilities for the right-hand side of (2) that also satisfy (1).

So how can we compute this approximation? We first have to identify $\theta^*(i)$. This involves forming the matrix $\hat{P}(\theta, i)$ for various θ , computing its maximum eigenvalue, and then seeking the positive value $\theta^*(i)$ for which that eigenvalue = 1, i.e., $r(\theta^*(i), i) = 0$. This is a root-finding problem that can be efficiently solved using a binary search. We implemented this approach in Java, computing the eigenvalues using the power method with 200 iterations, which came to within 0.002% of the eigenvalues as computed using MATLAB.

The remaining question is how to compute $\hat{P}(\theta, i)$ for a given θ . Recall that $\hat{P}(\theta, i)$ is a $d \times d$ matrix, where $d = 1764$. For a given value of θ we need to fill in the values of this matrix. This is algorithmically the most difficult part of the calculation. Our approach hinges on the fact that we can compute the LT of the total outgoings of Player 1 due to Players 1 and 2 each taking a single turn by computing the LT of the total outgoings of Player 1 due to Player 1's turn, then the LT of the total outgoings of Player 1 due to Player 2's turn separately, and then multiplying these LTs. (Recall that the LT of a sum of independent random variables is the product of the individual LTs.) We do not provide further details here.

So our estimator based on the asymptotic approximation (2) is as follows. Recall that U_k is the indicator that the game reaches Stage 2. Define \tilde{V}_k to be the minimum of 1 and the sum of the probabilities (2) for the two players. The estimator is then

$$1 - \frac{1}{n} \sum_{k=1}^n [1 - U_k + U_k \tilde{V}_k] = \frac{1}{n} \sum_{k=1}^n U_k (1 - \tilde{V}_k).$$

This estimator is a sample average of i.i.d. observations, so can be analyzed using standard confidence-interval methodology. A 95% confidence interval for the probability that games will go on forever based on $n = 3100$ replications is 0.12 ± 0.01 . This agrees with the previous estimators.

This approach is appealing in that it captures the infinite-time behaviour of the wealth process, and in contrast to the Brownian-motion estimator is relatively efficiently computed, and conditional on the state of the board at the start of Stage 2 contains no simulation noise. It suffers from an important disadvantage, however. The approximation (2) is one of many possible approximations that is consistent with the logarithmic asymptotic (1) and another consistent approximation might yield very different results. We now address that concern using importance sampling.

7 IMPORTANCE SAMPLING

In this section we present the key ideas behind an importance sampling scheme for estimating the probability that the game goes on forever, again relying heavily on the theory in [Lehtonen and Nyrhinen \(1992\)](#). If one is familiar with that theory then our approach is conceptually exactly as they suggest. The difficulty really lies in the algorithmic implementation, because the discrete distributions involved have many points of support.

The overall simulation procedure is as follows. We start by simulating a game until either the game ends or Stage 2 is reached, whichever comes first, exactly as before. This is done under the usual probability distribution associated with the simulation. Upon reaching Stage 2 we compute the drift of the player wealths as described in the previous section, and if the drift of either player's wealth is not strictly positive then the conditional probability that the game ends is estimated to be 1. Otherwise, we need to compute the conditional probability that the game will end, conditional on the initial configuration at the onset of Stage 2. At this point we change the transition probabilities of the Markov chain (and the corresponding payments) so as to ensure that both players will go bankrupt eventually with probability 1. This biases the probability estimates, so to correct the estimates we have to multiply by the likelihood ratio.

