EDUCATION AND EXAMINATION COMMITTEE

OF THE

SOCIETY OF ACTUARIES

INVESTMENT AND FINANCIAL MARKETS STUDY NOTE

MEASURES OF INVESTMENT RISK, MONTE CARLO SIMULATION,

AND EMPIRICAL EVIDENCE ON THE EFFICIENT MARKETS HYPOTHESIS

by

Toby A. White, FSA, CFA, PhD

Measures of Investment Risk, Monte Carlo Simulation,

and Empirical Evidence on the Efficient Markets Hypothesis

Toby A. White, FSA, CFA, PhD

Version: February 1, 2018

Contents

1   Introduction

The purpose of this study note is to supplement the corporate finance and investment material in the Berk and DeMarzo (2017) textbook. Although the great majority of the Investment and Financial Markets (IFM) Exam learning objectives from Domains I-V (i.e., the half of the exam on Corporate Finance and Investment) were addressed adequately in Berk and DeMarzo, there are a small number of topics included on Exam IFM that required more elaboration.

Below is a list of the IFM learning objectives that are addressed in this study note:

- Topic 4-Learning Outcome a): Discuss the advantages and disadvantages of different measures of investment risk.
  o Define the following measures of investment risk:
        variance, semi-variance, Value-at-Risk (VaR) and Tail Value-at-Risk (TVaR).
  o Explain the advantages and disadvantages of the risk measures listed above.
  o Calculate the risk measures listed above in order to compare investment opportunities.
- Topic 4-Learning Outcome b): Conduct risk analysis.
  o Understand the following methods to conduct risk analysis: Monte-Carlo simulation.
- Topic 3-Learning Outcome a): Explain the three forms of the efficient market hypothesis
  o Identify empirical evidence for or against each form of the efficient market hypothesis.

Note that all of Topic 4-Learning Outcome a) is covered in this study note, but only a portion of both Topic 4-Learning Outcome b) and Topic 3-Learning Outcome a) is covered here, with the remaining portion covered in the Berk and DeMarzo text. The three learning outcomes listed above are covered in Sections 2, 3, and 4 of this study note below, respectively. These three sections are largely mutually exclusive and do not have to be read in any particular order.

2   Investment Risk Measures

Risk measures can be applied in several contexts. Insurance companies are interested in how high claim liabilities might be. Banks care about profit and loss, specifically how large losses might be. Investment professionals seek to describe the probability distribution of an individual asset or portfolio of assets, and worry about return outcomes that may be especially low. In any case, one must determine an amount of assets to set aside to protect against a specific adverse outcome of the quantity of interest.

Risk measures are needed for both internal purposes, as within a corporate risk management department that informs executives of exposure pockets, or for external purposes, as when capital requirements must be satisfied for regulatory purposes. When designing risk management products or policies, the analyst is most concerned with forecasting downside risk, and thus needs to be able to quantify the probability of a particular negative outcome, or the expected loss if it surpasses some negative threshold. One can address such issues by relying on certain characteristics of the probability distribution used to describe the phenomenon of interest.

2.1   Variance

The first two risk measures discussed relate to variability. In mean-variance portfolio theory, the most popular measure of risk is variance or standard deviation, which can be employed in either an individual asset or a portfolio context. In this study note, we assume we are interested only in the distribution of a single asset instead of an entire asset portfolio.

For an individual asset, using the notation of Berk and DeMarzo, let R be a random variable that reflects the per-period return of an individual security. Then, the expected return is $E(R) = \mu$ and the variance of returns is $Var(R) = E[(R - \mu)^2] = E(R^2) - [E(R)]^2 = \sigma^2$. Correspondingly, the standard deviation is $SD(R) = \sqrt{Var(R)} = \sigma$. Naturally, the higher the value of σ^2 (or σ, since $\sigma \geq 0$), the more variable the asset's return distribution is, and the more risky the asset will be.

2.2   Semi-Variance

Because we may only care about the downside risk rather than the upside variability when designing risk management strategies, an alternative risk measure that provides such information is semi-variance, which is also known as the downside semi-variance. Now, we only look at what happens below the mean return, μ.

More specifically, let $\sigma_V^2 = Var(R) = E[(R - \mu)^2]$, and then define the semi-variance as $\sigma_{SV}^2 = E\{[\min(0, R - \mu)]^2\}$. Thus, if $R < \mu$ (which is the case we care about), then $R - \mu < 0$, and $\min(0, R - \mu) = R - \mu$. However, if $R \geq \mu$, then $R - \mu \geq 0$, and $\min(0, R - \mu) = 0$. Correspondingly, the downside standard deviation is $\sigma_{SV} = \sqrt{\sigma_{SV}^2}$, i.e., the square root of the semi-variance.

The semi-variance can be estimated by the sample semi-variance. Let n be the number of observations, and $\bar{r} = \frac{1}{n} \sum_{i=1}^{n} r_i$ where r_i is an observed outcome. Then, the sample semi-variance is $\hat{\sigma}_{SV}^2 = \frac{1}{n} \sum_{i=1}^{n} [\min(0, r_i - \bar{r})]^2$. A similar formula can be used to estimate the variance with the minimum function removed from the numerator above. Also note that the semi-variance must be non-negative, even though all of the terms in $\min(0, r_i - \bar{r})$ are non-positive due to squaring.

