Being Warren Buffett: A classroom simulation of financial risk

Dean P. Foster & Robert A Stine
Statistics Department, The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6340

October 29, 2004

Abstract

Students in business and other areas who are new to Statistics have a hard time making the connection between variance and risk. To convey the connection, we developed a classroom simulation. In the simulation, groups of students roll three colored dice that determine the success of three "investments". The simulated investments behave quite differently. The value of one remains almost constant, another drifts slowly upward, and the third climbs to extremes or plummets. As the simulation proceeds, some groups have great success with this last investment – they become the "Warren Buffetts" of the class. For most groups, however, this last investment leads to ruin because of variance in its returns. The marked difference in outcomes shows students how hard it is to separate luck from skill. The simulation also demonstrates how portfolios, weighted combinations of investments, reduce the variance. In the simulation, a mixture of two poor investments is surprisingly good. Rather than use arbitrary properties, we calibrated the returns on two simulated investments to mimic returns on US Treasury Bills and stocks.

1. Introduction

The definition of variance as the expected squared deviation from the mean often strikes students as capricious. Why square the deviations from the mean rather than use the absolute value? Why average the values? Without the machinery of maximum likelihood or concepts of asymptotic efficiency, one is left to vague, heuristic explanations. When dealing with money, however, the definition of variance is just right. Rather than make this connection with formulas and theorems, we have found it more useful and memorable to let students experience the effects of variance first-hand. After defining means and variances with some basic examples, we use this 'dice game' to show the importance of these concepts. The discussion of the simulation requires only basic properties of means and variances, with the most sophisticated property being that the variance of a sum of independent quantities is the sum of the variances.

The three investments in this simulation have different characteristics. One investment resembles an old-fashioned savings account whose interest has been adjusted for the effects of inflation. At the other extreme, a second investment matches our intuitive definition of being very risky. A third lies between these extremes.

We have students simulate the changing value of these investments by rolling three differently colored dice. We label the three investments *Red*, *White,* and *Green* because it is easy to find dice in these colors. Though we have tried to save class time by letting a computer roll the dice (it's easy to program the simulation in Excel, say), we have found that students find the results more impressive when they roll the dice themselves. Although the simulation and ensuing discussion consume only an hour and 20-minute class (it also works well divided into 2 one-hour classes), the lessons of the simulation are among those that students take away from our course. After this simulation, everyone appreciates the importance of variance when looking at data.

The following section describes the dice simulation. The third section describes the origins of the simulated investments and explains how portfolios improve investments by reducing variation. This section also introduces the notion of volatility drag to quantify the effects of variation. The concluding section returns to the theme of distinguishing luck from skill.

2. The Dice Simulation

2.1 Getting Started

Before describing the simulation, we break our students into teams and describe the three investments. Teams of 3 or 4 students seem about right. On each team, one person plays the role of nature (or the market) and rolls the dice. Another keeps track of the dice and reads off their values. The third records the outcomes; hopefully every team has a responsible member. Others can help out. We pass out a sheet like that suggested in Figure 1 to each group. This record-keeping page organizes the results of the simulation in a format useful in later steps. The two unlabeled columns provide space to compute the returns on a portfolio later in the exercise. We collect these sheets at the end of the simulation so that we can review the results in the next class.

	Multiplier				**Value**			
Round	**Green**	**Red**	**White**		**Green**	**Red**	**White**	
Start	1	1	1	1	1000	1000	1000	1000
1								
2								
3								

Figure 1. First rows of the data collection form used to record the value of the three investments simulated by rolling a red die, a white die, and a green die.

Once we have the class divided into teams, we pose the following question. We've found it useful to elicit a written preference from each team before starting the simulation. This gets them talking about the simulation and avoids too many "Monday morning quarterbacks" in the subsequent discussion. If a team has chosen an investment before starting the simulation, the team members seem more interested in following their choice as the simulation evolves.

Question 1. Which of the three investments described in the following table seems the most attractive to the members of your group?

Investment	*Expected Annual Return*	*SD of Annual Return*
Green	7.5%	20%
Red	71%	132%
White	0%	6%

Table 1. Expected value and standard deviation of the annual return on three investments to be simulated in the dice game.

Table 1 might require more explanation, depending on what has been taught. For a basic class that has not spent much time with random variables, we describe the table in this way. Suppose that you invest $1000 in one of these choices, say *Red*. From the table, you can expect the value of your investment to be 71% larger at the end of the first year, up to $1,710. Students find this sort of calculation quite reasonable, but have little intuition for how to think about the standard deviation – other than to know that the presence of a large standard deviation means that the results are not guaranteed.

