SCHULTZ COLLINS LAWSON CHAMBERS, INC.
INVESTMENT COUNSEL

Black Swans and Albino Crows

Patrick J. Collins Ph. D., CFA, CLU

This paper reflects the opinion of the author only and does not necessarily represent the view of Schultz Collins Lawson Chambers, Inc. or any of its principals or employees.

Schultz Collins Lawson Chambers, Inc
455 Market Street
Suite 1450
San Francisco, CA 94105

P415-291-3000
F 415-291-3015

May 2014

In October of 2008, the *Wall Street Journal* reported that Alan Greenspan called the recent market collapse a once-in-a-century crisis. Investors saw the S&P 500 Stock Index retreat from a high of 1,565.15 on October 9, 2007 to a low of 676.53 on March 9, 2009. This essay explores two topics of interest to trustees providing investment management services:

1. Are large declines in stock prices truly rare events?
2. What is a reasonable assessment of the recent market decline and recovery?

Let's define a large stock price decline as a 50% or more decrease in the value of the S&P 500 Stock Index from a price peak to a price trough. As noted, in the recent decline, the peak was 1,565.15 and the trough was 676.53 for a total decline of approximately 57%.[1] Of course, portfolio withdrawals would have pushed the decline in dollar value beyond 57%. This essay, however, does not consider the impact of withdrawals, taxes, fees, trading costs, or loss of purchasing power. [2]

Historically (1973 through 2013) the average annual return of the S&P 500 is 11.88% with an annual standard deviation of 18.20% (this translates into an annual compound return of 10.26%).[3] An average annual return of 11.88% over both good and bad economies is good, and is one reason why investors are attracted to the stock market. During this period, investors suffered through the OPEC oil shock, the Reagan recession, the Market Crash of October 1987, the Bush One recession, the Asian Banking Crisis, the NASDAQ tech-stock meltdown, and the real estate price collapse / global recession.[4]

Often, when investors consider the likelihood of significant declines, they think in terms of a normal distribution which has a well-known bell curve shape. The normal distribution is symmetric about its arithmetic mean. The symmetry indicates that the likelihood of a significantly bad event equals the likelihood of an equally significant good event where good and bad are defined as equal distances above and below the mean. Quite simply, in a bell curve distribution, risk is fully described by standard deviation because this statistic measures the probability of both significant upside and downside deviations from the mean.[5] The following graph depicts these concepts:

---

[1] Nate Silver in The Signal and the Noise, Penguin Press (New York, 2012), p.33 reports: "In December 2007, economists in the *Wall Street Journal* forecasting panel predicted only a 38 percent likelihood of a recession over the next year. This was remarkable because, the data would later reveal, the economy was *already in recession* at the time."

[2] For an analysis of trading costs, inflation, taxes, fees and expenses, see Collins, Patrick J., Lam, Huy & Stampfli, Josh, "Monitoring and Managing a Retirement Income Portfolio: Sufficiency, Solvency, and Investor Utility," (forthcoming). See also, Zhow, Guofu & Zhu, Yangzi, "Is the Recent Financial Crisis Really a 'Once-in-a-Century' Event?' Financial Analysts Journal, (January/February, 2010) pp. 24-27. This study concludes that, as the expected long-term stock market value increases, so also the probability of realizing large declines in market value. Under the authors' model, there is an accumulating probability of significant price declines.

[3] Each dollar invested in 1973 grew by 10.26% per year through 2013. The final value of each dollar initially invested on January 1, 1973 is $54.94 on December 31, 2013.

[4] Although the 11.88% was certainly not a free lunch, at the risk of stating the obvious, if you avoid all risk, you can expect to earn only the risk-free rate of return. If you take risk, you expect to earn a premium above the risk-free rate but the expectation is not a guarantee. During this period, the compound return on a constant-maturity, one-year U.S. T-Bill is 5.79%. The annual compound reward (risk premium) for investing in the S&P 500 is 4.47% (10.26% - 5.79%). Although it is easy to jump to the conclusion that an expected equity risk premium must become a certainty given a long-term planning horizon, this conclusion is not warranted. A dollar invested today in a stock is not worth more than a dollar invested in a 30-year bond. It is unlikely that a bond owner would pay a premium for swapping the bond return for the stock return. Despite these facts, one often hears the old chestnut that stocks must always outperform bonds in the long run.