2.3   Value-at-Risk

Often abbreviated by VaR, Value-at-Risk is defined based upon α, where α is the given quantile or percentile. For example, if we care more about the high end of the distribution (e.g., with insurance losses), if $\alpha = 95\%$, then the VaR would be the 100αth, or 95th, percentile. However, if we care more about the low end of the distribution (e.g., poor investment returns or poor profit/loss results), if $\alpha = 5\%$, then the VaR would be the 100αth, or 5th, percentile.

Let X be a random variable that represents the investment *gain* in a security over one period so that very low (typically negative) values of X represent losses, while positive values of X represent profits. In all examples below, we assume that X is continuous, so that we find $VaR = \pi_\alpha$ as the solution to $\Pr(X \leq \pi_\alpha) = \alpha$. For example, if $\alpha = 0.05$, then we want to solve for $\pi_{0.05}$ such that $\Pr(X \leq \pi_{0.05}) = 0.05$. Other common values for α, if looking for especially adverse scenarios, might be 0.025, 0.010, or 0.005. If $\alpha = 0.005$, only 1 out of 200 times will the observed value of X fall below the VaR, assuming the given distribution is correct.

Note that if high values are more detrimental than low values, as with insurance losses, then $VaR = \pi_\alpha$ will still be the solution to $\Pr(X \leq \pi_\alpha) = \alpha$, except that now the α values will be closer to 1 (like 0.950, 0.975, 0.990. or 0995). In this case, Value-at-Risk could be conceived as the amount of capital required, with a high probability of certainty, to protect against insolvency.

To determine quantiles for probability distributions, it may be enough to rely on the cumulative distribution function, F, if it is in closed-form. Then we have

$$\Pr(X \leq \pi_\alpha) = \alpha$$
$$F(\pi_\alpha) = \alpha \quad \text{(definition of cumulative distribution function)}$$
$$F^{-1}[F(\pi_\alpha)] = F^{-1}(\alpha) \quad \text{(applying the inverse function)}$$
$$\pi_\alpha = F^{-1}(\alpha).$$

Thus, we need to apply the inverse cumulative distribution function, which may be available from tables or from software. Note that if the variable of interest is normally distributed, then after converting the variable to a standard normal (i.e., $Z \sim N(0,1)$) by subtracting its mean and dividing by its standard deviation, normal tables can be applied. However, very few distributions in finance or insurance behave in a "normal" manner.

2.4 Tail Value-at-Risk

A primary weakness of the VaR measure is that if the adverse outcome is worse than the assumed quantile cutoff, there is no way to measure how bad the outcome actually is. We initially assume that adverse outcomes occur at the lower end of the distribution, as with investment gains or profitability. Here, if $a = 0.05$, and the 1-in-20 worst-case scenario occurs, we do not necessarily know how bad things are. Returning to our X variable, for example, are losses just barely below the 5th quantile, or are they closer to the 0.01th quantile (where losses are near the theoretical maximum of 100%)?

To address this obstacle, Tail-Value-at-Risk, or TVaR, was created to focus on what happens in the adverse tail of the probability distribution. More specifically, TVaR, also known as Conditional Tail Expectation (CTE) or Expected Shortfall, is the expected loss *given* that the gain/loss is in the lowest 100α% of the probability distribution.

Then, for a given confidence level α, where α is closer to 0 than to 1,

$$TVaR_\alpha = E(X \mid X \le \pi_\alpha) = \frac{1}{\alpha} \int_{-\infty}^{\pi_\alpha} xf(x)dx$$

where $\alpha = \Pr(X \le \pi_\alpha)$. If the lower bound for losses is limited (e.g., to the initial investment amount), replace $-\infty$ above with the theoretical minimum. Note that if instead, adverse outcomes occur at the upper end of the distribution, as with insurance claims, our revised definition of TVaR, for a given confidence level α, where α is closer to 1 than to 0, is

$$TVaR_\alpha = E(X \mid X > \pi_\alpha) = \frac{1}{1-\alpha} \int_{\pi_\alpha}^{\infty} xf(x)dx .$$

Note that it may be computationally challenging to calculate TVaR, and in fact, in some cases, it may even be difficult to calculate VaR. In such cases, Monte Carlo simulation techniques (see Section 3 below) may be of use. One can also estimate these quantities based on the data if the sample size is large enough to supply a reliable estimate. For example, if there are 1,000 gain/loss observations, the VaR could be estimated by ordering the observations (as order statistics) $X_{(1)} \le X_{(2)} \le \cdots \le X_{(999)} \le X_{(1000)}$, and then choosing either $X_{(50)}$, $X_{(51)}$, or some weighted-average of the two as a proxy for the 5th percentile. Correspondingly, the TVaR could be estimated as $(X_{(1)} + X_{(2)} + \cdots + X_{(50)})/50$.