Let $p_{xy}(c, i)$ be the probability, under the usual probabilities that govern Monopoly, that from state x we transition to state y , and in the process player i incurs outgoings c . The importance sampling scheme changes these probabilities upon reaching Stage 2 to $q_{xy}(c, i)$ say. If we let $\tau(i)$ be the time at which player i goes bankrupt, then conditional on reaching Stage 2 and the configuration of the board and so forth at that time, and taking the time of reaching Stage 2 to be time zero, then the conditional probability that Player i goes bankrupt is

$$EI(\tau(i) < \infty) = \tilde{E} \left[I(\tau(i) < \infty) \prod_{k=1}^{\tau(i)} \frac{p_{X_{k-1}, X_k}(C_k(i), i)}{q_{X_{k-1}, X_k}(C_k(i), i)} \right], \quad (3)$$

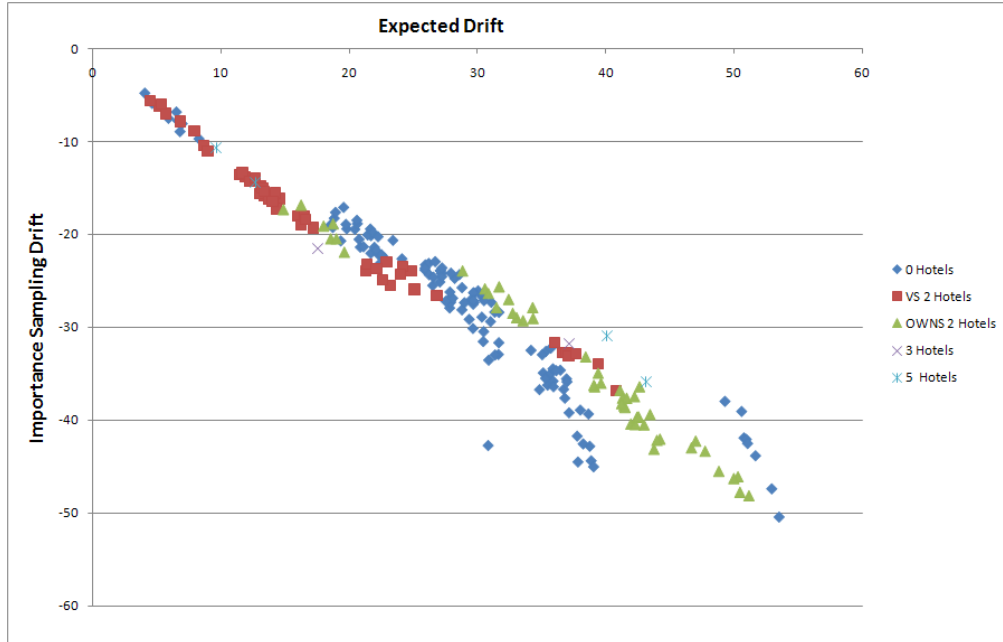


Figure 5: The expected drift of the wealth process under the standard and importance-sampling distributions. The effect of importance sampling is approximately to reverse the sign of the drift. The points are broken down into games where there are no hotels, where one player plays against a player with 2 hotels, where one player owns 2 hotels and is playing someone with no hotels, and when there are more than 2 hotels on the board.

where \tilde{E} denotes the fact that we simulate the chain under the *new* dynamics q rather than p in Stage 2.

The theory in Lehtonen and Nyrhinen suggests that in order to estimate $P(A(i))$, we should simulate the chain under the particular choice of transition probabilities

$$q_{xy}(c, i) = \frac{\hat{p}_{xy}(\theta^*)h(y) e^{\theta^*c} p_{xy}(c)}{\lambda(\theta^*)h(x) \hat{p}_{xy}(\theta^*)},$$

where we write θ^* for $\theta^*(i)$ and so forth for brevity. We have used this form of the expression to show that if one first adds over all possible values of c , the second fraction sums to 1, and then adding over all possible y the first fraction sums to 1, so that this is, indeed, a set of transition probabilities. However, these transition probabilities can also be written

$$q_{xy}(c, i) = \frac{e^{\theta^*c} p_{xy}(c) h(y)}{h(x)}, \quad (4)$$

since $\lambda(\theta^*) = 1$. This second expression better reflects how we generate transitions using the new probabilities.