[5] The mean is the central return tendency of the distribution. Economists refer to the mean as "expected return." Expected return, however, is a statistical concept ***not an investment prediction***. If the expected return over a 20-year period is 10%, that means that if you had thousands of independent 20-year periods each earning a unique return, most of the 20-year returns

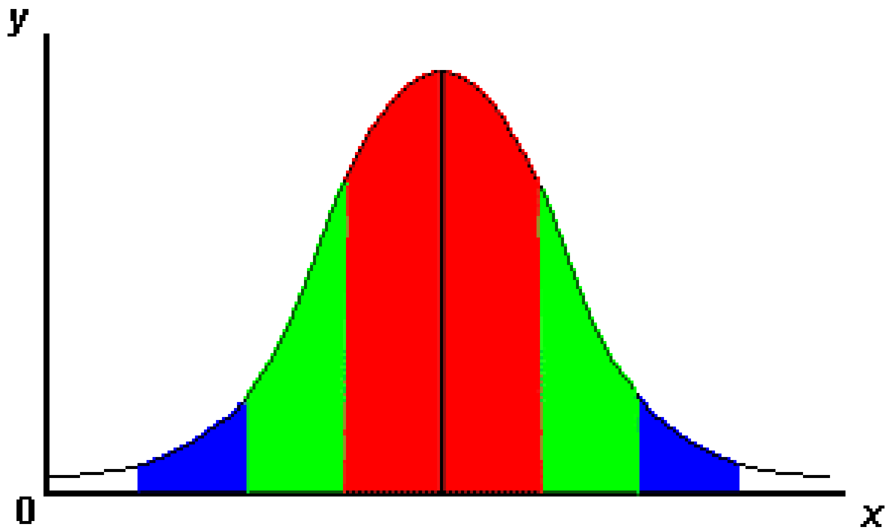

In a normal, or 'bell-shaped,' return distribution, one standard deviation away from the average in either direction on the horizontal axis (the red area on the chart above) includes roughly 68 percent of monthly returns. Visually, this means that the red area under the bell curve is approximately two/thirds of the curve's total area. Stated otherwise, there is a two/thirds probability that the return in any single period will fall into the red-colored region of the distribution. Two standard deviations away from the mean (the red and green areas together) include roughly 95 percent of monthly returns. Three standard deviations (the red, green and blue areas) account for over 99 percent of monthly returns. Given the symmetrical nature of the distribution, here is how trustees commonly use standard deviation as a guide to assessing S&P 500 stock index risk:

Expected Annual Return of the S&P 500 Stock Index: 11.88%
Annual Standard Deviation of the Index: 18.20%
Expected Range of Annual Returns at a 68% Probability: 11.88% ±18.20%x1 = -6.32% to +30.08%
Expected Range of Annual Returns at a 95% Probability: 11.88% ±18.20%x2 = -24.52% to + 48.28%
Expected Range of Annual Returns at a 99% Probability: 11.88% ±18.20%x3 = -42.72% to +66.48%
Probability of Annual Return less than -42.72%: ½ of 1%
Probability of Annual Return greater than +66.48%: ½ of 1%.

Under the normal distribution model of stock price risk, a yearly decline of 50% or more in the S&P 500 is an extremely low probability event and, having experienced one of these "black swan" moments, we should not expect a reoccurrence for many years.[6]

---

should be close to 10%. It is a horse of a different color to predict the actual return that an investor will earn over any single 20-year period. The statement that the average annual return of the S&P 500 is 11.88% over the previous 41 years is, of course, empirically correct. It is based, however, on only a single realized return path. Investment expectations based *solely* on past results are, in fact, expectations derived from sample of one historical sequence.

[6] Given a ½ of 1% chance of a 42.72% decline in any year, the odds of such a decline over any holding period are $[1 - (.995)^x]$ where 'x' is the length of the holding period. For example, the odds of a 42.72% decline over a twenty-year holding period are 10%; over a hundred-year period are approximately 40%. As the length of the holding period increases, so does the likelihood of low probability events. The term "Black Swan" event was popularized by Nassim Taleb's book The Black Swan: The Impact of the highly Improbable, (New York: Random House), 2007. It should be noted that a Black Swan event is not an outcome subject