When comparing TVaR to VaR, observe that TVaR is more conservative as it provides a greater buffer against adverse outcomes. Put differently, the expected value of X, the gain/loss random variable, given that it is below the VaR, must be smaller (i.e., representing a larger loss) than the VaR itself. Thus, while the VaR only determines where the "tail" of a distribution begins, the TVaR quantifies the tail behavior.

2.5   Coherent Risk Measures

Now that we have identified four distinct risk measures, it is useful to develop criteria that seek to compare which ones perform the best. There are four desirable characteristics for a risk measure. If all are present, the risk measure is labelled "coherent." Let X and Y be two random variables, each representing a loss, with risk measures $g(X)$ and $g(Y)$,[1] these criteria are:

(1) Translation Invariance: for any positive constant c, $g(X + c) = g(X) + c$;

(2) Positive Homogeneity: for any positive constant c, $g(cX) = cg(X)$;

(3) Subadditivity: $g(X + Y) \le g(X) + g(Y)$; and

(4) Monotonicity: if $X \le Y$ for all possible outcomes, then $g(X) \le g(Y)$.

---

[1] Note that any risk measure can be expressed as a function of a random variable that maps the random variable to a number that depends on the random variable's distribution.

Translation invariance implies that adding a constant amount to a risk adds an equivalent amount to the risk measure. Positive homogeneity means that changing the units of loss will only change the amount of the risk measure in a proportional manner. Subaddivity suggests that it is not possible to reduce the capital required to manage a risk by splitting it into separate parts, or equivalently, there is some diversification benefit from combining risks as long as the two risks are not perfectly correlated. Monotonicity just states that if one risk always exceeds another, then the corresponding risk measures must be similarly ordered.

It can be shown that VaR is not subadditive and therefore is not coherent. Because of this, the TVaR risk measure, which is coherent, has gained in popularity in recent decades. However, many still prefer VaR since it is easier to calculate and understand, and in many circumstances, it satisfies the subadditivity property anyway. For example, if loss distributions are assumed to be normal, the VaR measure can be shown to be coherent.

3   Monte Carlo Simulation

For a model with several variables, recall from Berk and DeMarzo, Section 8-5, that sensitivity analysis measures how the quantity of interest (e.g., NPV for capital budgeting) changes when one model variable is changed, keeping all other variables at their base level. With break-even analysis, we determine the level of each model variable that makes the NPV = 0, keeping all other variables at their base level. With scenario analysis, we explore how NPV changes when various desired subsets of the complete set of model variables are changed simultaneously.

Monte Carlo simulation is the most thorough method for project analysis because it allows us to consider all possible combinations of model variables in an aggregate framework. Now, we can inspect the entire distribution of project outcomes, which in theory, provides the most broad-based, realistic description of what might happen. Note that there is a tradeoff between obtaining the best depiction of reality and the necessary complexity of the method, but with recent advances in statistical computing, this is no longer a great concern.

The more complex an insurance product, or an investment portfolio, the more beneficial simulation methods will be. For example, if the insurance product's cash flows depend either on interest rates or the stock market, or if the investment portfolio contains a combination of stocks, bonds, and options (see McDonald (2013)), Monte Carlo simulation may prove especially useful in forecasting the entire probability distribution of profitability (or gain/loss).

3.1   The General Methodological Approach

The simulation process generally consists of the following five steps:

> (1) Build the model of interest, M, which is a function of several variables $(X_1, X_2, \ldots, X_v)$, where v is the number of variables in the model. Note that M is also a random variable and measures the quantity of interest overall (e.g., NPV or profitability). Assume that the joint probability distribution for the set of variables is known and specified.

(2) Generate n random values of the vector of v variables. Thus, for y_1, generate the vector $(x_{11}, x_{12}, \ldots, x_{1v})$ from the joint distribution. For y_2, generate $(x_{21}, x_{22}, \ldots, x_{2v})$, and so on. Note that there will be nv distinct random values generated. The easiest way to generate a single random value (for an individual model variable) is to make a random draw, u, from the uniform distribution from 0 to 1, and then use the result of this draw to determine the corresponding percentile from the probability distribution of the model variable. For example, consider simulating one observation from an exponential distribution with mean 10. Suppose you draw 0.63 for u. The simulated value is then the 63rd percentile from that distribution, which is the solution to $0.63 = F(x) = 1 - e^{-x/10}$, which is $x = -10\ln(1 - 0.63) = 9.94$. This is called the inversion method, because the simulated value is found by inverting the distribution function. Note that if the v values are dependent, after simulating the first value, the second one must be simulated from the inverse of the conditional distribution given the first value.

(3) For $j = 1, 2, \ldots, n$, determine the value of the quantity of interest, m_j. Thus, m_j is the aggregate model value based on $(x_{j1}, x_{j2}, \ldots, x_{jv})$. In the end, there are n simulated values for m. Typically, n is a large enough number (1,000, 10,000, or 100,000 are common choices) to form a good approximation of the distribution of the aggregate model.

