FACT OR FICTION: STOCKS AND BONDS ARE LESS RISKY FOR LONG-TERM INVESTORS

EVANSTON CAPITAL MANAGEMENT

MAY 2015

Peter Hecht, Ph.D.
Managing Director, Senior Investment Strategist
phecht@evanstoncap.com

EXECUTIVE SUMMARY

- Conventional wisdom suggests that long-term investors face less investment risk, thus impacting their strategic asset allocation. While there are many reasons for risk to vary with investment horizon, the typical explanations (e.g., time diversification) are misguided.

- Two commonly referenced investment environments that describe how returns evolve over time are random walk and mean reversion. Understanding both of these environments is critical to addressing questions regarding time horizon and risk.

- In a random walk environment, returns from one period are *uncorrelated* with returns from future periods. Thus, cumulative return risk scales linearly with time, and *risk per period* is the *same* across investment horizons.

- In a mean reversion environment, unexpected returns from one period are *negatively correlated* with future expected returns. Thus, cumulative return risk grows slower than the random walk case, and *risk per period declines* with investment horizon.

- The degree of mean reversion in the historical data is *large* for stocks, bonds, and balanced portfolios. A 10-year 78%/22% stocks/bonds investor experiences the same risk per period as a 1-year 60%/40% investor. However, historical measures of mean reversion require estimates of means, variances, and correlations, and we inevitably estimate them with error. Investors need to address this reality and model parameter uncertainty. Once parameter uncertainty is incorporated, the degree of mean reversion is *dramatically* reduced – the 10-year 78%/22% stocks/bonds allocation referenced before drops to 64%/36%.

- In sum, don't allow another strategic asset allocation study to be completed without a thorough discussion on investment horizon, how returns are assumed to evolve over time, and parameter uncertainty. Even if it's *unintentional*, most asset allocation work *implicitly* assumes a one-year investment horizon, a random walk, and/or no parameter uncertainty. These material assumptions do not typically reflect reality.

EVANSTON CAPITAL MANAGEMENT

EVANSTON CAPITAL MANAGEMENT, LLC
1560 SHERMAN AVENUE, SUITE 960
EVANSTON, IL 60201
P: 847-328-4961
F: 847-328-4676
www.evanstoncap.com

INTRODUCTION

At a conference I recently attended, the topic "Do you have too much equity-related risk?" came up during a risk parity[1] discussion. One of the investors said his portfolio wasn't over allocated to equity risk since he had a long-term investment horizon. He claimed that equities are less risky for long-term investors: with a long enough investment horizon, the good and bad return years cancel each other out leaving the investor with an almost guaranteed return. This logic resonated with many of the other conference attendees.

When I heard this line of reasoning, it reminded me of similar comments made over the years by other well-respected investors and asset managers:

1) **Time Diversification:** Adding uncorrelated sources of return to a portfolio lowers risk. Similarly, if returns are uncorrelated through time, a long-term investor benefits more from time diversification and, thus, equities are lower risk for long-term investors.

2) **Average Returns:** As investment horizon increases, the average (or annualized) stock return becomes less risky, justifying a larger stock allocation. This is a simple application of the law of large numbers.

3) **Sharpe Ratio:** Expected returns grow linearly with investment horizon (i.e., the two-year expected stock return equals 2 times the one-year expected return). Volatility, a measure of risk, grows with the square root of time (i.e., the two-year volatility equals the one-year volatility multiplied by the square root of 2). As a result, Sharpe Ratios (i.e., return per unit risk) increase with investment horizon, making it worthwhile for long-term investors to hold more risky assets, such as equities.

4) **Permanent Loss of Capital:** Volatility (or variance) is a poor risk measure for long-term investors who care more about permanent loss of capital. Investors with these risk preferences should hold more equities, all else equal.

At first read, all of the statements above make intuitive sense, which is probably why they continue to have traction within the investment community. While there are many reasons for risk to vary with investment horizon, the logic above is misguided or, at least, materially incomplete.

While there are many reasons for risk to vary with investment horizon, the logic above is misguided or, at least, materially incomplete.

The combination of an important investment problem (i.e., do long-term investors face less risk, and thus, should they hold more equity-related assets?), continued investor confusion, and ample academic research provides an opportunity for someone like me to attempt to fill the gap[2]. The goals of this paper are threefold. First, I will present a simple framework to think about "horizon effects" in asset risk. Topics, such as proper risk measures (e.g., volatility or variance), random walks, mean reversion, and parameter uncertainty, will be discussed. Secondly, I will apply this framework to a simple model calibrated to real data. What does the historical data suggest about "horizon effects" in stock and bond risk? Lastly, I will provide key takeaways for investors.

FOCUS ON VARIANCE, NOT VOLATILITY

Although there are many legitimate risk measures available, this discussion will focus on volatility and variance (which is just volatility squared). These risk measures are the easiest to understand and are the "work horses" of applied modern portfolio theory. As I've said in previous research, all models are false by definition. However, models, such as those focusing on volatility (and variance), can still provide valuable portfolio management insights and reasonable rules of thumb to follow[3].

Rarely quoted in discussions regarding risk, variance, not volatility, is actually the more relevant measure for determining optimal portfolio allocations. A way to understand this better is to split the allocation decision into two components: risk-adjusted returns and risk. Sharpe Ratio (expected excess return divided by volatility), a measure of risk-adjusted return, should influence an asset's *risk* allocation. Intuitively, assets with higher risk-adjusted returns should have higher risk allocations, all else equal. For example, if stocks and bonds have the same Sharpe Ratio, 50% of the risk allocation would go to stocks and 50% to bonds. If a 10% total portfolio volatility target were desired, half of it would be attributed to stocks (5%) and the other half to bonds (5%). However, one would NOT hold a 50%/50% portfolio of stocks/bonds since stocks are much more risky than bonds.