Although this asset price risk model is comforting, the story cannot end at this point. What if the model is incorrect in the sense that it fails to capture important econometric aspects of the S&P 500's empirical price series? How incorrect could it be?[7] For any distribution, irrespective of its shape, the theoretical bound on possible deviations from the mean is set by a mathematical expression known as Chebyshev's inequality. Specifically, the inequality sets an upper bound limit of 1/(#standard deviations)^2 for the probability of deviating 'n' standard deviations above or below the mean. For the lower tail alone, the probability of a major decline cannot exceed 1/[(#standard deviations)^2+1]. Under the bell curve risk model, the chance of a loss greater than -42.72% (a three standard deviation event) is 1 in 200 (½ of 1%). Under Chebyshev's inequality, the chance of a loss greater than -42.72% is $1/(3^2 + 1)$, or 1 in 10. This suggests that the real risk of the S&P 500 Stock Index may lie somewhere in the middle.

The normal curve risk model is primarily useful as a single period model. However, one drawback of using the model to assess multiperiod risk is that the normal distribution does not allow for volatility clustering over time. Studies of the S&P 500 time series often note that although volatility tends to subside over time, nevertheless it exhibits various degrees of persistence. That is to say, high volatility months tend to be followed by high volatility months, and low volatility months tend to be followed by low volatility months. When volatility manifests itself in downside returns, the bad returns have a tendency to keep-on-coming. Market participants become nervous, commentators start asking investment gurus to predict the market bottom, heightened volatility mandates a higher discount rate for stock price valuation, and so forth. Stock prices exhibit serial correlation (bad returns in the next month are positively correlated to bad returns in the last month), and, if too many investors head for door at once, downside risk turns into a market crisis.

To recap, trustees would like to have an adequate picture of risk prior to committing money to uncertain ventures. Historical returns provide an accurate measure of realized risk, but are only a useful guide to the future if you believe that future economies will be exactly like past economies. The attempt to parameterize risk by assuming that the possible distribution of future returns fits into a normal bell-curve shape (fully determined by two parameters—mean and standard deviation) is a good way to break out of the "groundhog day" risk-evaluation syndrome wherein the past incessantly repeats. However, the bell curve risk model fails to capture important stock price behaviors. For example, the worst day of the October 1987 market crash poses an explanatory dilemma under this risk model. A *daily* standard deviation to the downside greater than 20% translates into a yearly standard deviation of approximately 382%. Assuming a normal distribution an event of this magnitude is unlikely to occur within the total amount of time from the beginning of the universe to the present. Additionally, the bell curve risk model cannot begin to predict the wave of economics shocks from 1973 through the present.[8] It fails in that it cannot offer a risk model that incorporates dynamic volatility clustering,

---

to a probability measure. That is to say, such an event is both rare and unpredictable. In previous history, however, events like the appearance of a comet were "Black Swans." Today, however, we can predict such events with great accuracy. Thus, there is a lively debate over whether Black Swan event arrivals are unknowable or whether we simply lack sufficient predictive skills at this moment in time.

[7] The normal distribution incorporates several assumptions that do not apply to stock price series. These include: (1) time-invariant mean and standard deviation parameters, (2) period returns that are independent and identically distributed, and (3) symmetrical returns above and below the mean. The model is parsimonious and easily communicated, however, and these reasons undoubtedly contribute to its widespread use. In reality, extreme returns occur at a frequency far greater than that predicted by a normal distribution model. However, in a highly mathematical branch of finance known as "continuous time finance" a log-normal distribution approximation is a more credible model of asset price risk. The second part of this essay refers to several theorems commonly used in continuous time finance studies.

[8] This is also true for Monte Carlo simulations based on the assumption of normal or joint multivariate normal return distributions.

reasonable downside risk frequency prediction, or other elements important for risk assessment and control.

Advanced risk metrics generally resolve the above-listed deficiencies by use of two approaches that are more dynamic than either historical track record or parameterized curve fitting:[9] (1) Vector Auto-Regression Models [VAR], or (2) Markov Regime Switching Models. The VAR approach demands that the model user pre-specify the variables that are key drivers of stock price change. Variables can be macroeconomic variables (inflation, GDP, Industrial Production, etc.), accounting variables (investment fundamentals such as income statement and balance sheet ratios), statistical variables (lead/lag relationships), and so forth. A VAR model posits a linear relationship between past values for each variable and future stock price returns and volatilities. Relationships are expressed mathematically (e.g., multiple equations outline the relationship between lagged values of a variable and its likely future evolutions). The mathematical “blueprint” for future return evolutions is purely a-theoretical in that, unlike the bell-curve risk model, it does not impose any return distribution on the data. Although VAR models produce a less distorted picture of investment risk, good VAR models are largely confined to university economic departments and are rarely found in the investment practitioner community.