(4) Using $m_1, m_2, \ldots, m_n$, approximate the cumulative distribution function for M, $F(m)$, by $F_n(m)$, the empirical cumulative distribution function. That is,

$$F_n(m) = \frac{\text{number of } m_1, m_2, \ldots, m_n \text{ that are} \leq m}{n}.$$

(5) Use the empirical cumulative distribution function or the raw values $m_1, m_2, \ldots, m_n$ to compute quantities of interest. Examples include the sample mean to estimate $E(M)$, or the sample variance to estimate $\text{Var}(M)$. Risk measures such as the semi-variance, VaR, and TVaR can be similarly estimated.

3.2   Examples of Monte Carlo Simulation
We now consider three examples of how Monte Carlo simulation may be implemented. Example #1 is introductory, and consists of spinning a wheel. Example #2 applies capital budgeting techniques to estimate a project's NPV. Example #3 involves an aggregate loss distribution for an insurance company that incorporates both claim frequency and severity.

3.2.1   Example #1: Playing a Game of Chance

Consider a wheel, where it costs 1 to spin, that has 3 regions of varying size and outcome:

- 60% of the time, it lands on 0, which creates a loss of 1.
- 30% of the time, it lands on 1, which creates no gain (i.e., the player gets the entry fee back).
- 10% of the time, it lands on 6, which creates a gain of 5.

You plan to spin the wheel exactly 50 times, and care about the distribution of your total gain or loss at the end of your 50 spins. Although this example is simple enough to be fleshed out using basic probability theory, we will apply the techniques of Monte Carlo simulation for comparative purposes.

For each spin, use a random number generator (e.g., the RAND function in Excel or the runif function in R) to select a number u between 0 and 1.[2] If $0 \leq u < 0.6$, the spin's outcome is 0; if $0.6 \leq u < 0.9$, the spin's outcome is 1; if $0.9 \leq u < 1$, the spin's outcome is 6. It should be clear that if the generator is truly simulating from the uniform distributions, the three outcomes will occur with the desired probabilities. Note that this algorithm can be implemented in Excel using IF statements or by writing a simple program in R. After obtaining all 50 random draws, observe the outcomes of the 50 spins, convert the outcomes to gains or losses by subtracting the spin fee of 1, and calculate the sum of the 50 gain/loss results. Let's call this G, where G is a single simulated value of total gain/loss from spinning the wheel 50 times.

This entire process (of 50 spins) must be repeated now a large number of times to adequately approximate the probability distribution of G. Thus, if we were to do all the steps in the above paragraph 100 times (a total of 5,000 spins), we would have 100 simulated values for G, and thus a crude approximation of its distribution. For example, in one application of this in Excel, the mean of the 100 simulated G values was –4.7, compared to the theoretical mean of –5 = 50×[(0.6)(–1)+(0.3)(0)+(0.1)(5)]. The standard deviation of the 100 simulated G values was 13.43, compared to the theoretical standard deviation of 12.43. 37% of the simulated G values were positive, creating a total gain across 50 spins, compared to the theoretical rate of 34.37%. If we did 1,000 or 10,000 trials of 50 spins each, there would likely be even smaller discrepancies between the simulated estimates above and the theoretically exact quantities.

3.2.2   Example #2: Determining a Project's NPV

Now, assume a local cupcake manufacturer is considering a franchise expansion to other locations outside the region. They plan to do a capital budgeting forecast to estimate this project's net present value (NPV) so they can decide whether or not to implement the expansion. They estimate the base-case, best-case, and worst-case levels for six different variables that affect NPV. The three levels for each variable are given in the table below:

---

[2] We assume the generator can produce a value of 0, but not a value of 1. Generators are not consistent in this regard.

Variable Name	Worst-Case Level	Base-Case Level	Best-Case Level
Initial Investment (at = 0) to Expand	5,500,000	5,000,000	4,500,000
Pre-Tax Cash Flow at End-of-Year 1	900,000	1,000,000	1,100,000
Growth Rate in Cash Flows (after Year 1)	2%	3%	4%
Tax Rate	35%	25%	15%
Cost of Capital	7%	6%	5%
Project Lifetime	8 years	10 years	12 years

Assume that the project has a finite time horizon and no shutdown costs (so that when the project lifetime expires, all cash flows cease). Also, assume that the NPV is calculated as the PV(after-tax cash flows) less the initial investment. Finally, assume that the six variables are mutually independent. One can easily verify that the NPV, using all of the base-case levels for each of the six variables, is 1,239,103, which would indicate the project should be approved (since NPV > 0).

Using Monte Carlo simulation, though, we can approximate the entire probability distribution of estimated NPV values. First, we have to assign probabilities to each of the three cases (for each of the six variables). To keep things simple, let's assume these probabilities are the same across all six variables; more specifically, the probability the worst case occurs is 25%, the probability the base case occurs is 50%, and the probability the best case occurs is 25%. We can use the same technique as in the spinner example. First, generate 6 random numbers between 0 and 1, one for each variable, and then based on those draws, read in the appropriate value from the table. More specifically, if $0 \leq u < 0.25$, we pull from the worst-case column, if $0.25 \leq u < 0.75$, use the base-case column, and if $0.75 \leq u < 1$, use the best-case column. As with the spinner, there is a finite number of variable combinations possible (since all six have a discrete distribution); in particular, here we have 3^6, or 729, possible combinations. A deterministic approach to estimating NPV overall would be to take a probability-weighted average of the NPVs that result for each of the 729 possible combinations. It would be useful to compare this result to the estimated NPV from the stochastic approach detailed below.