Rarely quoted in discussions regarding risk, variance, not volatility, is actually the more relevant measure for determining optimal portfolio allocations.

But how does one go from a risk allocation to an actual *portfolio* allocation? One needs to divide by a measure of risk such as volatility. Dividing a risk allocation by a risk number properly sizes the actual portfolio allocation. For a given Sharpe Ratio, it makes sense that a higher (lower) volatility should translate to a lower (higher) portfolio allocation, all else equal. In the stock/bond example above, stocks will end up with a much lower portfolio allocation given their greater volatility. Bringing it all together, Sharpe Ratio divided by volatility, or, equivalently, expected excess return divided by variance, is the relevant metric for determining optimal portfolio allocations. This is why portfolio allocations should focus on variance, not volatility[4].

$$\frac{\textit{Sharpe Ratio}}{\textit{Volatility}} = \frac{\textit{Expected Excess Return/Volatility}}{\textit{Volatility}} = \frac{\textit{Expected Excess Return}}{\textit{Volatility}^{2}} = \frac{\textit{Expected Excess Return}}{\textit{Variance}}$$

Let's go through a simple example to demonstrate the point. Assume a 20% portfolio allocation to the S&P 500 is optimal (i.e., Sharpe Ratio/Volatility = 20%). However, the S&P 500 is no longer available to trade. A new security called "2X S&P 500" is introduced instead. As the name suggests, it offers two times the return on the S&P 500[5]. How much of the new security should be held? The new security's expected return and volatility are both doubled, leaving the Sharpe Ratio (SR) unchanged. However, with the volatility (Vol) doubled, the new portfolio allocation should be cut in half to 10%.

$$\textit{New Portfolio Allocation} = \frac{\textit{New SR}}{\textit{New Vol}} = \frac{\textit{Old SR}}{2(\textit{Old Vol})} = \frac{20\%}{2} = 10\%$$

RANDOM WALK VS. MEAN REVERSION

Now that the relevant risk measure is identified, let's apply it to two common environments that describe how returns evolve over time: random walk and mean reversion. Understanding both is critical to addressing questions regarding time horizon and risk.

In order to keep things simple, the focus is on the risk characteristics of a one-year return versus a two-year cumulative return (i.e., longer horizon). We'll work with continuously compounded returns since that allows for the two-year cumulative return to equal the sum of the annual returns in year one and year two[6]. Furthermore, in order to understand the core concepts described here, it is important to break up the realized return into its two natural components: expected return ($E(Return)$) and unexpected return (ε).

$$Return_{1+2} = [E(Return_1) + \varepsilon_1] + [E(Return_2) + \varepsilon_2]$$

What's uncertain (i.e., risky) to the investor standing at the beginning of year one? Let's start with the easy answers. The Year 1 expected return is assumed to be known, e.g. the expected stock return next year is 8%. By definition, the year one and two unexpected returns are unknown and uncorrelated with each other. That's why they're called "unexpected". The year two expected return could be either known or unknown depending on the assumed environment. Now let's get back to the different return environments.

Random Walk

The random walk is one of the most popular assumptions in financial economics. Most of the first generation finance models implicitly or explicitly assumed a random walk. Even the popular press and practitioner audience fell in love with the random walk - Burton Malkiel's "A Random Walk Down Wall Street" was a national bestseller.

In order to better understand the mechanics of a random walk, consider a bingo bin containing a bunch of numbered balls. Each number represents a realized annual return. If each annual return over time gets drawn from the same bingo bin (with replacement), that's the definition of a random walk. Equivalently, returns are independently and identically distributed. Each year the return is drawn from the *same* distribution, which is represented by the bingo bin. If the distribution is the same each year, the annual expected return and variance remain constant over time. Thus, the realized return in year one has *no impact* on year two's expected return; in other words, the year 1 unexpected return is *uncorrelated* with year two's expected return.

Under the random walk assumption, the variance of the two-year cumulative return is equal to 2 times the one-year variance (i.e., variance grows linearly with investment horizon). While the two-year cumulative return variance is larger than the one-year variance[7], the variances on a "per year basis" are identical - the two-year cumulative return variance per year (i.e., divided by 2) equals the one-year variance[7]. Therefore, there are no "horizon effects" in asset risk under the random walk assumption. A two-year investor faces the *same* risk per period as a one-year investor (Figure 1).

There are no "horizon effects" in asset risk under the random walk assumption. A two-year investor faces the *same* risk per period as a one-year investor.

Figure 1

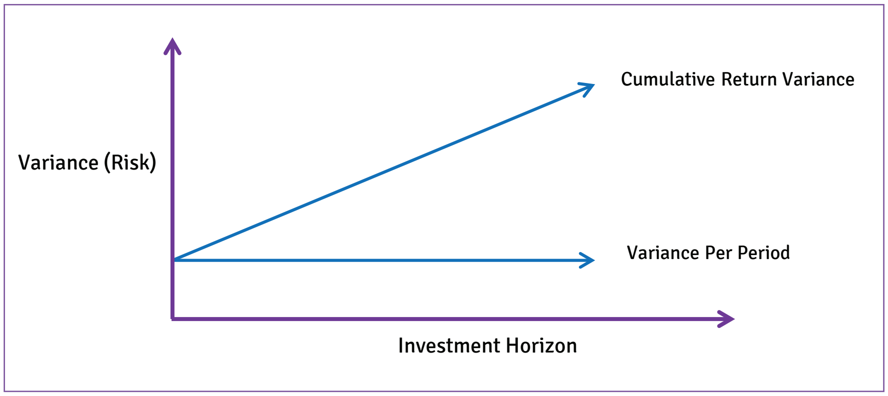

At this point, it should be clear why the logic presented at the conference (and its related variants) is misguided/incomplete. Under a random walk, returns are uncorrelated through time (so called "fallacy of time diversification"), *average* (not cumulative) returns have lower risk with longer horizons, and Sharpe Ratios increase with investment horizon. Yet, the variance of cumulative returns (the relevant risk measure) grows with investment horizon, and the variance per period is constant. The logic from the introduction does not warrant a larger equity allocation for long-term investors, all else equal. End of story. We'll address the "permanent loss of capital" issue in a moment. Hold tight.