We describe the table more precisely if students are familiar with discrete random variables. Random variables are a natural way to represent the uncertainty of the value of investments that, unlike bank accounts, can increase or decrease in value. The random variable that is most natural in this context is the return on the investment. For example, if the random variable R_j denotes the return on *Red* in round j, then at the end of the first year, the value of the initial $1000 in *Red* is

$$\$1000(1 + R_1)$$

If *Red* goes up by 10% in the first year, then $R_1 = 0.10$ and the $1000 grows to $1,100. The summary table shows that $E(R_j) = 0.71$ and $SD(R_j) = 1.32$. The table includes the properties of two more random variables, namely those that describe the returns on *Green* and *White*. At the end of one year, if we start with $1000 in each of these, we'd expect to have $1,075 in *Green* and $1,000 in *White* at the end of the year. Investing in *White* resembles putting money into a mattress.

Red is obviously the best choice among these three if we only consider the expected values. Because the expected value of a product of independent random variables is the product of expectations, we can easily find the expectations for each investment over a longer horizon. Over 20 years, we expect the initial $1000 invested in *Red* to be worth an astonishing

$$\begin{aligned} E\big(1000\,(1+R_1)(1+R_2)\cdots(1+R_{20})\big) &= 1000\,E(1+R_1)\,E(1+R_2)\cdots E(1+R_{20}) \\ &= 1000(1.71)^{20} \\ &= \$45,700,000 \end{aligned}$$

By comparison, the initial investment in *Green* grows on average to $2,653 and *White* remains at $1,000.

The standard deviations might lead students to question the wisdom of investing in *Red*, but having compared expectations it is difficult for most to see how to trade off the large expected return for the variation. The standard deviation of the return on *Red* is the largest of the three, SD(R_j)=1.32. The annual return on *Red* is 10 times the return on *Green*, but its SD is also 6.5 times larger. Few students appreciate the bumpy ride promised by *Red*.

To help the class appreciate the role of uncertainty, we use this simple example. Suppose that a graduate lands a good job that pays $100,000 per year. In the first year, the company does well and her salary grows by 10% to $110,000. The next year is leaner, and she has to take 10% cut in pay, reducing her salary down to $99,000. The average percentage change in her salary is zero, but the net effect is a loss of 1% of the starting salary over the two years. Figured at an annual rate, that's a loss of 0.5% per year. It turns out that this simple example is a special case of a more general property that captures how variance eventually wipes out investments in *Red*.

2.2 Running the Simulation

After this introduction, we start the simulation. We distribute three dice to each team. Each roll of all three dice represents a year in the simulated market, and the outcomes of the dice determine what happens to the money held in each investment.

The following table shows how the outcomes of the dice affect the values of the three investments. The number in each cell of the table gives the value of $1 invested in each at the end of the round. We find it helpful to project this table on a screen visible to the class as the simulation proceeds.

Outcome	**Green**	**Red**	**White**
1	0.8	0.06	0.9
2	0.9	0.2	1
3	1.05	1	1
4	1.1	3	1
5	1.2	3	1
6	1.4	3	1.1

Table 2. Value multipliers for the amount invested in each of the three simulated investments.

An example of the calculations clarifies the calculations of the returns from Table 2. Each investment begins with an initial value of $1000. As an example, suppose that on the first roll the dice show these outcomes:

(Green 2) (Red 5) (White 3)

Then the three investments after the first year are worth

Green: $\$1000 \times 0.9 = \$900$

Red: $\$1000 \times 3 = \$3000$

White: $\$1000 \times 1 = \$1000$

For the next roll, the values are compounded, starting from the amounts at the end of the first year. If the second roll gives

(Green 4) (Red 2) (White 6),

then the three investments are worth

Green: $\$900 \times 1.1 = \$990$

Red: $\$3000 \times 0.2 = \$600$

White: $\$1000 \times 1.1 = \$1100$

after the second round. Figure 2 shows the recording sheet after the first two rolls.

	Multiplier				Value			
Round	**Green**	**Red**	**White**		**Green**	**Red**	**White**	
Start	1	1	1	1	1000	1000	1000	1000
1	0.9	3	1		900	3000	1000	
2	1.1	0.2	1.1		990	600	1100	

Figure 2. Data table recording the outcomes of the first two rounds of the dice simulation of three investments.