One alternative to VARs is a Markov Regime Switching Model. Analysis of stock returns indicates that critical parameter values (mean returns, volatility, and correlations) differ significantly depending on whether the stock market is in a Bull or Bear regime. A regime switching model simulates a Markov transition process with a probability [P] of remaining in a given regime conditional on being in that regime currently, and a probability [1-P] of transitioning to the other regime. This model allows both the equity risk premium and return volatility to vary over time. The model also produces simulated investment results that reflect dynamic correlations and, most importantly, time varying returns and volatilities.

This essay employs a Regime Switching Model to assess the downside risk of owning the S&P 500 stock index.[10] This experiment examines holding periods of one through ten years as well as a twenty-year holding period. Specifically, the risk model simulates 5,000 portfolio evolutions for each of the 11 holding periods to determine the likelihood of a peak-to-trough decline in portfolio value equal to or greater than 50% in dollar value.[11] The following table summarizes the results:

Holding Period	Likelihood of at least a 50% decline in peak-to-trough portfolio dollar value.
1 Year	0.3%
2 Years	3.6%
3 Years	6.9%
4 Years	9.7%
5 Years	12.6%
6Years	16.4%
7 Years	19.1%
8 Years	21.5%

[9] In addition to the normal and log-normal distribution curves, financial economists use a variety of distributions to model risk. These include the more “fat-tailed” distributions like the Student’s t-distribution and the Levy stable distribution models. More advanced models incorporate a mixture of distributions. In general, single distributions and distribution mixtures have difficulty modeling risk in a multiperiod context in that they do not adequately reflect the time-varying nature of important parameters.

[10] It does not make any adjustments for inflation or investment expenses. Index returns are based on the period 1973 through 2013.

[11] Under all risk models, over a 100-year holding period, a peak-to-trough decline of 50% or more is likely. Specifically, the bull/bear risk model indicates a likelihood of 95.9% for such an event over a century of simulated investment experience. The normal curve risk model suggests a likelihood of approximately 40%. At some point in every century, the market may turn nasty and a lot of money will be taken off the table.

9 Years	24.0%
10 Years	26.4%
20 Years	47.2%

The good news is that it is not too likely that a portfolio will suffer a 50% decline in value right out of the box. It's never fun to see a bear market eat into profits; but it's downright brutal to see it destroy the original investment principal. The eye-catching information is the rapid increase in the probability of a virulent bear market as the length of the planning horizon increases. A bull/bear Markov regime switching risk model suggests that investors might expect a one-in-three chance of such a market over any decade and a better-than-even chance over any twenty-year holding period. These results are in stark contrast to those suggested under the normal curve risk model which pegged the chance of a decline of 42% or greater over a twenty-year holding period at approximately 10%.

We cannot say for sure whether the Markov regime switching model underestimates or overestimates risk because no one knows the true return generating process. We only know the outcome of the process as it manifests itself in realized returns over a specific historical period. What we can say is (1) our confidence in the Markov model is higher because it more closely emulates the observed behavior of financial asset time series, and (2) it is probably not a good idea to bet your fortune entirely on the performance of the S&P 500 despite the time honored maxims about the safety of owning a "blue-chip" stock portfolio. The geometric distribution with parameter p = .58 for each 20-year period is a "back-of-the-envelope" approximation[12] of the odds of a 50% or greater peak-to-trough decline throughout a span of 100 years. This distribution shows the increase in the likelihood of a significant S&P 500 stock decline as time unfolds. We define the likelihood of a decline as: [1 - the probability of not realizing a decline]. Thus, for the first 20-year period, the probability of a 50%+ decline is [1 - .42 = .58]; for the first and second 20-year periods the probability of a 50%+ decline is [$1 – (.42)^2$ = .82]; for the first, second, and third 20-year periods the probability of a 50% decline is [$1 – (.42)^3$ = .93]; and so forth. By the end of the fourth period, the likelihood increases to 97%, and, by the end of the fifth period (the full 100-year time span), the likelihood approaches certainty (99%). Although the geometric distribution does not tells us the length of time required for market recovery, it is obvious that concentrating trust portfolios in the equity of a single nation is hazardous to wealth.[13]

Trustees seek to use current financial assets to discharge future financial liabilities. A liability can be as narrow as a contractual obligation (a loan repayment on a trust-owned home) or as broad as a general financial objective (pay for a child's education, support lifetime income, etc.). Sound investment policy, however, presupposes some understanding of the underlying process of asset price change.