Mean Reversion

So, how does a mean reversion environment contrast with a random walk? There are two key differences. First, standing at the beginning of year 1, the year 2 expected return is not known. Second, and most importantly, the year 2 expected return *depends* on the year 1 unexpected return. When the year 1 unexpected return is negative (positive), the year 2 expected return is higher (lower) than normal (i.e., when a small number is drawn from the bingo bin, the *average* number in the *next* bingo bin increases). In other words, the year 1 unexpected return is *negatively* correlated with the year 2 expected return. This creates a tendency for unusually good (bad) return years to be followed by unusually bad (good) return years.

Under the mean reversion assumption, there is a key new term in the two-year cumulative return variance expression: the correlation between the year 1 unexpected return and the year 2 expected return. We know from the above description that the correlation term is negative. As shown in Figure 2, the two-year cumulative return variance in a mean reversion environment is *lower* than in the random walk case. Additionally, the cumulative return variance per period is no longer constant. It decreases with time horizon, creating a rationale for long-term investors to possibly hold more of a risky asset (relative to short-term investors)[8].

...[under mean reversion] the cumulative return variance per period is no longer constant. It decreases with time horizon, creating a rationale for long-term investors to possibly hold more of a risky asset...

Figure 2

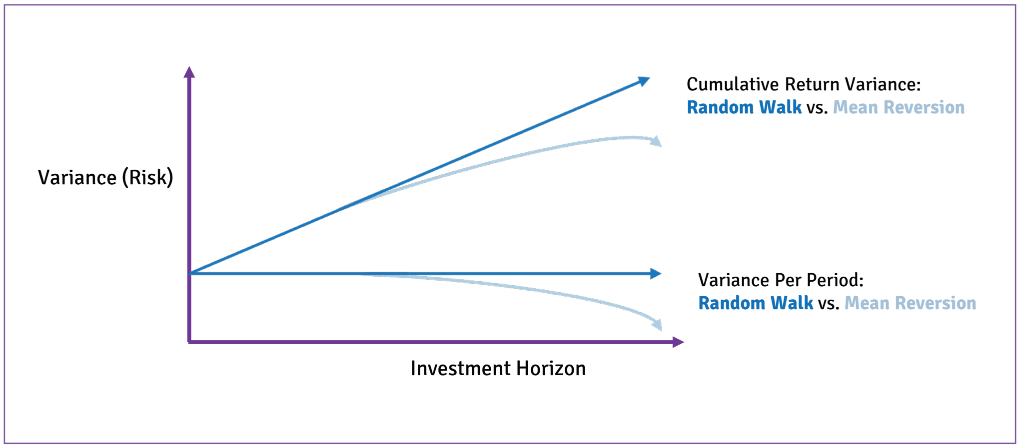

Considering how many times I've heard this from investors and asset managers, it's important to note that forecasting a lower (higher) return than the current period's unusually high (low) realized return is NOT necessarily evidence of mean reversion. It could be, but it also could be consistent with a random walk.

Let's go through an example first assuming a random walk. Say the one-year expected return is 8%. The realized return turns out to be 30% (i.e., an unusually great year). After observing the year one return, the investor forecasts a year two return that is lower than 30% (the current period's realized return). This is perfectly consistent with a random walk! Under a random walk, the year two expected return is still 8% (the same as year 1), and 8% is lower than 30%. This is NOT consistent with mean reversion. On the other hand, a statement consistent with mean reversion would be the following: after realizing a 30% return in year one, I revised my year 2 forecast (i.e., expected return) down from 8% to 3%. This demonstrates the necessary negative correlation between year one's unexpected return and year two's expected return – the defining property of mean reversion.

Permanent Loss of Capital

Before leaving this section, it's important to tie the random walk and mean reversion concepts to an often-cited risk metric, permanent loss of capital. Many long-term investors care only about permanent loss of capital and, thus, could care less about short-term, temporary price fluctuations that drive up volatility. Under a random walk, since unexpected returns are *uncorrelated* with future expected returns, all price movements are *permanent*. When a stock unexpectedly drops in price, one does NOT expect to get any part of that loss back in the future. Short-term volatility is *identical* to long-term volatility. There is no notion of a temporary price movement under a random walk. Thus, volatility, whether short-term or long-term, is a valid measure of permanent loss of capital in a random walk environment.

In contrast, mean reversion is all about temporary price movements. This is why today's unexpected return is *negatively* correlated with future expected returns. When a stock drops in price, one expects to get part of that loss back in the future through higher, forward-looking expected returns. In other words, part of the stock price drop was *temporary*. Thus, under mean reversion, *short-term* volatility overstates permanent loss of capital for long-term investors. In the end, volatility is NOT necessarily a bad risk measure for long-term investors focused on permanent loss of capital, but it is flawed when mean reversion is present AND a short-term volatility measure is used. Long-term volatility (or variance), the focus of this study, works as a proxy for permanent loss of capital under both return environments, random walk or mean reversion.

...volatility is NOT necessarily a bad risk measure for long-term investors focused on permanent loss of capital, but it is flawed when mean reversion is present AND a short-term volatility measure is used.