At this point, we turn the class lose and let the simulation begin. We generally run the simulation for 20 or 25 “years” in order to have the long-term patterns emerge. If the simulation runs much longer, the *Red* investment becomes less and less likely to do well. Stopping after 20 or 25 rounds leaves a good chance that some team will be doing very well with *Red*.

An aside

Some students might question whether the multipliers in Table 2 really correspond to the means and variances shown in Table 1. When we have had time to develop random variables, we use a homework exercise that has students check that the multipliers do indeed match up to the prior means and standard deviations.

Some may not recognize that the multipliers in Table 2 are 1 plus the returns. For example, consider *Green*. If a roll of the green die gives the value 5, the multiplier 1.2 means that every dollar invested in *Green* grows to $1.20, a 20% increase. All of the values in the table of multipliers are just one plus the return on each investment for that roll of the corresponding die.

Roll	Probability	Return
1	**1/6**	**–0.2**
2	**1/6**	**–0.1**
3	**1/6**	**0.05**
4	**1/6**	**0.10**
5	**1/6**	**0.20**
6	**1/6**	**0.40**

Table 3. Probabilities for the returns on the Green investment.

Table 3 gives the distribution of the discrete random variable *G*, the return on *Green*. To recover the mean shown in Table 1,

$$E(G) = \frac{-0.20 - 0.10 + 0.05 + 0.10 + 0.20 + 0.40}{6}$$
$$= 0.075$$

The rest of Table 1 follows similarly.

2.3 Pink

As the class runs the simulation, we browse the room to see how the different teams are doing and make sure that they are doing the calculations correctly. Generally, the room gets a little noisy, particularly if there's a group for which *Red* is working nicely. *Red* triples in value half of the time, so there's a good chance that some team will do well with *Red* if the simulation is run 20 rounds.

After letting the class run the simulation for 20 or 25 rounds, we interrupt the chatter and pose another task. This task does not require more rolling of the dice. For this part, the students consider a hybrid investment that mixes the previous results for *Red* and *White*. We call this investment *Pink*.

To compute the value of *Pink*, we instruct the students to use the *previously recorded rolls* of the red and white dice. It's easiest to describe how to figure out what happens to the value of *Pink* with an example. *Pink* also begins the simulation with $1000. For the first round, using the same dice rolls as in the previous example (*Red*=5 for a multiplier of 3 and *White*=3 for a multiplier of 1), the value of the *Pink* becomes

$$\$1{,}000 \times \frac{3+1}{2} = \$2{,}000$$

Compounded in the second round (which had values *Red*=2 for a multiplier 0.2 and *White*=6 with multiplier 1.1), the result is

$$\$2{,}000 \times \frac{0.2+1.1}{2} = \$1{,}300$$

Figure 3 shows the data recording form with values for *Pink* added. It is important that students average the multipliers, not the values, for *Red* and *White*.

	Multiplier				**Value**			
Round	Green	Red	White	**Pink**	Green	Red	White	**Pink**
Start	1	1	1	1	1000	1000	1000	1000
1	0.9	3	1	2	900	3000	1000	2000
2	1.1	0.2	1.1	0.65	990	600	1100	1300

Figure 3. Data recording with the values for the mixture Pink added to the calculations.

Before turning them lose again, we pose some questions to make them think before calculating. Because *Pink* mixes the returns of *Red* and *White*, most students expect it to be a mix of bad and boring and not do well.

> ***Question:*** **Before you compute your outcomes, discuss *Pink* with your team. How you expect *Pink* to turn out? Do you think it will be better or worse than the others?**

We have found it useful to circulate through the class as students figure out the results for *Pink*. A common mistake is to average the final values for *Red* and *White*. This error gives a very different answer than obtained by averaging the returns. Often, as they do the calculations, teams suspect that they have done something wrong because *Pink* does so well!

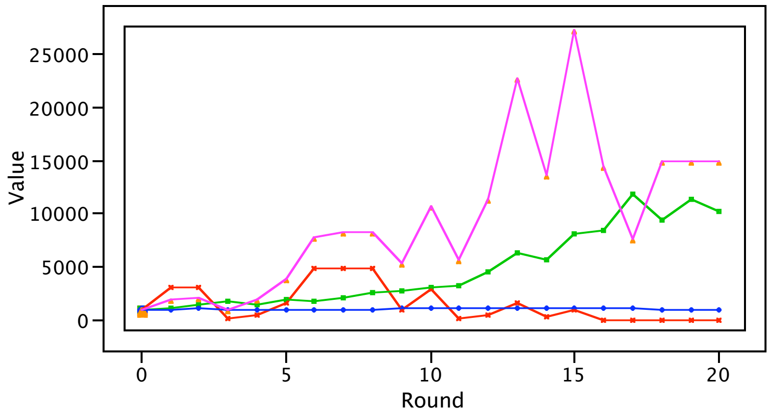

(a)