Consider, for a moment, some commonly used models of asset price change. The following chart depicts four price-change processes: (1) White Noise, (2) Deterministic, (3) Random Walk, and (4) Drift with Diffusion.

[12] The Geometric distribution assumes independence of events and constant probabilities.
[13] The May 2009 level of Japan's Nikkei Stock Index was approximately equal to the index's value in 1982.

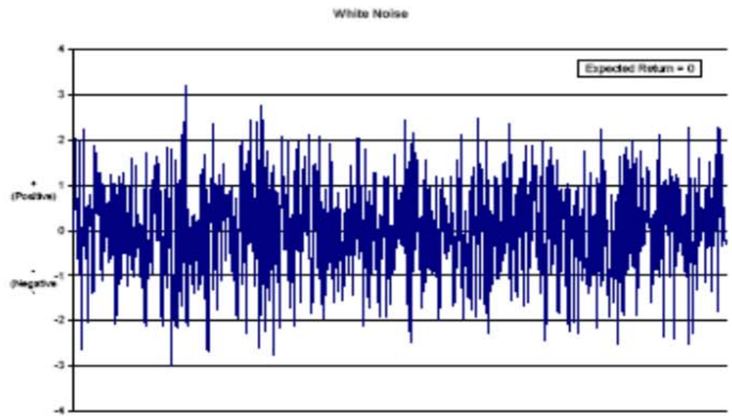

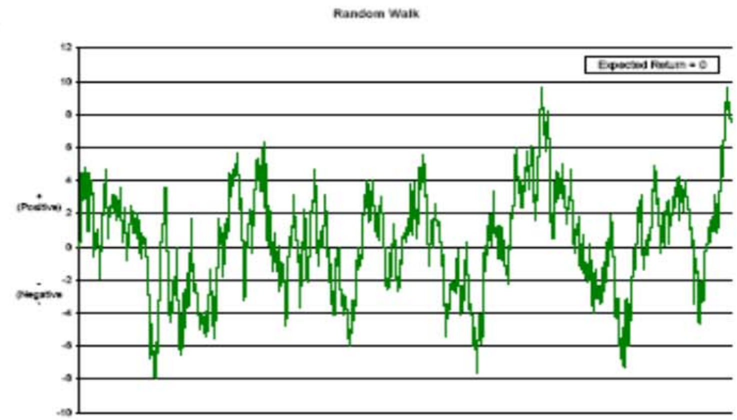

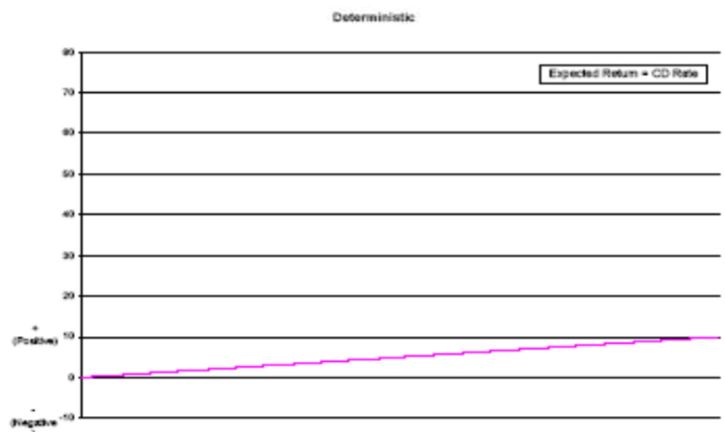

- Returns are:
    1) Riskless
    2) Predictable
- Expected returns are **not** a function of volatility.