CALIBRATING A SIMPLE MODEL WITH REAL DATA[9]

Now that we have a conceptual framework for thinking about investment horizon and risk, the next natural step is to estimate a simple model with real data and measure the "horizon effects" (i.e., degree of mean reversion) in asset risk for stocks, bonds, and a balanced portfolio. Given the popularity of these indices, the analysis focuses on the annual returns for the S&P 500, the Barclays Aggregate, and a 60%/40% S&P 500/Barclays Aggregate portfolio[10]. Due to data availability for the Barclays Aggregate index, the time period studied is 1976 through 2014.

As discussed in the last section, the degree of horizon effects in asset risk will depend on how expected returns vary through time (if at all) and whether they are correlated with the past unexpected returns. In the spirit of simplicity, we model expected stock returns as a function of the beginning of period dividend yield (i.e., regress stock returns on lagged dividend yield) and model expected bond returns as a function of the beginning of period yield to maturity (i.e., regress bond returns on lagged yield to maturity). And for completeness, the candidate predictor variables, the dividend yield and yield to maturity, are modeled too. Both variables are assumed to follow a simple first order autoregressive process (AR1)[11]: the dividend yield (yield to maturity) is regressed on the lagged dividend yield (yield to maturity). In total, there are four regressions estimated, two for stocks and two for bonds. The results for the 60%/40% balanced portfolio can be implied from the individual stock and bond results. Descriptive statistics for the relevant data and the detailed regression results are provided in the appendix.

Qualitatively, the key regression results are the following: first, both the dividend yield and yield to maturity positively predict future annual stock and bond returns, respectively. Second, both the dividend yield and yield to maturity are highly persistent with autocorrelation coefficients above 0.9. Third, the unexpected dividend yield is negatively correlated with unexpected stock returns. Finally, the unexpected yield to maturity is negatively correlated with unexpected bond returns. These results are consistent with a mean reversion environment. Let's walk through the logic for stocks (the same logic applies to bonds):

1) Expected annual stock returns are positively related to the beginning of year one dividend yield. Assume the beginning of year one dividend yield is normal. Thus, the year one expected stock return is normal.
2) Year one plays out, and stocks do unusually well. In other words, stocks experience a positive unexpected return in year one.
3) Given the *negative* correlation between unexpected stock returns and the unexpected dividend yield, the end of year one (i.e., the beginning of year two) dividend yield is now likely to be below normal.
4) A below normal beginning of year two dividend yield translates into a below normal year two expected stock return.
5) The feedback loop is complete: unexpected stock returns in year 1 are negatively correlated with year two's expected stock return – again, the defining property of mean reversion.

With this type of risk reduction, a "balanced" ten-year investor could hold a 78%/22% stocks/bonds portfolio and achieve the same level of risk per year as a "balanced" one-year investor holding a 60%/40% stocks/bonds portfolio.

Exactly how much mean reversion is exhibited in the historical data, and what are the implications for "horizon effects" in stock and bond risk? Table 1 and Figure 3 on the next page report the cumulative return variance per period (and related statistics) implied by the regression model as a function of investment horizon. Using this simple model, the degree of mean reversion in the historical data is *large* for stocks, bonds, and balanced portfolios. A stock investor with a ten-year investment horizon experiences 32% (= 100% - 68%) less risk (cumulative return variance per period) than an investor with a one-year investment horizon. The results for bonds are even more extreme. A bond investor with a ten-year investment horizon experiences 59% (= 100% - 41%) less risk (cumulative return variance per period) than an investor with a one-year investment horizon. Given that a 60%/40% "balanced" portfolio is dominated by stock risk, a "balanced" investor with a ten-year investment horizon experiences a risk reduction close to that observed for stocks only: 35% (= 100% - 65%). With this type of risk reduction, a "balanced" ten-year investor could hold a 78%/22% stocks/bonds portfolio and achieve the same level of risk per year as a "balanced" one-year investor holding a 60%/40% stocks/bonds portfolio[12]. It is important to note that the general result (lower risk with longer investment horizons) is robust across various model specifications – the time period examined exhibited clear signs of mean reversion, and intuitively, any degree of mean reversion will reduce long-term risk.

Table 1

	Investment Horizon	Cumulative Variance	Variance Per Period	Volatility Per Period	Variance Per Period As % of Yr 1	Risk Reduction Per Period
S&P 500	1 Year	0.026	0.026	16.10%	100.00%	0.00%
	10 Years	0.177	0.018	13.31%	68.29%	31.71%
Bar Agg	1 Year	0.005	0.005	6.79%	100.00%	0.00%
	10 Years	0.019	0.002	4.34%	40.95%	59.05%
S&P/Bar Agg	1 Year	0.011	0.011	10.65%	100.00%	0.00%
	10 Years	0.074	0.007	8.58%	65.02%	34.98%