The generation of the 6 random numbers, the conversion to appropriate cells in the table, and the calculation of after-tax NPV, comprises a single trial. The next step is to repeat this process a large number of times (say 10,000), which will form a simulated probability distribution for NPV. We can then answer questions like:

- What is the expected NPV? What is the variability of our estimate of the expected NPV?
- What is the probability that NPV will be positive?
- What is the probability that NPV will be above 1,000,000? Below 1,000,000?
- What is the 90th percentile of the estimated NPV? 10th percentile? 1st percentile?

3.2.3   Example #3: Aggregate Loss Distribution

The two examples above use only discrete distributions for all model variables. Here, we incorporate a compound distribution, which consists of a discrete distribution for the frequency or number of claims and a continuous distribution for the severity or resulting loss amount from each claim.

Consider a large family that has an auto insurance policy that covers five different cars.

- In a particular calendar year, assume that the number of auto claims incurred by this family, N, is Poisson distributed with mean 1.2.
- Also, assume that conditional on $N = n$, claim amounts, given by $X_1, X_2, \ldots, X_n$, are independent and have the exponential distribution with mean 3,000.
- Assume there is no deductible or upper limit, and that there is no dependence structure between the value of N and any of the X values.

We are interested in knowing about the distribution of $S = X_1 + X_2 + \cdots + X_N$, that is, the aggregate distribution of total losses in the calendar year by the entire family. Here, it can be computationally difficult to derive the distribution of S using introductory probability facts.[3]

The first step in a simulation is to obtain a value from the frequency distribution. Begin once again by drawing a random number from a uniform (0,1) distribution, and matching that draw to the Poisson probabilities, which are known. In this case, the probability of having 0 claims is $e^{-1.2} = 0.3012$, of 1 claim is $1.2e^{-1.2} = 0.3614$, of 2 claims is $,1.2^2 e^{-1.2} / 2 = 0.2169$ and so on. Then we use the uniform intervals [0, 0.3012), [0.3012, 0.6626), [0.6626, 0.8795), to simulate 0, 1, and 2 claims respectively. Suppose the uniform draw is 0.7327. In that case, we have simulated an outcome of two claims.

Continuing the example, we now must make two additional random draws from the severity distribution. Using the inversion method, for a given value of u, we must solve $u = F(x) = 1 - e^{-x/3000}$ for $x = -3000 \ln(1-u)$. For this example, suppose the two uniform draws were 0.2301 and 0.8752. This leads to the simulated claims $-3000 \ln(1-0.2301) = 784.48$ and $-3000 \ln(1-0.8752) = 6243.13$. Then, our simulated aggregate loss is S = 784.48 + 6243.13 = 7027.48. We then repeat this entire process a large number of times, as before, to develop an entire, simulated probability distribution for S.

4   Market Efficiency

In the discussion of market efficiency below, we focus on equities and the stock market. An efficient market is one in which stock prices fully reflect information. As stated in Berk and

---

[3] Other approaches to obtaining the distribution of *S* are covered elsewhere in actuarial examinations. The focus here is on providing an illustration of the simulation approach.

DeMarzo, there are three levels of market efficiency, distinguished by the type of information reflected in stock prices:

(1) Weak-Form: Prices reflect information contained in the record of past prices. Thus, it would be impossible to make consistent profits from analyzing past returns.

(2) Semi-Strong Form: Prices reflect not just past prices but also all other publicly available information, such as that found in easily accessible financial statements or in the financial press. Then, prices will adjust immediately upon the release of any public announcements, such as earnings reports, new issues, or mergers. Subsequently, there would be no further opportunities to profit from this information.

(3) Strong Form: Prices reflect not just past prices and publicly available information, but also all information obtainable from a painstaking analysis of the company, industry, and economy, or any information acquired from private sources. Then, there are only lucky and unlucky investors, but none who can consistently beat the market based on skill.

There is a spectrum of beliefs in the investment profession about the efficient markets hypothesis (EMH). One extreme is to believe in the strong form (and thus, all three forms), so that any active investment strategy would be futile. The other extreme is to believe in none of the three forms (and thus, not even the weak form) so that any of the previously discussed forms of information can be used to produce a winning strategy over the long run. In between these two extremes are those who believe in some but not all of the three forms of the efficient markets hypothesis.

In the remainder of this section, we consider two opposing viewpoints. First, we make the case that markets are efficient and provide evidence in support of the EMH. Then, we cite several market anomalies that suggest that perhaps markets are not always efficient, and that there may be systematic opportunities to beat the market over time.

4.1   Empirical Evidence supporting the Efficient Markets Hypothesis

There have been several studies that support the conclusions for each of the three levels of the EMH. We consider each of the aforementioned levels in turn, beginning with the weak form.