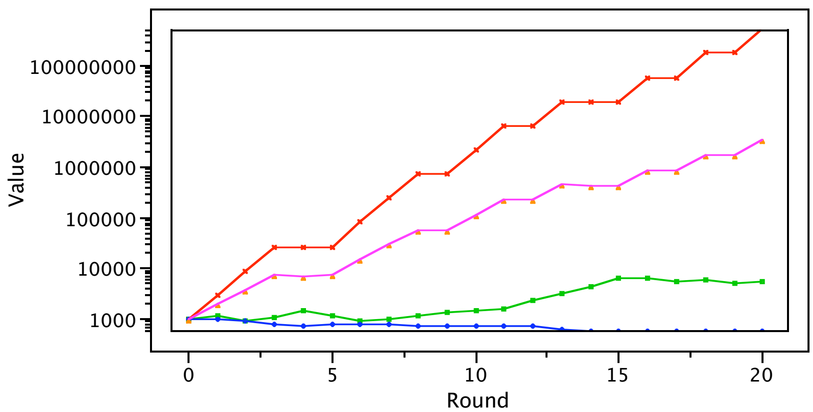

(b)

Figure 4. Timeplots of the values of 4 simulated investments in the dice game. The outcomes in (a) are typical whereas those in (b) show what happens when a team becomes the Warren Buffetts of the class.

2.4 Collecting the results

Once it seems that most have finished the calculations for *Pink*, we query the teams for their results. To maintain flow of the discussion, we find it simplest to track the outcomes for the investments on a transparency that we augment as teams announce results.

For most, *Red* becomes nearly worthless as in Figure 4a. *Red* looks great at first, drops, then recovers before collapsing. *Green* rises slowly, but steadily. *Pink* is more volatile than *Green*, but closes with a larger value. As we poll the class, most say that *Red* fell to pennies. For a few, however, *Red* does very well as in Figure 4b. For a simulation with 20 rounds, direct calculation shows that the probability that *Red* is worth $10,000 or more is about 20%,

$$P\left[(1+R_1)(1+R_2)\cdots(1+R_{20})>10\right]\approx 0.19$$

Similar calculations show that the chance for becoming a millionaire with *Red* is a bit larger than 5%. It comes as quite a surprise to the rest of the class when a team announces that their value in *Red* is, say, $10,000,000. We call this team the "Warren Buffetts" of the class. Business students generally know the name by reputation.

As for the other two original investments, *White* generally drifts downward but remains close to the initial $1,000 stake. *Green* shows a steady return and is usually the best of the original alternatives.

Pink presents the students with their second surprise. Across the class – with the exception of the Warren Buffetts – *Pink* usually results in the highest value at the end of the simulation as in Figure 4a. Though *Pink* mixes two investments that are individually poor choices, this simple mix of *Red* and *White* works very well. That frequently seems impossible to the students, leaving many to question how the average of two poor investments can become so valuable.

3. Discussing the Simulation

3.1 Why these multipliers?

We open our discussion of the dice game by linking the simulated investments to real investments. *Green*, which does the best for most teams until they discover *Pink*, performs like the US stock market when adjusted for inflation. *White* represents the inflation-adjusted performance of US Treasury Bills, the canonical "risk-free" investment. We made up *Red*. We don't know of any investment that performs like *Red*. If you know of one, please tell us.

The timeplot in Figure 5 summarizes the history of stocks and Treasury Bills in the US from 1926 through the end of 2003. Both series are monthly. For stocks, this plot tracks the value of one dollar invested in January 1926 in a value-weighted portfolio of the US stock market. (A value-weighted portfolio, such as the S&P 500, buys stock in proportion to their capitalized value. Alternatives such as the Dow-Jones Index simply buy one share of each.) This plot also shows the value of a $1 investment in 30-day Treasury Bills. The Y-axis in the timeplot uses a log scale. When plotted on a log scale, geometric growth appears as a straight line.