Figure 3

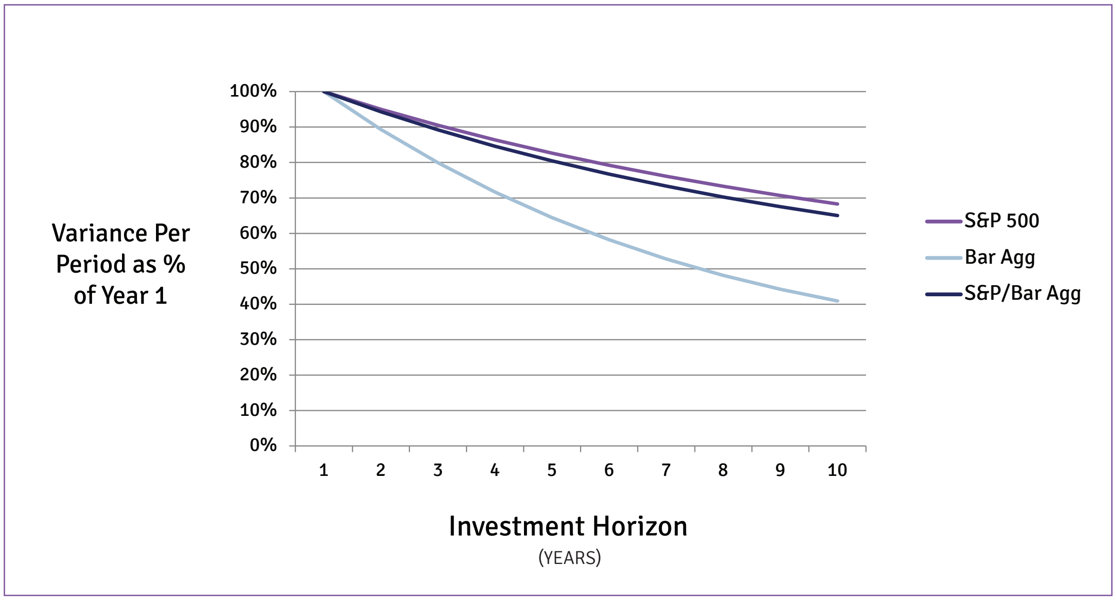

INTRODUCING PARAMETER UNCERTAINTY[13]

I have to admit it. I pulled a fast one on you during our discussion of "bingo bins" a few sections ago. While we never assumed we could predict the exact ball (i.e., return) that would be drawn from the bingo bin, we assumed we knew the exact contents of the bingo bin. In other words, we assumed we knew the *exact distribution* of outcomes and how that distribution *evolved* over time. The expected return, the return variance, the correlation between year one's return realization and the shape of year two's distribution, etc. were all assumed to be known with certainty. If only investing were that easy. Unfortunately, we don't know the shape of the return distribution and what exactly drives the evolution of that distribution over time. We need to estimate means, variances, correlations, and other parameters, and we inevitably estimate them with error.

So what do we do? Incorporate the degree of parameter uncertainty into the analysis. By doing so, one accurately reflects an *additional layer of risk* faced by real investors. When making forward-looking investment decisions, investors need to probabilistically model 1) the set of possible parameters (means, variance, correlations, etc.) and 2) the set of possible end outcomes *for each set of possible parameters.*

Since this sounds overwhelming, let's go through two random walk binomial tree examples: one with and one without parameter uncertainty. For the random walk case with known parameters (Figure 4), assume the expected return and volatility are 8% and 20% (variance = 0.2^2), respectively.

Figure 4 **2-Period Random Walk with Known Mean**

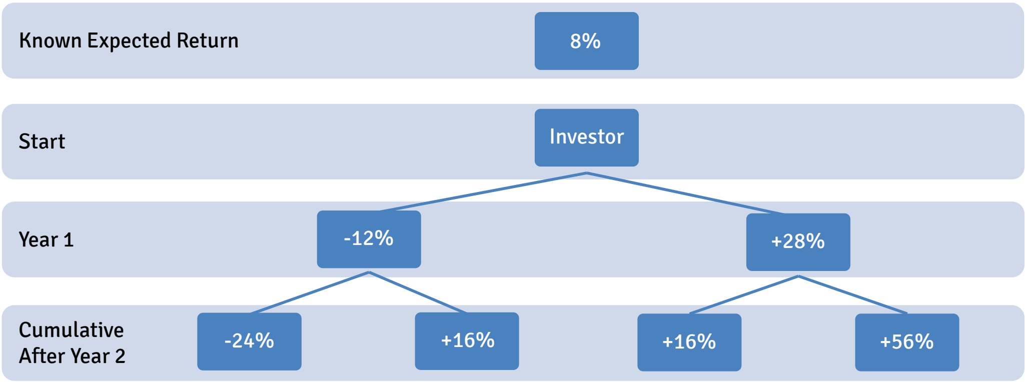

Figure 5 **2-Period Random Walk with Uncertain Mean**

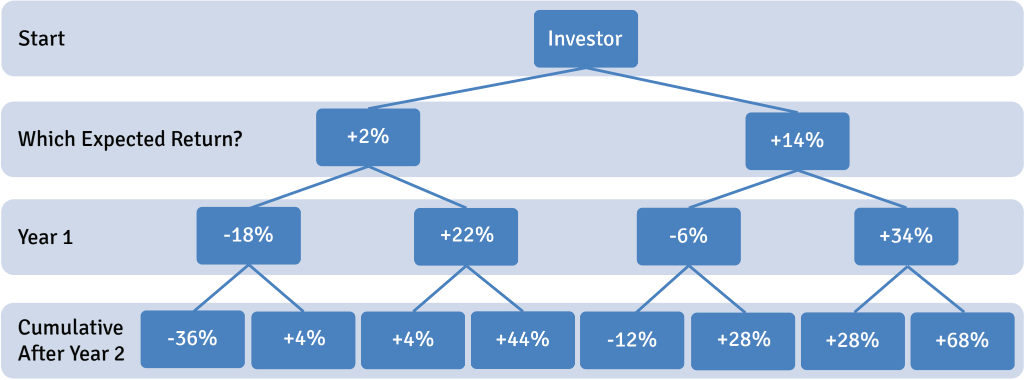

For the parameter uncertainty case (Figure 5), assume there is only uncertainty about the expected return (i.e., mean), and that there's wide dispersion in the possible *expected* returns - a 50% chance the *expected* return is 2% or 14%. Probability weighted, the expected "expected return" is 8% (= 50%*2% + 50%*14%), which is identical to the "known parameters" case. All other parameters, such as the volatility (or variance), are known with certainty and identical to the "known parameters" case. The relevant risk and return metrics are reported in Table 2.