4.1.1   Weak Form

Kendall (1953) set out to discover regular cycles in stock prices, but upon examining the data, he found that prices behaved more like a random walk. This means that price series wandered in no particular pattern, as if someone simply picked a random number (which could be positive or negative) to add to the current stock price to arrive at the price at the end of the next period. Successive changes in the value of the time series are independent, so that whatever price patterns have occurred in the past do not affect what might occur in the future.

A more formal way to test for dependence between adjacent entries in a time series is to calculate the autocorrelation coefficient. In Brealey, Meyers, and Allen (2017), they examine the price changes for each of four major stocks on successive days over a 23-year period from

1991-2014. They plot the return on day t on the horizontal axis versus the return on day $t+1$ on the vertical axis. If there was a tendency for increases (i.e., positive returns) to be followed by additional increases, one would expect to see more points plotted in the first quadrant (NE region) than in the fourth quadrant (SE region). However, there was no evidence of such a momentum effect. Similarly, if increases were more typically followed by reversals (i.e. decreases, or negative returns), then one would expect more points in the fourth quadrant than in the first quadrant. Furthermore, there was also no evidence of a mean-reverting effect; that is, conclusions were similar when beginning with *decreases* on day t (and then comparing the second and third quadrants).

In fact, the scatterplots for all four stocks reflected a high density of points around the origin with no discernible pattern or ordering among any of the four quadrants. Furthermore, when the autocorrelation coefficients were calculated, they were all very close to 0, suggesting little-to-no relationship between price returns on successive days. Upon reflection, this is not altogether a surprising result, since if past price changes could systematically be used to predict future price changes in a competitive market, it would be too easy for investors to make profits, and everyone would follow similar trading rules.

Poterba and Summers (1988) took the notion of random walks one step further by examining the proportionate variability of returns. More specifically, if successive returns are completely random, then the variance of these returns should increase in direct proportion to the interval over which the returns are measured. For example, the variance of annual returns should be 12 times the variance of monthly returns. They found that this theoretical guideline worked well as an approximation, although there were slight deviations in both directions. When comparing short time periods of differing length, they found the variance of the longer period to be slightly larger than the rule would suggest, indicating some short-term momentum in stock prices. Correspondingly, when comparing long time periods of differing length, they found the variance of the longer period to be slightly smaller than the rule would suggest, indicating a tendency of price changes to reverse.

Still, the majority of these findings are consistent with the weak form of the EMH in that the current stock price already reflects all information obtainable from the record of past prices. Furthermore, this past record cannot be used to further predict future prices.

4.1.2 Semi-Strong Form

If the semi-strong version of the EMH were true, we should expect no discernible pattern in stock price data *after* the market initially reacts to the public release of any new information about the company that issues the stock. Ideally, this information is released simultaneously to the investing community so that all parties have the same chance to react to the new information. However, when examining what happens to prices *prior* to the release of such information, there may be actual differences in levels of information access or leakage.

Researchers have examined what happens to stock prices after various news releases, the most common of which are earnings or dividend announcements, merger or takeover intentions, or the release of macroeconomic conditions. It is not typically the absolute numbers contained in the news that drive price changes, but rather the updated forecasts relative to previous expectations. To isolate the effect that a news announcement will have on a stock's price, one needs to calculate both the adjusted stock return and the abnormal stock return. The adjusted stock return controls for market movements, since we only care about the portion of the stock return that is not just due to market returns. Thus,

Adjusted stock return = actual stock return – actual market return (using some index like the S&P 500 as a market proxy)

The abnormal stock return also incorporates a stock's estimated α and β coefficients, and thus, reflects firm-specific news only. Here, α is the average amount a stock price changes when the market index remains unchanged, and β is the % change in the stock price for a 1% change in the market index. Thus,

Abnormal stock return = actual stock return – expected stock return, where

Expected stock return = α + β ×(actual market return).

For example, Keown and Pinkerton (1981) examined a sample of 17,000 firms that were targets of takeover attempts. Because acquiring firms typically must pay a "takeover premium," or above-fair-market-value-price, to win the deal, the target firm's stock price increases once the takeover bid is publicly announced. They explored the average abnormal stock returns before, at, and after the announcement date. In the 3 months before the announcement date, there was a gradual increase in price levels as investors sensed that a takeover was coming. At the exact time of the announcement, there was a one-time, instantaneous increase in prices as expected. However, after the announcement, the price-increase trend stopped immediately and there was no significant further positive drift. This result is consistent with the semi-strong version of the EMH; that is, once the announcement occurs, and prices react nearly instantaneously, it is not possible for the tardy investor to then buy in and make further profits.

4.1.3 Strong Form

When testing for the strong form of the EMH, we now look not just at past prices or publicly available information, but at any remaining information that only the experts may be able to gather (or infer) after doing an exhaustive review of the financial state of the company, its industry, and the economy at large. If this sort of information could provide a competitive advantage relative to one's peers who do not have, or cannot exploit, this news as efficiently, then one would expect the superior analyst, or one that exploits inside information, to realize superior returns more often than not, relative to the market (or their peer group).