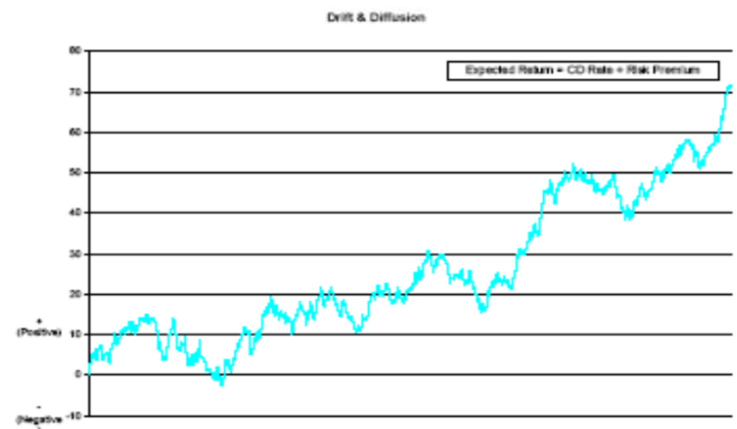

- Enhance returns through:
    3) Diversification
    4) Rebalancing
- Expected returns are a function of volatility.

Prior to deciding how much stock or bond market risk it is prudent for a trust to assume, it is critical to have a clear idea about the nature of the return-generating process. We distinguish among price-change processes in which:

1. returns are pure "noise" exhibiting no propensity to head off in any long-term direction (top left quadrant);
2. a "random walk" where returns wander far away from the long-term mean(in this case, zero) and exhibit an equal propensity to produce long periods of profits and long periods of losses with little discernable prospect for real long-term growth (top right quadrant);
3. a deterministic time series with no volatility--like a return on a long-term certificate of deposit (bottom left quadrant); or,
4. a drift-with-diffusion process that has a long-term equilibrium drift process (in this case with a positive coefficient for the mean)[14] along with a possibly mean-reversionary diffusion process.

Clearly, no one wants to invest money in a process that is only white noise. Volatility remains unrewarded and future investment growth prospects are bleak. Technically, in a white noise process, there is no signal-to-noise ratio because there is no signal—no one knows if the market has a direction.

The random walk process is interesting in that it can produce long runs of returns (just as random coin flips can produce long runs of heads or tails) that are either positive or negative. The simple zero-mean random walk model seems to cry out for a market timing strategy. However, if the simple random walk model incorporates constant probabilities such as those found in a fair coin-flip game, the trustee is as helpless at predicting the market's future direction as he is in predicting the next coin toss result.[15]

If a trustee can discharge future liabilities by investing in a deterministic price-change process, he may wish to avoid some or all market risk. Most trustees, however, are faced with a dilemma in that the deterministic (risk-free) rate of return may be far less than the return required to fund a trust's liabilities.

Finally, you may notice that the drift-with-diffusion process is, in certain respects, analogous to a combination of the deterministic time series and the random walk time series and, at high frequency intervals, a white noise process. Returns can wander up and down (sometimes with great velocity and magnitude) but have a positive central direction over the long run. If the stock return process is one with a positive drift component (drift = the long-term steady state or equilibrium value that represents long-term expected return) combined with a diffusion process that, in the short term, generates either above-equilibrium returns or below-equilibrium returns, then two results follow:

1. Long-term investors can expect to be compensated via a risk premium for tolerating the ups and downs of the market; and,

[14] The coefficient's value may be the underlying real long-term growth rate of the economy.
[15] The zero-mean random walk model incorporates Martingale mathematics in which the best prediction of future values is the current value (expected price change is zero, actual price change is zero ± white noise). The S&P 500 stock index seems to be better represented by a "Sub martingale" mathematical equation wherein the best prediction of futures values is zero ± white noise + positive long-term growth constant.

2. The long-term compensation from holding stocks will be greater than the long-term compensation from holding risk-free investments--where the long-term may be very long, indeed.

An infinitely risk averse investor with low return requirements will hold only the safe asset (cash if the planning horizon is short; or inflation-adjusted bonds if the horizon is long)[16]; and will not be attracted by the higher expected return potential of stock investments. Long-term investors with moderate to high return requirements will want to own both stocks and bonds. Although we noted earlier that serial correlation (bad returns follow bad returns / good returns follow good returns) make stocks more risky for short term investors, long-term investors focused on wealth accumulation (or on periodic wealth decumulation) still have a positive demand to hold stocks especially if stock returns are mean reverting. Mean reversion means that periods of high stock returns are likely to be followed by periods of lower returns and vice-versa. In today's economic climate, with the S&P 500 attaining the 1,960 level, a mean reversionary tendency suggests that unexpectedly high returns may be followed by low future returns.[17] Is this, as some financial advisors suggest, an historic selling opportunity?

