

Classes 7 & 8

Time-Varying Expected Returns and Volatility

The following was written for the master-level class I taught at MIT Sloan a few years ago. If time permits, I'll revise the content to make it more suitable for a PhD level class.

1 Return Predictability and Market Efficiency

We will start with the efficient market hypothesis, using it as a framework to help us understand what it means to be able to predict the market. People often believe that market efficiency means that returns are unpredictable. This is not true. In an economy with time-varying business condition or time-varying risk appetite, the expected returns are time-varying and, most likely, persistent. As a result, you will see return predictability. The more relevant question is: How strong is the predictability? We will look at some of the empirical evidences.

Follow the Information: The financial markets are an information central, where people bring their information to trade. If it is a correct and useful piece of information, which has not yet been incorporated into the price, then there is room for profit. But as soon as the market price adjusts to the news, the information loses its usefulness and there is no longer any profit to be made with this piece of information.

So when it comes to predicting the market, one should follow the flow of information. A trader who wants to make a profit from predicting the market should always ask himself: Am I good at collecting information? If so, then I have all the incentive to do so because I will be rewarded for bringing this information to the financial markets. The next question is: What kind of information am I good at collecting, macro-level for the entire economy, or micro-level for individual stocks? Depending on your talent, the nature of your trading strategy will be very different. Global macro funds place directional bets on the overall market: interest rate, foreign exchange, and maybe the stock market. Long/short equity funds or fixed-income arbitrage funds avoid taking any directional bet. Instead, they focus on the relative mis-pricing between groups of stocks or bonds. At the super high-frequency

domain, where the life span of information is on the order of milliseconds, market making funds and statistical arbitrage funds populate this space to facilitate trades and provide liquidity.

All of these market players are motivated by a common goal: making a profit. And they are able to do so by bringing information to the markets. As a result of these efforts, new and relevant information gets incorporated into the prices. And the markets become more efficient.

The Efficient Market Hypothesis: It's impossible to talk about market predictability without bringing up the efficient market hypothesis, or the question about market efficiency. So let me spend some time clarifying some of the confusions.

First of all, in my personal opinion, the efficient market hypothesis is simply a statement that defines what it means to have an efficient market: when market prices incorporate information. It is like saying, being happy means to have peace of mind.

Second, without a proper asset pricing model, there is no way to test the efficient market hypothesis. With a proper model, then we are simply testing the model (e.g., the CAPM) which usually assumes market efficiency. So there is really no point in sweating over it. To be more specific, the efficient market hypothesis is not a stand alone test on market efficiency. It is always a joint test. Market efficiency can only be tested in the context of an asset pricing model that specifies equilibrium expected returns. For example, market efficiency implies zero predictability only if the expected returns that investors require to hold stocks are constant through time (or at least serially uncorrelated). Otherwise, if expected stock returns are time-varying and persistent, then there will be predictability in stock returns and it does not imply at all market inefficiency.

Third, Finance in general and efficient market hypothesis in particular is really not a system of beliefs. What we can offer in Finance are tools. Tools for clear thinking. Don't believe, don't don't believe. Use the tools, apply them to the data and to your own experiences, make an honest and sincere effort to figure things out for yourself.

Orange Juice: Since we are on the topic of market efficiency, let me tell one story that impressed me the most over the years. It is about orange juice, written in a 1984 paper by Prof. Roll from UCLA. It is the kind of paper I've always wanted to write: simplicity at its best; maximum power with minimum fluff.

Cold weather is bad for orange production. Orange trees cannot withstand freezing temperatures that last for more than a few hours. The central Florida region around Orlando, which accounts for more than 98 percent of U.S. production of frozen concentrated orange juice, occasionally has freezing weather. During the 6 and 1/4 year period studied by Prof. Roll, there

were four periods when the temperatures were below 30°F, each accompanied by significant price increase in orange juice futures prices.

Overall, the most important determinant in the pricing of orange juice futures is weather in central Florida. Quoting Prof. Roll, “So if the OJ futures market is an efficient information processor, it should incorporate all publicly available long-term and short-term weather forecasts. Any private forecasts should be incorporated to the extent that traders who are aware of those forecasts are also in command of significant resources. The futures price should, therefore, incorporate the predictable part of weather in advance.”