However, studies have repeatedly shown that superior returns do not necessarily result from this type of expert analyses. For example, consider the universe of actively managed mutual

funds from 1971-2013 (Brealey, Meyers, Allen (2017)). In some years, the diversified equity-based mutual funds beat the market, as measured by the Wilshire 5000 index, but about 60% of the time they underperformed the market. Of course, many mutual funds are more specialized and focus only on a single sector or industry, but even then, they are observed to either underperform the respective benchmark index, or do just as well at best.

When fees and expenses are factored into the analysis, the results are even less encouraging for active investment strategies; that is, without any fees, the funds may roughly match their benchmarks, but with fees netted out of gross returns, the funds clearly underperform their reference indices. When studying the performance of top ranking active fund managers in one year, they only have a 50% chance to beat the market in the following year (Malkiel, 1995).

These findings are relatively well known in the investment community, and are consistent with the strong form of the EMH. As a result, there has been a strong movement out of actively managed funds recently and into more passive strategies. For example, investors are increasingly just buying index funds that track the market (or a particular segment of it), which not only alleviates worries about underperformance, but is also much cheaper in terms of transaction costs and service fees.

4.2 Empirical Evidence against the EMH: Market Anomalies

Despite the large body of evidence that supports the notion that markets are efficient, many still believe there are opportunities available to achieve superior differential returns. In fact, in the academic community, the prevailing opinion has shifted over time. Formerly, it was assumed that the EMH was an accurate description of reality, but in recent decades, there have been a plethora of market anomalies discovered that refute this notion. A market anomaly is a violation of market efficiency whereby, if true, a relatively simple trading strategy that relies either on the record of past prices, publicly available information, or certain market factors, can generate superior returns.

The number of market anomalies that have been discovered and published is too large to include here. Thus, we focus only on a representative sample of anomalies. It is necessary, though, to mention a few caveats. First, although these anomalies were found to have worked using large data sets from the past, there is no guarantee that they will continue to work in the future; in fact, it is often the case that they disappear over time. Second, for many of the anomalies known to exist today, one may need access to computing algorithms that can perform high-frequency trading to exploit them; if too slow to react to an arbitrage opportunity, your technologically advantaged competitors may have reaped all the benefits available before you have taken any action. Third, skeptics might say that anomalies are just artifacts of data mining; that is, if you can dig deeply enough into a large data set, you are bound to find some interesting relationships, but it is possible that such conclusions could be spurious, and not replicable in general.

4.2.1   Calendar/Time Anomalies

There may be certain times of the year, week, or even day, where investors have more of an advantage than at other times. Here are three such calendar, or time-based, anomalies:

(1) January effect: returns were observed to be higher in January (and lower in December) than in other months. The theory is that investors sell off a portion of their portfolio for both tax reasons and spending needs at end-of-year, pushing prices down, and then buy back in at the beginning of the year when prices are at a discount.

(2) Monday effect: returns were observed to be lower on Monday (and higher on Friday) than in other days. There is not a legitimate explanation for this, and in fact, the finding is counterintuitive, as one would expect Monday returns (which encompass the full three-day weekend) to be higher than Friday returns (which reflect only a single day).

(3) Time-of-Day effect: returns are observed to be more volatile at times close to both market open (9:30 a.m. E.S.T. in the U.S.) and market close (4:00 p.m. E.S.T.). In other words, the majority of returns in a particular day occur close to the opening and closing bell, along with relatively higher trading volumes around these times.

4.2.2   When Investors either Underreact or Overreact to new information

Here, we revisit the semi-strong form of the EMH, but cite several phenomena where returns have drifted even further after the announcement date, which implies that investors either underreacted or overreacted to the news initially, and needed additional time to process it. We offer three of the most commonly cited anomalies, or puzzles, of this sort:

(1) New-Issue Puzzle: When firms issue stock to the public for the first time, as in an initial public offering, there is often a mad scramble to subscribe to the new issue, since there can be significant capital gains in the first few days (or hours/minutes) the stock is traded. However, five years later, the returns on such stocks have lagged comparable benchmarks by 3-4%. Here, participating investors overreacted to the initial news.

(2) Earnings Announcement Puzzle: Between 1972 and 2001, the 10% of all stocks of firms with the most unexpectedly good earnings outperformed the 10% with the most unexpectedly bad earnings by about 1% per month over the 6 months following the earnings announcement. Here, investors underreacted to the initial news.

(3) Momentum Effect v. Reversal Effect: Many studies have cited evidence that disproves the random walk model and thus violates the weak form of the EMH. These studies are split, though, on the direction of the effect (i.e., whether there is a momentum effect or a reversal effect). The momentum effect suggests that there is positive serial correlation in stock prices, so that investors underreact to any new information. The reversal effect, however, suggests the opposite, whereby there is negative serial correlation, and that stock prices follow a mean-reverting process, so that investors overreact to any new information.

4.2.3 Other (Non-Intuitive) Anomalies

Briefly, here are some other prominent anomalies that are potentially harder to explain:

(1) Siamese Twins: one would expect that two securities with claims on the same cash flows, which still trade separately, should trade at the same price, but this is not always the case.