With this particular choice of importance-sampling transition kernel, $I(\tau(i) < \infty) = 1$ a.s., and so (3) simplifies to

$$\tilde{E} \left[e^{-\theta^* \sum_{k=1}^{\tau(i)} C_k(i)} \frac{h(X_T)}{h(X_{\tau(i)})} \right], \quad (5)$$

and this is our estimator of the conditional probability that Player i goes bankrupt. As with many other estimators of ruin probabilities in the light-tailed setting, we have seen that the effect of importance sampling is approximately to reverse the sign of the drift. Figure 5 shows this effect, where we plot the expected drift as computed by the Markov chain model under the original probability dynamics versus the expected drift under importance sampling for a number of games.




Figure 6: Scatter plots of the estimated conditional probability of games that have reached Stage 2 ending in finite time. The three plots give, from left to right, scatter plots of the Brownian estimator versus the estimator based on the asymptotic approximation (2), the Brownian estimator versus the importance-sampling estimator, and the asymptotic-approximation estimator versus the importance-sampling estimator.

The full estimator of the probability that the game goes on forever is then as follows. As before, let U_k be the indicator that Stage 2 is reached on the k th simulation replication. Upon reaching Stage 2, we switch from the full Monopoly simulation under the usual probabilities to simulations of the Markov chain under the importance-sampling transition probabilities (one for each player), and simulate each path until the players go bankrupt. We can then compute the quantity inside the expectation in (5) as a single-replication estimator of the conditional expectation of bankruptcy for each player $i = 1, 2$. Next, define \hat{V}_k to be the minimum of 1 and the sum of these two estimators of the conditional expectations in Stage 2. The importance sampling estimator of the probability that the game goes on forever is then

$$1 - \frac{1}{n} \sum_{k=1}^n [1 - U_k + U_k \hat{V}_k] = \frac{1}{n} \sum_{k=1}^n U_k (1 - \hat{V}_k).$$

This estimator is a sample average of i.i.d. observations, so can be analyzed using standard confidence-interval methodology. A 95% confidence interval for the probability that games will go on forever based on $n = 3100$ replications is 0.12 ± 0.01 .

8 COMPARISON, DISCUSSION AND FUTURE RESEARCH

All four of our estimators yield confidence intervals that suggest that the probability that the game goes on forever is close to 12%. This was the conclusion we obtained using the simplest estimator of all, in Section 3. However, we could not be certain, based on that estimator alone, that the probability would not actually be significantly lower than this value. The other three estimators were motivated by this issue and the fact that we had fun computing them. Using those estimators we confirmed that, in fact, the potential for the game ending after Stage 2 is reached is very slight. The three estimators based on Stage 2 calculations are all biased due to the use of Boole's inequality to bound the probability that either player goes bankrupt by the sum of those probabilities. Beyond that bias, the Brownian motion estimator and the asymptotic-approximation estimator both possess additional bias, but the importance-sampling estimator does not.

It is interesting to consider how these three estimators compare. We conducted a small experiment where we identified approximately 100 games that reach Stage 2, and that have conditional probabilities that the game will eventually end that are in the range of 0.01 through 0.99. We then computed each of the estimators, and produced scatter plots of their values in Figure 6.

There is close agreement between the asymptotic-approximation estimator and the importance-sampling estimator, but the Brownian estimator differs. This difference does not exhibit itself in our final estimators of the probability that the game goes on forever since the probability of reaching a point in Stage 2 where these differences occur is so small that it cannot significantly influence the results.

In future work we might see how strategies can influence the length of the game. Our assumption that players do not trade severely limits the possibility of developing properties, which in turn limits the probability that the game will end. When players have hotels, or many houses, the variance in the player wealth process is much higher than when there is only

limited development, and so the chance that the game will end would likely be higher. In that setting we would expect that the three estimators that are based on Stage 2 calculations would provide a benefit beyond the simple estimator, and that the probability that the game goes on forever would be smaller than twelve percent.