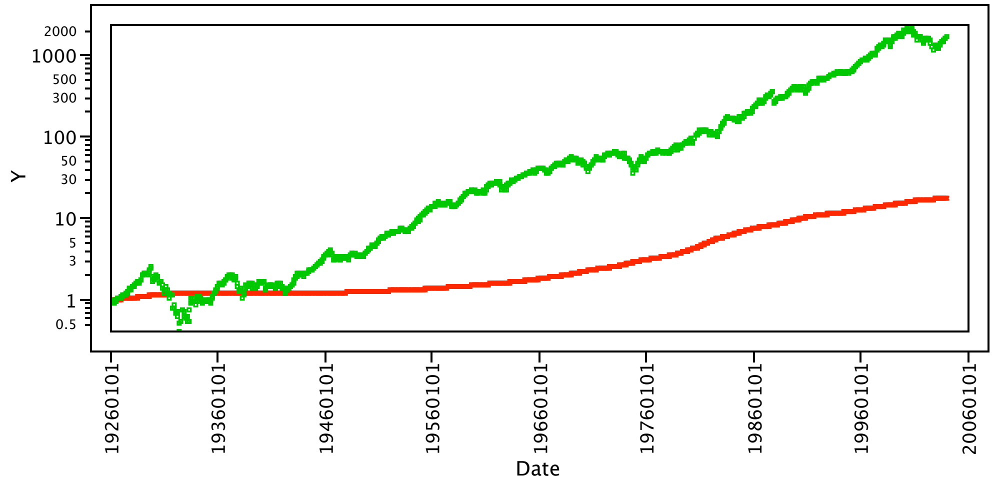

Figure 5. Return on a 1926 investment of $1 in the stock market and in Treasury Bills. (red x = $1 in 30-day Treasury Bills, green o = $1 in stocks).

This plot is a little misleading because it ignores inflation. Although inflation has recently been quite low, it exceeded 15% or more annually in the past. To adjust for inflation, Figure 6 shows the cumulative values after subtracting the rate of inflation from the growth of $1 investments in the stock market and Treasury Bills. To measure inflation, we used month-to-month changes in the Consumer Price Index in the US. Once adjusted for inflation, Treasury Bills do not appear so risk-free. The inflation-adjusted value of the investment in Treasury Bills declined for several long periods. Net of inflation, the $1 invested in Treasury bills ends at $1.65. The $1 invested in the stock market ends up at $150, even allowing for the Great Depression and the dot-com bust.

Returns are the key random variables in the dice simulation. The third timeplot in Figure 5 shows the monthly returns for stocks and Treasury Bills, net of inflation. (Again, we subtracted out the rate of change in the Consumer Price Index). The month-to-month variation of the returns on Treasury Bills is so much smaller than the variation in returns on stocks that this sequence becomes almost invisible in the plot. Several familiar events are also apparent. On the left, starting in the late 1920s and running through the 1930s, is the Great Depression. Returns on the market were incredibly volatile during that period. In 1933, the market dropped almost 30% in one month. Less well known is that about a year later, the market increased by about 40% in each of two months. Following WW II, the returns on stocks became rather stable

– at least in comparison to the volatility during the Depression. You can also identify other big shocks such as the drop in October 1987.

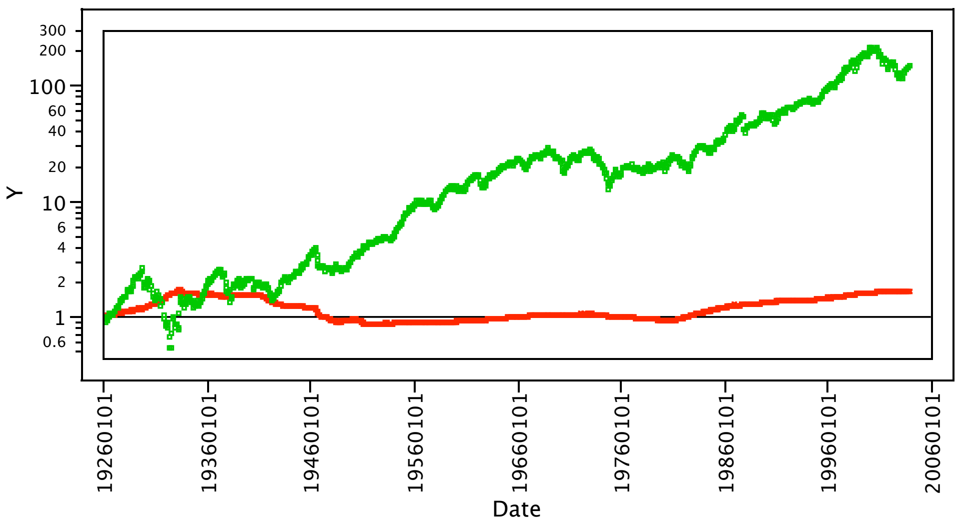

Figure 6. Return on a 1926 investment of $1 in the stock market and in Treasury Bills after adjusted for inflation by subtracting the rate of change in the Consumer Price Index.