If you subscribe to the notion that 'what goes down must come up,' then the strict-contrarian market-timing strategy of buying on the dips / selling on the highs looks pretty good.[18] However, should a trustee bet the farm on successfully timing an investment's recovery? For example, Warren Buffett placed some large investment bets under the assumption that the market was fairly valued when the Dow declined from 14,165 to approximately 10,000. Unfortunately, the Dow continued its fall to a low of 6,547.05. Unlike Mr. Buffett, many trust beneficiaries face a severe threat to their standard of living should the portfolio dip below some critical value. For example, a current beneficiary may count on the investment portfolio to sustain a threshold standard of living throughout an uncertain life span. Once the bet is made, if the investment position fails to bounce back within a reasonable time frame, the economic viability of trust may be put at risk.

If trustees base financial strategy on the old maxim: "all's well that ends well," then they may have a problem because there remains a positive probability that things may not end well (or, that the beneficiary may run out of time prior to the happy ending). Conversely, if the decision is to maximize "safety," the trustee may be hamstrung by an unproductive conservatism. The result is akin to a refusal to drive an automobile because there is a positive probability of a fatal accident. Here the decision process has the driver weighing the risks of traveling by car against the opportunity costs incurred by remaining a shut in. If a "safe" portfolio has a high probability of failing to generate the return required

[16] Taxation of trust income may alter this assertion.

[17] There is some econometric evidence that suggest that equity markets are mean reverting. The nature of the evidence is technical but the main points include:

- Shocks to returns are not persistent over the long term (volatility exhibits a rate of decay);
- There is a feedback loop from volatility to the expected equity risk premium. This feedback loop suggests a time-varying risk premium where higher than expected returns are followed by lower than expected returns;
- Conditional variance of long-term stock returns is "predictable" [not in the sense that you can formulate profitable trading strategies], and is less than unconditional variance. This results in a faster long-term convergence of stock returns towards their equilibrium values; and,
- The volatility of stock returns is itself time varying, with market volatility tending to manifest itself in "clusters" across long-term horizons. Periods of high volatility make future stock returns more uncertain with the inevitable result that stock prices decrease; when such periods are followed by low volatility, stocks prices adjust upwards as the market becomes less uncertain.

[18] Unless, of course, you are trying to market time Enron, British Petroleum, or the Czarist stock market shortly before 1917. Perhaps it is more prudent to follow the advice given by Will Rogers: Don't gamble; take all your savings and buy some good stock, and hold it till it goes up, then sell it. If it don't go up, don't buy it."

to discharge critical financial objectives, is it really a “safe” portfolio? Should the trustee reverse market time—leave the market and go to cash during high volatility periods (sell low) and, when things calm down, fully return to market to make up lost ground (buy high)? [19]

It is tempting to side-step these difficult questions by reducing the decision making process to dogma--markets always bounce back; the more severe the decline, the quicker the recovery.[20] Under a dogmatic approach, the investment process consists of one-part courage, one-part patience, one-part discipline, and one-part faith. Discipline is the new Prudence.

But the fact remains that many beneficiaries do not have the luxury of waiting, perhaps over an infinite planning horizon, for a successful investment outcome. An intolerable drop in portfolio value may result either from an accumulation of small losses or a from a single significant market drop. The magnitude of the loss may, in fact, be so large that the expectations for recovery may require an unreasonably long investment horizon. In such instances, trustees may not be able to maintain the trust’s strategic asset allocations.[21] As the trust faces the prospect of the portfolio’s value dropping below a critical threshold, the trustee can remain “disciplined,” or can recognize that a change in economic circumstances might require a change in investment exposures. This is the difference between static investment policy in which the only virtue is patience and dynamic investment policy which seeks asset management elections that help the investor survive economic damages.[22]

Disciplined investing sounds like a good thing, but, as Goldman Sachs' Abby Joseph Cohen notes: "discipline sometimes does not give the right answer. It just gives a formulaic answer and can intensify the consequences of an incorrect answer." [23] Let’s examine the potential consequences of an incorrect answer by returning to the zero-mean random walk investment model. We are now leaving the world of discrete time finance—e.g., month-to-month price changes in the S&P 500 stock index—to enter the complex domain of continuous time finance.