But primarily we hope to use the tools we have developed to this point to identify highly effective player strategies. We cannot hope to approach the complexity of the strategies outlined in [Lehman and Walker \(1975\)](#), since those strategies involve extremely complicated and subtle trades that go well beyond what we envisage coding. However, we can certainly investigate questions such as whether one should build through to hotels, or to halt development earlier so as to “soak up” houses. An important question is how to use some kind of automated learning process to identify the best of a family of strategies. The reason this is important is not just that we are enthusiasts and see it that way, but also because Monopoly is a microcosm of many important systems in real life, where sequential decisions under uncertainty are made in the face of competition from other entities.

True enthusiasts of estimating the probability that the game goes on forever might also implement importance-sampling estimators that do not rely on the use of Boole’s inequality, so as to obtain completely unbiased estimators of the probability that the game goes on forever. However, our results suggest that such an effort would have to identify the probability extremely accurately in order to provide additional information beyond the confidence intervals produced in this paper.

ACKNOWLEDGMENTS

This work was partially supported by National Science Foundation Grant Numbers CMMI-0758441, ITR-0325453 and CDI-0835706. We would like to thank current and former students Alex Aidun, Raghu Chandrasekaran, Sean Choi, Dennis Li, and Andy Tuchman for simulation programming.

REFERENCES

- Collins, T. 2009. Probabilities in the game of Monopoly. <<http://www.tkcs-collins.com/truman/monopoly/monopoly.shtml>> [accessed Aug 4, 2009].
- Darling, T. 2009. How to win at Monopoly – a surefire strategy. <<http://www.amnesta.net/other/monopoly/>> [accessed Aug 4, 2009].
- Darzinskis, K. 1987. *Winning Monopoly: A complete guide to property accumulation, cash flow strategy, and negotiating techniques when playing the best-selling board game*. HarperCollins.
- Fritzlein 2009. Monopoly/official rules. <http://en.wikibooks.org/wiki/Monopoly/Official_Rules> [accessed Aug 4, 2009].
- Lehman, J., and J. Walker. 1975. *1000 ways to win at Monopoly*. Dell Books.
- Lehtonen, T., and H. Nyrhinen. 1992. On asymptotically efficient simulation of ruin probabilities in a Markovian environment. *Scandinavian Actuarial Journal*:60–75.
- Meyn, S. P., and R. L. Tweedie. 1993. *Markov chains and stochastic stability*. London: Springer-Verlag.
- Ross, S. M. 1996. *Stochastic processes*. 2nd ed. New York: Wiley.
- Stewart, I. 1996a. How fair is Monopoly? *Scientific American* 274:104–105.
- Stewart, I. 1996b. Monopoly revisited. *Scientific American* 275:116–119.

AUTHOR BIOGRAPHIES

ERIC J. FRIEDMAN is an associate professor in the School of Operations Research and Information Engineering at Cornell University. His main research interests are in game theory, computer networks and other problems arising on the Internet. He also has a strong interest in combinatorial games and their connections to dynamical systems theory and physics. He prefers dogs to cats. His web page can be found via <<http://www.orie.cornell.edu>>.

SHANE G. HENDERSON is a professor in the School of Operations Research and Information Engineering at Cornell University. He is the simulation area editor at *Operations Research*, and an associate editor for the *ACM Transactions on Modeling and Computer Simulation* and *Operations Research Letters*. He co-edited the handbook *Simulation* as part of Elsevier’s series of Handbooks in Operations Research and Management Science, and also co-edited the Proceedings of the 2007 Winter Simulation Conference. He likes cats but is allergic to them. His research interests include discrete-event simulation and simulation optimization, and he has worked for some time with emergency services. His web page can be found via <<http://www.orie.cornell.edu>>.

Friedman, Henderson, Byuen and Gutiérrez Gallardo

THOMAS BYUEN is an undergraduate in the School of Operations Research and Information Engineering. He can be reached at tb287@cornell.edu.

GERMÁN GUTIÉRREZ GALLARDO is a Master of Engineering student in the School of Operations Research and Information Engineering at Cornell University. He can be reached at gg92@cornell.edu.