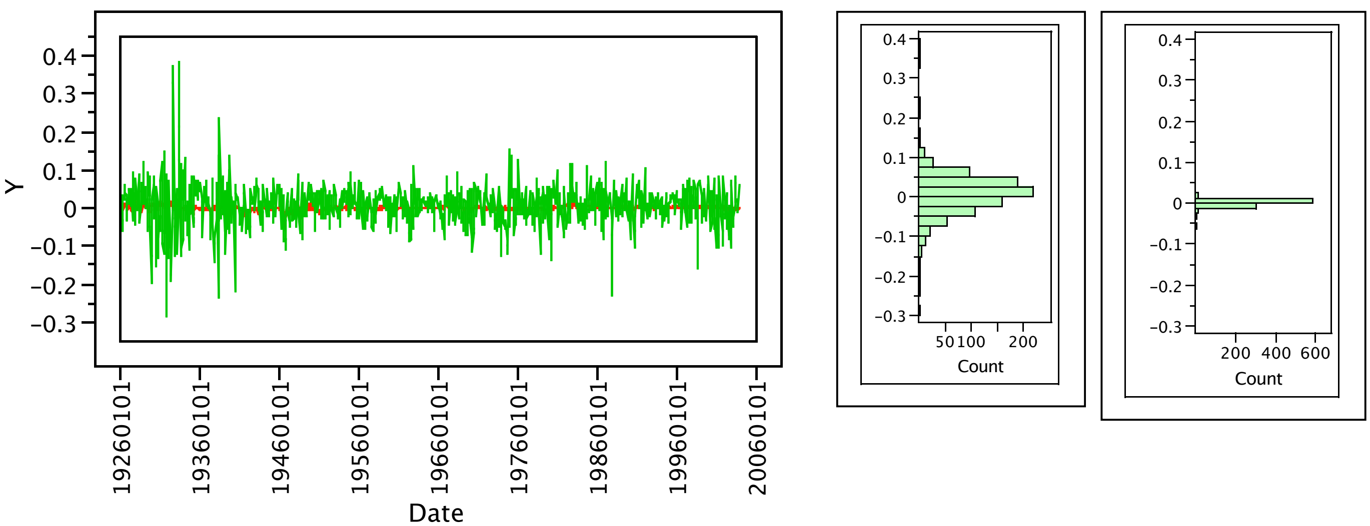

Stock Returns *T-Bill Returns*

Figure 7. Timeplot and histograms of the inflation-adjusted monthly returns for stocks and Treasury Bills in the US.

Both sequences of returns resemble series of independent observations. After converting to returns, histograms seem like reasonable summaries of the data, at least if one ignores the bunching of periods of high volatility. While important in modern financial modeling, we avoid the complications of conditional volatility when introducing these ideas. The two histograms beside the timeplot in Figure 7 compare the returns on the stock market to the returns on Treasury bills, both adjusted for inflation and shown on the common scales from –0.3 to 0.4 (monthly returns from –30% up to 40%). The monthly returns on the stock market need that range; the monthly returns on Treasury bills never venture far from zero. (We like the placement of the histograms beside the timeplot to emphasize that the histograms simply count the number of points in different horizontal slices in the time plot.)

	Stocks	*T-Bills*
Mean	0.0069	0.00055
Std Dev	0.0552	0.00545
Variance	0.0031	0.00003
N	936	936

Table 4. Means, standard deviations and variances of the monthly inflation adjusted returns on US stocks and Treasury Bills.

The summary statistics in Table 4 show that the excess annual return on investments in the stock market above inflation from 1926 through 2003 averaged about 12(.0069) = .0828, slightly above 8%. Returns on Treasury Bills have been essentially flat, just keeping pace with inflation. The average net return above inflation for T-bills has been 12(.00055) = .0066, about 2/3 of one percent. The two returns obviously have very different variation. The month-to-month standard deviation for stock returns is 0.0552. Assuming independence over time, this monthly standard deviation implies an annual standard deviation near $\sqrt{12}(0.0552) = 0.192$, about 19%.