(2) Political Cycle Effect: the first and last years of a political administration have higher returns versus other years, perhaps as the market reacts to expected upcoming policy changes.

(3) Stock Split Effect: a stock split often occurs when a stock's price is too large (i.e., to improve the salability of the stock). Although, the total market cap remains unchanged, returns are observed to be higher both before and after the announcement of the split.

(4) Neglected Firm Effect: firms that are very small, and often neglected by analysts, and thus have low liquidity in the market, often experience abnormally high returns.

(5) Super Bowl Effect: When the Super Bowl winning team is from the AFC (of the NFL), the stock market for the upcoming year is likely to do worse than if the winning team is from the NFC. Although this relationship has been remarkably well supported by the data, there is no theoretical basis for it. This, this is likely to be an example of a caveat mentioned earlier whereby data mining can sometimes lead to spurious conclusions.

Note that Berk and DeMarzo (2017), in Section 13-6, discuss further examples of market anomalies for the size effect, the value effect (in terms of book value to market value ratio), and the momentum effect (also discussed earlier). These factors are incorporated further in the Fama-French-Carhart multifactor model discussed in Section 13-7.

4.2.4 Bubbles and Irrational Exuberance

Although bubbles are not market anomalies *per se*, they do serve as examples of market inefficiency. A bubble occurs when investors make decisions, perhaps irrationally, more caught up in a speculative frenzy than subject to traditional financial fundamentals. This causes asset prices to reach lofty levels that cannot be justified by earnings and dividends, as other investors join the herd mentality under the assumption that prices can only rise further.

For a period of time, the bubble is self-sustaining, but these supernormal rates of growth cannot continue forever. Eventually, the bubble will burst, creating potentially huge losses for those who did not time the market well, (i.e., buying in too late and/or selling out too late).

A key characteristic of a bubble is that investors lose confidence in benchmarks for what price levels should be. Thus, traders who normally act rationally may be temporarily confused, which in turn pumps dangerously high volatility and trading volume into the market, that is, until a price equilibrium can be re-established.

Briefly, there have been three leading examples of bubbles in the last 30-35 years:

(1) The Japanese stock and real estate bubble: in the late 1980s, the Nikkei 225 rose about 300%, but after interest rates increased sharply in the early 1990s, stock prices soon thereafter lost about 50% of their value within 9 months. By the late 2000s, stock prices were down 80% and real estate prices had fallen 87% from their peak levels.

(2) The Technology or dot.com bubble: in the late 1990s, the Nasdaq Composite Index rose about 580%, peaking in early 2000. By late 2002, prices were down 78% from their peak levels. Many of the most volatile stocks during this time were high-tech or dot.com stocks, many of which no longer exist today.

(3) The U.S. housing bubble and crisis: From 1996-2006, U.S. housing prices increased nearly 300%, as many investors assumed that prices could only continue to go up, and banks extended credit to individuals beyond what normal lending standards would allow. By March 2009, U.S. home prices were down more than 30% from their peak.

As this note is being written, we may be in the midst of another bubble, this one concerning the cryptocurrency known as bitcoin (Krugman, 2018). Time will tell if it is indeed a bubble.

5   References

- Berk, J. and DeMarzo, P. (2017), *Corporate Finance*, 4th edition, Boston: Pearson.
- Brealey, R., Myers, S., and Allen, F. (2017), *Principles of Corporate Finance*, 12th edition, New York: McGraw-Hill Education.
- Brealey, R., Myers, S., and Allen, F. (2014), *Principles of Corporate Finance*, 11th edition, New York: McGraw-Hill Irwin.
- Hardy, M. (2006), "An Introduction to Risk Measures for Actuarial Applications," (Former Exam C Study Note) from *The Casualty Actuarial Society and the Society of Actuaries.*
- Kendall, M. (1953), "The Analysis of Economic Time Series, Part I. Prices," *Journal of the Royal Statistical Society*, 96, p.11-25.
- Keown, A. and Pinkerton, J. (1981), "Merger Announcements and Insider Trading Activity," *Journal of Finance*, 36, p.855-869.
- Klugman, S., Panjer H., and Willmot, G. (2012), *Loss Models: From Data to Decisions*, 4th edition, New York: Wiley.
- Krugman, P. (2018), "Bubble, Bubble, Fraund and Trouble," *New York Times*, January 29, 2018, accessed online at https://www.nytimes.com/2018/01/29/opinion/bitcoin-bubble-fraud.html?action=click&pgtype=Homepage&clickSource=story-heading&module=opinion-c-col-left-region&region=opinion-c-col-left-region&WT.nav=opinion-c-col-left-region
- Malkiel, B. (1995), "Returns for Investing in Equity Mutual Funds," *Journal of Finance*, 50, p.549-572.
- McDonald, R. (2013), *Derivative Markets*, 3rd edition, Boston: Pearson.
- Poterba, J. and Summers, L. (1988), "Mean Reversion in Stock Prices: Evidence and Implications," *Journal of Financial Economics*, 22, p.27-60.