Table 2

Cumulative Return Horizon	Avg. Cumulative Return	Avg. Return Per Period	Cumulative Return Variance	Variance Per Period
1 Yr. Known Mean	8%	8%	.040	.040
2 Yr. Known Mean	16%	8%	.080	.040
1 Yr. Uncertain Mean	8%	8%	.040	.040
2 Yr. Uncertain Mean	16%	8%	.094	.047

There are two main takeaways from the above analysis. First, when parameter uncertainty exists, the range of possible return outcomes at the end of year two exhibits more dispersion. In other words, parameter uncertainty increases risk irrespective of investment horizon. Secondly, even though we assumed a random walk world when parameter uncertainty exists, cumulative return variance per period in year two is *higher* than year one. Thus, parameter uncertainty makes assets look *more risky* to long-horizon investors. Why is this? The uncertainty in the expected return is perfectly correlated over time (i.e., if the expected return turns out to be 2%, all future expected returns will be 2% too), allowing the estimation error to accumulate quickly at longer horizons.

Thus, parameter uncertainty makes assets look *more risky* to long-horizon investors.

What does parameter uncertainty do to the empirical results from the last section? Before answering this, let me confess - I pulled another fast one on you. In the last section, we spent no time talking about the standard errors associated with all of the estimated regression parameters. We pretended all of the estimated regression parameters were known with certainty. Oops. Let's redo the analysis from the previous section, but incorporate the regression model's parameter uncertainty[14]. Table 3 and Figure 6 below report the cumulative return variance per period (and related statistics) implied by the regression model with parameter uncertainty as a function of investment horizon.

Table 3

	Investment Horizon	Cumulative Variance	Variance Per Period	Volatility Per Period	Variance Per Period As % of Yr 1	Risk Reduction Per Period
S&P 500	1 Year	0.027	0.027	16.32%	100.00%	0.00%
	10 Years	0.253	0.025	15.91%	95.10%	4.90%
Bar Agg	1 Year	0.005	0.005	6.88%	100.00%	0.00%
	10 Years	0.027	0.003	5.21%	57.50%	42.50%
S&P/Bar Agg	1 Year	0.012	0.012	10.79%	100.00%	0.00%
	10 Years	0.105	0.010	10.25%	90.20%	9.80%

Figure 6

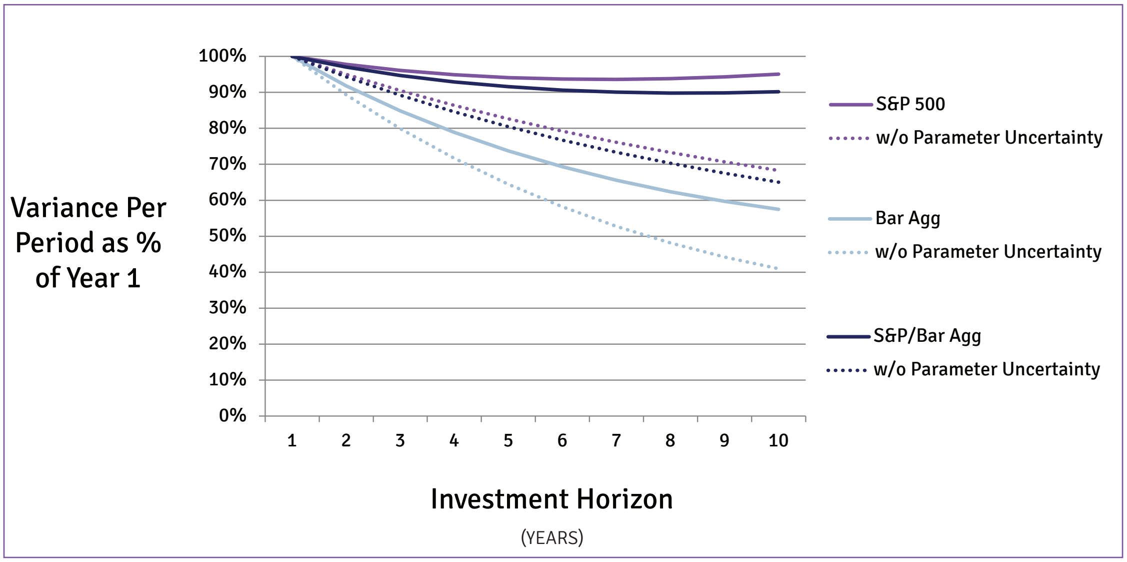

Once parameter uncertainty is incorporated, the degree of mean reversion in the historical data is *dramatically reduced*, especially for stocks and the balanced portfolio. Instead of a 32% reduction in risk, a stock investor with a ten-year investment horizon now only experiences a 5% (= 100% - 95%) reduction relative to an investor with a one-year investment horizon. The bond results change from a 59% to 42% (= 100% - 58%) risk reduction. Mean reversion for bonds is lower than before, but still meaningful. A "balanced" investor with a ten-year investment horizon continues to experience a risk reduction close to the one observed for stocks, 10% (= 100% - 90%), which is much lower than the original 35% reduction reported in the last section. With this muted type of risk reduction, a "balanced" ten-year investor could hold a 64%/36% stocks/bonds portfolio and achieve the same level of risk per year as a "balanced" one-year investor holding a 60%/40% stocks/bonds portfolio...not that big of a deal. Clearly, even in this simple model, parameter uncertainty is real and material for investors. Don't ignore it!

With this muted type of risk reduction, a "balanced" ten-year investor could hold a 64%/36% stocks/bonds portfolio and achieve the same level of risk per year as a "balanced" one-year investor holding a 60%/40% stocks/bonds portfolio...