With this idea in mind, Prof. Roll uses the OJ futures prices to predict the weather. Not surprisingly, you will find a relationship between the two. The ingenious design of Prof. Roll’s regression is to find out if the OJ futures prices can predict weather more accurately than the National Weather Service. On the left hand side of his regression is the temperature forecast error, which is the percentage difference between the actual temperature and the forecast temperature provided by the National Weather Service. On the right hand side of his regression is the returns on orange juice futures.

What did he find? Orange juice futures prices are better at predicting the weather than the National Weather Service. This predictability is especially strong for the P.M. temperature forecast because of the sensitivity of orange trees to freezing temperatures.

The Value of Millisecond: Let me tell you another story that fascinated me. It borders on craziness, but is a good story. This is the first paragraph of *Flash Boys*, a recent book by Michael Lewis.

“By the summer of 2009 the line had a life of its own, and two thousand men were digging and boring the strange home it needed to survive. Two hundred and five crews of eight men each, plus assorted advisors and inspectors, were now rising early to figure out how to blast a hole through some innocent mountain, or tunnel under some riverbed, or dig a trench beside a country road that lacked a roadside – all without ever answering the obvious question: *Why?* The line was just a one-and-a-half-inch-wide hard black plastic tube designed to shelter four hundred hair-thin strands of glass, but it already had the feeling of a living creature, a subterranean reptile, with its peculiar needs and wants. It needed its burrow to be straight, maybe the most insistently straight path ever dug into the earth. It needed to connect a data center on the South Side of Chicago to a stock exchange in northern New Jersey. Above all, apparently, it needed to be a secret.”

All of these efforts just so the speed of information transmission can be improved in the order of ... millisecond. Let me quote Lewis again, since he is a much better writer.

“One way to price access to the line, Tabb thought, was to figure out how much money might be made from it, from the so-called spread trade between New York and Chicago – the simple arbitrage between cash and futures. Tabb estimated that if a single Wall Street bank were to exploit the countless minuscule discrepancies in price between Thing A in Chicago and Thing A in New York, they’d make profits of \$20 billion a year. He further estimated that there were as many as four hundred firms then vying to capture the \$20 billion.”

Market Efficiency not a Marble Statue: In telling the previous two stories, I would like to impress upon you the process through which markets become efficient. Market efficiency is not really a doctrine for you to believe or disbelieve. It is a process, a process of arbitrageurs participating in the markets with the objective of making a profit. Sometimes, this process works; sometimes it fails. It is an organic process, not a marble statue.

After the 2008 financial crisis, many people were hard on the efficiency market hypothesis. Some people believed that the financial crisis was the result of a misguided faith in market efficiency that encouraged market participants to accept security prices as the best estimate of value rather than conduct their own investigation. Some wrote that among the causes of the recent financial crisis was an unjustified faith in rational expectations, market efficiencies, and the techniques of modern finance.

Seriously, I really don’t know how these people got their ideas. Rational expectation builds on the understanding that all players in the market are motivated to optimize their risk and return tradeoff; market efficiency does not happen in the vacuum; it happens only when investors bring their information to the market with the objective of making a profit; and techniques of modern finance do help reduce trading cost and improve risk sharing in the society.

As to 2008? The flow of information broke down at some point. Large banks were sitting on supposedly super safe tranches of CDO and CDO2 without realizing or the willingness to realize the real risk. The rest of the market had a very limited access to this kind of balance-sheet (or off balance-sheet) level information and the market prices failed to incorporate this information. But did the banks take these positions out of their belief of market efficiency? I really doubt it.

Market Inefficiency and Limits to Arbitrage: Since we touched upon the topic of market efficiency, I think it would be fair to mention the Behavior Finance literature on market inefficiency. It was an area of Finance that grew in popularity after the tech boom of 1990s. If you are interested in this topic, you can start with Prof. Shleifer’s book, “Inefficient Markets: An Introduction to Behavioral Finance.”

The efficient market hypothesis assumes that the market incorporates the new information right away. In practice, however, there is uncertainty surrounding the information and the process of price discovery itself involves uncertainty. In certain situations, a correct piece of information might not get incorporated into the price right away. If the price moves temporarily in the opposite direction of his information, the arbitrageur might in fact lose money trading this information. This argument, proposed by Shleifer and Vishny (1997) and often referred to as “limits to arbitrage” can help explain why bubble can keep building up even when many people are calling it a bubble.