Comparing these statistics for stocks and Treasury Bills to the summary table associated with the dice simulation, we see that *Green* resembles the market. In fact, the expected return on *Green* at 0.075 is very close to the *excess return* of the stock market, the return on the market minus the risk-free rate, 0.0828 – 0.0066 = 0.0762. If you borrow the money that you invest in the market (and you can borrow at the same rate of interest as the US government), this is the return that you would net after paying the interest on your loan. Similarly, *White* is the risk-free rate after inflation – zero. We left more variability in *White* than in the returns on Treasury Bills so that it would not remain constant in the simulation.

3.2 The Success of Pink: Volatility Drag

Before looking at the numbers, it is essential that students understand that the value of *Pink* is not a simple average of *Red* and *White*. The returns on *Pink* average those on *Red* and *White*, but one does not get this performance by starting with \$500 in *Red* and \$500 in *White* and leaving it there. *Pink* requires that the portfolio be *rebalanced* at the end of each period so that half of the current value is invested in *Red* and half in *White*. Returning to the illustrative calculations, at the end of the first round, the initial \$500 invested in *Red* grows to \$1500 and the \$500 invested in *White* holds its value. Before the next round, the portfolio needs to be put back into 50-50 balance; that is, we need to move \$500 from *Red* into *White*, so that each has \$1000 at the start of the next round. This "protects" some of the winnings from *Red* in the prior round from subsequent volatility. When the next roll wipes out 80% of the value of *Red*, it only reduces the \$1000 left in *Red* down to \$200. The other \$500 produced by *Red* in the first round remains safely in *White*.

As a first step in understanding the success of *Pink*, we need the mean and variance of its returns. Because the return on *Pink* is the average of those on *Red* and *White*, students willingly accept that the mean return on *Pink* is (0.71+0)/2 = 0.355. Finding the variance is a little harder and requires that students know two basic manipulations for variances:

(a) Constants factor out with squares, $\text{Var}(c\,X) = c^2\,\text{Var}(X)$, and

(b) For independent random variables X and Y, variances of sums are sums of variances, $\text{Var}(X+Y) = \text{Var}(X) + \text{Var}(Y)$.

Because we simulate the returns in this example using separate dice, it should be clear that the returns on *Red* and *White* are independent. (Along with the invention of *Red*, the independence of the returns is a simplifying aspect of the dice simulation that differs from the real world. Returns on real investments are usually correlated, complicating the analysis of a portfolio.) Using (a) and (b), the variance of returns on *Pink* are easily found to be (Table 5)

$$
\begin{aligned}
Var(Pink) &= Var\left(\frac{R+W}{2}\right) \\
&= \frac{Var(R+W)}{4} \\
&= \frac{Var(R)+Var(W)}{4} \\
&= \frac{1.32^2+0.06^2}{4} = 0.436
\end{aligned}
$$

It is worth mentioning that although *Pink* sacrifices half of the expected return of *Red*, it also reduces the variance by a factor of 4.

Color Die	**E Return**	**Variance**	**Avg–Var/2**
Green	0.075	$0.2^2 = 0.04$	0.055
Red	0.71	$1.3^2 = 1.69$	–0.135
White	0	$0.06^2 = 0.0025$	–0.002
Pink	0.355	$0.65^2 = 0.4225$	0.144

Table 5. Mean, variance and volatility adjusted return of the four simulated investments in the dice game.

A simple expression based on the mean and variance of the returns shows how variation eats away at the value of an investment. Because it reflects how variation (or volatility) eats away at the value of an investment, we refer to this adjusted return as the *volatility adjusted return*. This is also known as the long-run return on an investment. The formula for computing the volatility adjusted return is simple:

Volatility-adjusted return = Long-Run Return

= Expected Annual Return – (Variance of Annual Return)/2

This penalty for variation is sometimes called the *volatility drag*.

Before we take a closer look at this formula, Table 5 shows the calculations for the dice game. The last column in this table shows the volatility adjusted return on *Green, Red*, *White* and *Pink*. Not surprisingly, *Pink* is most attractive, with almost three times the volatility-adjusted return of the stock market. Even though *Red* is a big loser for most teams, mixing it with *White* reduces the variance and produces a huge win. (As a little follow-up exercise, you might want to have students consider the following: What is the optimal mix of *Red* and *White*? That is, what proportions of *Red* and *White* produce the highest volatility-adjusted return? When considering investments with independent returns, investors should purchase *some* of any investment known to have positive mean. It becomes a question of how much.)