CONCLUDING THOUGHTS

Putting it all together, are stocks and bonds less risky for long-term investors? Historically, both stocks and bonds have exhibited a material amount of mean reversion, making their risk per period lower for longer investment horizons. From a pure risk perspective, this validates the decision of long-term investors to hold a more equity-centric portfolio, all else equal. However, it's important to remember that a longer investment horizon *in itself* does not lower asset risk (i.e., the fallacy of time diversification). A longer horizon *coupled with mean-reverting returns* makes it safer for long-term investors to hold more "risky" portfolios.

On the other hand, investors need to make investment decisions today based on the *prospective* distribution of future returns, and the historical data doesn't necessarily reflect the future. In reality, the "sky is cloudy" when trying to model the degree of mean reversion in asset returns, and the process for making investment decisions needs to reflect this (i.e., incorporate parameter uncertainty). Even with a simple model of parameter uncertainty, the results can materially change.

But, as I said before, all models are false by definition…including my simple model. Nonetheless, the combination of my experience and the revealing empirical evidence on parameter uncertainty leads me to believe that stocks and bonds will exhibit some degree of mean reversion on a prospective basis, but the net impact on risk reduction will be much lower than what history and "first pass intuition" would suggest.

I leave you with a few important takeaways:

1) "Do as I say (mean reversion) and not as I do (random walk)."
Many investors claim to believe in mean reversion but consistently calculate risk statistics under the implicit assumption of a random walk environment. Have you ever reported an *annual* volatility by calculating a *monthly* volatility multiplied by the square root of 12? I see this all of the time in the industry, and it is only valid if a random walk view of the world holds. If you assume a mean reversion environment and only monthly data is available, you must model how monthly returns evolve over time to arrive at a proper annual risk statistic. When mean reversion is present, *monthly* volatility multiplied by the square root of 12 will overstate *annual* volatility.

2) Volatility, if measured properly, is a valid measure of long-horizon, permanent loss of capital. Volatility (or variance) is not a perfect risk measure, but it is unfairly criticized by long-term investors who are focused on permanent loss of capital. Using *short-term* volatility as a measure of long-horizon risk is NOT a failure of volatility. It's a failure of the investor using the wrong horizon to estimate volatility. Long-term volatility is a valid measure for *long-term* risk and, thus, a valid measure of permanent loss of capital.

3) Assessing risk at a point in time is much easier than assessing risk through time.
Be careful when doing the latter as intuition (e.g., the so called benefits of time diversification) can be dangerous. Irrespective of your risk measure of choice (variance, tail risk, probability of not achieving a minimum target return, etc.), make sure to understand how that risk measure behaves through time under the assumption of a random walk. The random walk case should be your benchmark. Why? Under a random walk, as we learned, there is nothing dynamic about how returns evolve through time. Each year, the investor faces the same prospective return distribution (i.e., bingo bin). The returns in past years have no impact on future returns. In other words, there are no "horizon effects" in risk under the assumption of a random walk. So, if you think something is less risky with a longer horizon, make sure you don't get "your result" under a random walk. Because if you do, you're probably thinking about risk and "horizon effects" incorrectly.

4) Investors/institutions "pay the bills" with cumulative realized returns, NOT average (or annualized) returns. Focus your risk discussions on cumulative realized returns. Don't fall for arguments based on averages and the law of large numbers.

5) Make sure you believe in mean reversion for the *right* reasons.
Mean reversion requires you to change your forecast from period to period as a result of an unexpectedly high or low realized return. Just saying that next periods return is forecasted to be lower (higher) than the current period's high (low) realized return is NOT necessarily mean reversion.

6) Don't just debate your return assumptions, model them.
When possible, attempt to incorporate the uncertainty around your assumptions. After a debate, there's a victor, but that doesn't mean there still isn't uncertainty about the true underlying assumptions. Assumption uncertainty is a reality, and this reality can have a material impact on the investment process.

7) Don't allow another strategic asset allocation study to be completed without a thorough, healthy discussion on investment horizon, how returns are assumed to evolve over time (random walk vs. mean reversion), and parameter uncertainty.
Even if it's *unintentional*, most asset allocation work that I've seen *implicitly* assumes a one-year investment horizon and/or a random walk. However, many investors don't have a one-year horizon and/or think a random walk view of the world is too simplistic. Additionally, in practice, I have yet to see a strategic asset allocation study that attempts to incorporate parameter uncertainty, and as discussed, parameter uncertainty can have a material impact on the results. This is a huge problem considering strategic asset allocation is probably the most important component of the investment process.

ABOUT THE AUTHOR

Mr. Hecht is a Managing Director, Senior Investment Strategist at Evanston Capital Management, LLC (“ECM”). Prior to joining ECM, Mr. Hecht served in various portfolio manager and strategy roles for Allstate Corporation’s $35 billion property & casualty insurance portfolio and $4 billion pension plan. He also had the opportunity to chair Allstate’s Investment Strategy Committee, Global Strategy Team, and Performance Measurement Authority. Mr. Hecht also served as an Assistant Professor of Finance at Harvard Business School. His research and publications cover a variety of areas within finance, including behavioral and rational theories of asset pricing, liquidity, capital market efficiency, complex security valuation, credit risk, and asset allocation. Mr. Hecht previously served at investment banks J.P. Morgan and Hambrecht & Quist, and as a consultant for State Street Global Markets. Mr. Hecht has a bachelor’s degree in Economics and Engineering Sciences from Dartmouth College and an MBA and Ph.D. in Finance from the University of Chicago’s Booth School of Business.