For any zero-mean process with the final results generated by a path-independent process (the portfolio’s final value is the same if it is achieved by a ‘down’ followed by an ‘up’ return or by an ‘up’ followed by a ‘down’ return), we may apply a theorem known as the *reflection principle*. The gist of the theorem is that, over any time interval, a realized return path has a different but equally probable return path that is its mirror image. [24] The implications of the reflection principle are profound. It suggests that each time a stock market (operating under martingale-like return evolutions) dips and then recovers it was equally likely that it could have dipped and then continued to decline along a vector that is the inverse of the recovery vector’s trajectory. Martingale mathematics provides no guarantee of a recovery—the best guess of future value is current value. Stated otherwise, it was equally likely on

---

[19] If you believe this is a prudent strategy, you must also believe that a portfolio invested 100% in cash 50% of the time and 100% in stocks the other 50% of the time is safer than a portfolio invested 50% stocks / 50% bonds for 100% of the time.

[20] Consider, for example, the following analysis: “As long-term investors, we constantly preach that riding out market downturns requires patience. We were curious as to how long it takes for that patience to be rewarded in the recovery phase following a decline. In the 12 bear markets we examined, on average, investors recovered their losses over a period of about 14 months….The median recovery time among the 12 markets was eight months. In nine instances, recovery took only a year or less.” Financial Planning Magazine (September, 2008), p. 149. Among the many defects in this argument are: (1) patience can eliminate risk, (2) a small sample statistical bias, and (3) the use of past relative frequencies prior to determining if historical conditions are sufficiently similar to the current economic environment to justify the relative frequency argument.

[21] See, for example, Elements of an Investment Policy Statement for Individual Investors (CFA Institute, 2010), §3c.

[22] Collins, Patrick J., & Stampfli, Josh, , "Managing Private Wealth: Matching Investment Policy to Client Risk Preferences," The Banking Law Journal (December, 2009), pp. 923 - 958.

[23] "Aristotle on Investment Decision Making," Financial Analysts Journal (July/August, 2005).

[24] The theorem is well known to students of algebra familiar with Pascal’s Triangle which is analogous to a recombining (i.e., path independent) binomial lattice.

March 9, 2009 that the S&P 500 index could have declined further and the probability of such a decline was exactly equal to the probability of the current recovery. Financial economists sometimes express the consequences of the reflection principle in such statements as "the market has no memory." [25]

But, as noted, the market operates under a sub-martingale-like process which is not a zero-mean random walk. However, in the somewhat strange world of continuous finance, the Doob-Meyer decomposition theorem permits a sub-martingale process to be decomposed into the sum of a strict martingale process and a continuously increasing stochastic process. Finally, the Girsanov theorem allows us to apply the reflection principle to non-zero mean stochastic processes unfolding in continuous time. Intuitively, the Girsanov theorem provides a method to understand how a price-change process evolves over time when such process is subject to a change of measure such as that which results from the Doob-Meyer decomposition. Over the long run, investors in major capital markets have the expectation of positive returns. But, as Keynes stated, "in the long run we are all dead." If the present value of wealth drops below the present value of future liabilities, financial objectives may become infeasible within the trust's time horizon.

What is the purpose of discussing these mathematical techniques? It should be evident that even the complicated math of continuous time finance cannot fully capture the complex behaviors of the S&P 500 stock index. We are tempted to dismiss these highly quantitative exercises as interesting but practically useless. However, either an unwitting or an explicit decision to restrict risk models to simple forms has a potentially catastrophic result: the trustee ceases to understand the risk/reward tradeoffs embedded in investment positions. To many investors, the world of continuous time finance is akin to quantum mechanics or string theory—highly technical approaches leading sometimes towards surprisingly unexpected conclusions. It is not yet a science that yields prescriptive advice or practical investment strategies. However, we also recognize that overly simplistic risk models are not a credible alternative on which to base an asset management strategy.

Although it seems as if trustees are surrounded by the Black Swans of discreet-time finance on one hand and by the Albino Crows of continuous-time finance on the other, there is still the attraction of the average annual return of 11.88%. However, it is clear that, like Odysseus, the trustee must avoid the destruction of the Sirens' song on the way to the financial objective. Depending on the needs, circumstances, distribution requirements, and risk-tolerance of the trust, the investment strategy may take a variety of forms including a "stay-the-course" strategy for trusts exhibiting constant relative risk aversion, or including dynamic asset management strategies for trusts with greater than average sensitivity to changes in investment wealth.

[25] This phase is most commonly used to describe negative exponential distributions where the hazard rate is not a function of elapsed time.