Depending on the level of the class, we spend more or less time describing the origins of the formula for the volatility adjusted return. In an introductory class, we can get this formula for the volatility drag from our simple example of volatile changes in salary. Think of the changes in salary as a random variable, with half of the probability on + 0.10 (up 10%) and the other half on –0.10 (down 10%). The expected value of this random variable is zero. Its variance is the average squared deviation from zero, simply $0.1^2 = 0.01$. Now look back at the example. The salary dropped by 0.5% per year, or 0.005. That's half of the variance. On average, each year of employment reduces the salary by half of the variance of the percentage changes. That is precisely the adjustment provided by the volatility drag.

For students who are familiar with the weak law of large numbers and Taylor series, we use a more rigorous argument for the volatility-adjusted return. Assume that the initial value of an investment is W_0. We set this to $1000 in the dice game. Label the return during year t as R_t. Thus, the value at the end of the first year is the initial value times one plus the return earned in the first year, or

$$W_1 = W_0\,(1+R_1)$$

In general, by the end of year T the value is

$$W_T = W_0\,(1+R_1)\,(1+R_2)\,\ldots\,(1+R_T)$$

Taking logs reduces this product to a sum that is easier to manipulate,

$$\begin{aligned}\log W_T &= \log W_0 + \log(1+R_1) + \log(1+R_2) + \cdots + \log(1+R_T) \\ &= \log W_0 + \sum_{t=1}^{T} \log(1+R_t) \\ &= \log W_0 + T \times \frac{\sum_{t=1}^{T} \log(1+R_t)}{T} \\ &\approx \log W_0 + T \times E\log(1+R_t)\end{aligned}$$

The last approximation only applies for large T by the weak law of large numbers, so that the average of the observed values of the random variable has settled in on the expectation. The expected value $E\log(1+R_t)$ is called the expected log return in Finance, yet another name for the long-run growth rate.

To get the expression for the volatility drag, we prefer the approximation

$$\log(1+x) \approx x - \frac{x^2}{2} \quad .$$

This allows us to argue that

$$E\log(1+R_t) \approx E\left(R_t - \frac{R_t^{\,2}}{2}\right)$$
$$= E(R_t) - \frac{E(R_t^{\,2})}{2}$$
$$\approx E(R_t) - \frac{Var(R_t^{\,2})}{2}$$

So long as the average returns are small, $E\,R_t^2$ is about the same as $E(R_t - E\,R_t)^2 = \text{Var}(R_t)$. Alternatively, one can avoid these approximations by making the strong assumption that returns follow a lognormal distribution. That argument, however, requires an assumption about the distribution of the returns that is hard to verify in practice.

Finally, in a very advanced class, one can use this discussion to motivate the importance of the Shannon-Brieman-MacMillan theorem (for example, see Chapter 15, Cover and Thomas 1991). But we'll not do that here!

4. Conclusion

What about those Warren Buffetts?

We developed this simulation to show off the importance of the variance in assessing the long-term value of investments. *Pink* offers a simple illustration of how one can gain positive long-run returns by using a portfolio that sacrifices expected returns to reduce the variation.

Having used this simulation in undergraduate and MBA classes for several years at Wharton, we have come to appreciate the important message conveyed by the Warren Buffetts of the class. These are the few teams that, unlike most others, end the simulation with *Red* reaching an astonishing large value. It comes as quite a surprise to the rest of the class to discover, as we collect the final values from the teams, that some of their classmates have had huge success with *Red*. The differences are not slight either. For a team whose $1000 in *Red* has shrunk a few pennies, it seems impossible that another team's investment in *Red* is worth $10,000,000 at the end of the game. We even used to run the simulation longer, hoping that volatility would wipe out these lucky winners. We have, however, come to realize that these anomalies allow us to present the students with an important question.

What makes them believe that the real Warren Buffett was not just lucky? After all, with millions of investors seeking profits from the stock market, could it be that Warren Buffett simply "got lucky." There is usually considerable resistance from fans of the "Sage of Omaha", but even they have to concede how difficult it is to separate a knowledgeable strategy from a lucky strategy. In the dice simulation, all of the teams use the same "strategy" for *Red* and rolled the dice themselves; nothing is hidden in a mysterious random number generator. They all start with $1000 in *Red*, but only a lucky few end the game appearing a lot smarter than the others. In the dice game, they can all see that it was simply luck that produced the Warren Buffetts.

We are careful not to say that Warren Buffett became successful by sheer luck. We simply point out the difficulty in separating skill from luck, a problem that bedevils investors in hedge funds and requires methods outside the scope of this paper that we plan to describe elsewhere.

References

Cover, T. M. and J. A. Thomas (1991). *Elements of Information Theory*, Wiley, New York.