EVANSTON CAPITAL MANAGEMENT

Founded in 2002, Evanston Capital Management, LLC is an alternative investment firm with approximately $5.5 billion in assets under management as of May 1, 2015. ECM has extensive experience in hedge fund selection, portfolio construction, operations and risk management. The principals collectively have more than 75 years of hedge fund investing experience, and the firm has had no turnover in senior-level investment professionals since inception. ECM strives to produce superior risk-adjusted returns by constructing relatively concentrated portfolios of carefully selected and monitored hedge fund investments.

END NOTES

[1] Risk parity promotes equal risk allocations across traditional beta sources. According to proponents of risk parity, traditional balanced portfolios (e.g., 60/40 stocks/bonds) have too much equity risk. The portfolio allocation to stocks may be 60%, but the risk allocation is over 90%.

[2] A sample of articles and books on this topic are the following: "Risk and Uncertainty: A Fallacy of Large Numbers" by Samuelson (1963), "Fallacy of the Log-Normal Approximation to Portfolio Decision-Making Over Many Periods" by Merton and Samuelson (1974), "Strategic Asset Allocation: Portfolio Choice for Long-Term Investors" by Campbell and Viceira (2002), and "Are Stocks Really Less Volatile in the Long Run" by Pastor and Stambaugh (2012).

[3] My dissertation advisor and Nobel Laureate, Eugene Fama, was a vocal proponent of the "all models are false by definition" view of the world.

[4] In reality, the Sharpe Ratio and volatility need to be "beta-adjusted" when determining optimal portfolio allocations. For simplicity, these details are suppressed from the main text. Please refer to my Appraisal Ratio (https://www.evanstoncap.com/docs/news-and-research/evanston-capital-research---appraisal-ratio.pdf) paper for more detailed information.

[5] Assume a zero risk free rate for simplicity.

[6] As the return time period shrinks, it can be shown that the conditional variance of continuously compounded returns equals the conditional variance of simple returns. In other words, when looking at risk, nothing material is lost by using continuously compounded returns. However, when looking at average returns, it is important to add one half the variance to the average continuously compounded return to make it comparable to the average simple return.

[7] Mathematically,

$$Return_{1+2} = [E(Return_1) + \varepsilon_1] + [E(Return_2) + \varepsilon_2]$$

$$Var(Return_{1+2}) = Var(\varepsilon_1) + Var(\varepsilon_2) = 2Var(\varepsilon)$$

$$\frac{Var(Return_{1+2})}{2} = Var(\varepsilon) = Constant$$

[8] Mathematically,

$$Return_{1+2} = [E(Return_1) + \varepsilon_1] + [E(Return_2) + \varepsilon_2]$$

$$Var(Return_{1+2}) = 2Var(\varepsilon) + Var(E(Return_2)) + 2Cov(\varepsilon_1, E(Return_2))$$

$$\frac{Var(Return_{1+2})}{2} = Var(\varepsilon) + \frac{Var(E(Return_2))}{2} + Cov(\varepsilon_1, E(Return_2)) < Var(\varepsilon)$$

$$when \quad \frac{Var(E(Return_2))}{2} < |Cov(\varepsilon_1, E(Return_2))|$$

[9] This section is motivated by the work of Campbell and Viceira.

[10] We use the Campbell and Viceira return approximation for the balanced portfolio results.

[11] It is common to model dividend and bond yields as an AR(1). See, for example, "Stock Returns, Expected Returns, and Real Activity" by Eugene Fama (1990).

[12] The degree of mean reversion is even stronger at longer horizons.

[13] This section is motivated by the work of Pastor and Stambaugh.

[14] In order to incorporate parameter uncertainty, a Bayesian framework is used. The residuals are assumed to be multivariate normal. The residual covariance matrix is assumed to be known. All other regression parameters are assumed to have a diffuse prior. Under these assumptions, the parameter posterior distribution is well known.

[15] In order to convert a log mean into a simple mean, add ½ the log variance.

APPENDIX: REGRESSION MODEL, 1976 – 2014, ANNUAL

	Nominal Log S&P 500 D/P	Nominal Log Bar Agg Yield	Real Log S&P 500 Return	Real Log Bar Agg Return	Real Log S&P/Bar Agg Return
Mean	2.81%	6.63%	7.30%	3.87%	6.25%
Volatility	1.25%	2.92%	16.17%	7.05%	10.41%
Mean + ½ Variance[15]	2.82%	6.67%	8.60%	4.12%	6.79%

S&P 500 Model*:

$$Return_{S\&P,\ t+1} = \underset{(0.07)}{0.01} + \underset{(2.14)}{2.06}\, Div\ Yield_{S\&P,\ t} + \varepsilon_{S\&P,\ t+1}$$

$$Div\ Yield_{S\&P,\ t+1} = \underset{(0.00)}{0.00} + \underset{(0.06)}{0.93}\, Div\ Yield_{S\&P,\ t} + \varepsilon_{DivYield,\ t+1}$$

Barclays Agg Model*:

$$Return_{Agg,\ t+1} = \underset{(0.03)}{-0.01} + \underset{(0.39)}{0.70}\, Yield_{Agg,\ t} + \varepsilon_{Agg,\ t+1}$$

$$Yield_{Agg,\ t+1} = \underset{(0.00)}{0.00} + \underset{(0.07)}{0.95}\, Yield_{Agg,\ t} + \varepsilon_{Yield,\ t+1}$$

Correlation between unexpected returns and unexpected predictor variables:

$$Correlation\ (\varepsilon_{S\&P,\ t+1},\ \varepsilon_{Div\ Yield,\ t+1}) = -85.45\%$$

$$Correlation\ (\varepsilon_{Agg,\ t+1},\ \varepsilon_{Yield,\ t+1}) = -94.28\%$$

** Standard errors for regression coefficients*