When Alan Greenspan, the then chairman of the Fed, gave the famous “irrational exuberance” speech in December 1996, the Nasdaq was around 1,300. Initially, the stock markets around the world dropped precipitously in reaction to the speech. But the markets soon shrugged off the warning and started the most spectacular upward trajectory in the history of Nasdaq. A little over three years after the speech, on March 10, 2000, the Nasdaq peaked at 5,048.62. Then it went down as fast as it came up, and bottomed near 1,140 two and half years later on October 4, 2002.

One person who shared the same view with chairman Greenspan was Prof Shiller, who later wrote a book titled “irrational exuberance.” Prof Shiller also shared the 2013 Nobel Prize with Prof Fama and Hansen. It was said that Prof Shiller, following his own prediction about the internet bubble, actually shorted the Nasdaq in the late 1990s, only to lose money because the market kept its upward trajectory for too long and crashed much too late.

2 Predicting the Market Returns

What We’ve Learned? Talking about market efficiency and market predictability at a hypothetical level is just not that interesting. Now that we are seven classes into the semester, maybe we can start from what we’ve learned so far.

By learning about the various quant strategies, we do recognize that the alpha generated by a quant strategy does come from a certain ability to predict the future. People might vary in their opinion on whether the alpha comes from market inefficiency (under/over-reactions) or systematic risk exposure. One observation I am sure that you’ve made is that quant investors do not take a stand on the market risk. If possible, they choose to avoid the market risk by taking long/short positions of two portfolios with similar beta exposures.

And yet, the market risk remains the most important and pervasive. You’ve probably noticed in our Assignment 1 that market-neutral hedge funds are not really market neutral. For example, the hedge fund index in long/short equity has a beta around 0.40. Even for

market-neutral hedge funds, the beta exposure is non-zero: around 0.20. Overall, the market risk is an important risk and let's try to understand it more.

In this class, we will focus on the “first moment” of the aggregate stock market and move on to the “second moment” in the next class.

How Good are Investors at Predicting the Market? You must have heard this famous story about Rockefeller and his shoe shine boy. After receiving unsolicited stock tips from his shoe shine boy in 1928, Rockefeller decided to get out of the stock market. His rationale: when a shoe shine boy started to give stock tips, the market probably was reaching its peak.

I don't know if the story actually happened to Rockefeller, but the gist of the story got repeated again and again in the history of financial markets. Last year, from July 2014 to July 2015, I was on sabbatical and spent most of my time in Shanghai with my parents. I was living a very simple life, far away from the financial establishments in Shanghai. Yet one can hardly avoid the hype and then the disappointment of the stock market. I had to tell my 80-year old father repeatedly that his optimal allocation to the stock market is zero, regardless of how much money other people were making out of the market.

The empirical evidence paints a similar story: investors have no ability to predict the future. In fact, their prediction is a response to the stock market. When the markets are doing well, their prediction is optimistic; when the markets are doing poorly, they become pessimistic. Moreover, their prediction affects their behavior. The flow to equity mutual fund is driven heavily by the recent stock market performance. The same pattern of flow chasing performance can also be found in bond mutual funds.

Use Past Returns to Predict: For anyone wanting to predict the stock market, probably the very first regression would be:

$$R_{t+1} = a + \rho R_t + \epsilon_{t+1}.$$

Given the time-series data, this is the easiest regression to run. The results are mixed, depending on the horizon over which this regression is run. At the monthly horizon, the autocorrelation ρ is generally positive and statistically significant. The magnitude is small for the value-weighted portfolio and becomes larger for the equal weighted portfolio. But this result is not very stable and could flip sign or become insignificant during sub-sample analyses. Overall, the R-squared of this predictive regression is very small, indicating that much of the future returns remains unpredictable.

In academic, there is a pretty large literature on this topic: If you are interested, you can

read “Permanent and Temporary Components of Stock Prices” by Fama and French (1988), who ran this regression over a horizon of 3-5 years and found large negative autocorrelations. In “Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test,” Lo and MacKinlay (1988) propose the innovative variance ratio test as an alternative to the regression analysis. In “When are Contrarian Profits due to Stock Market Overreaction?” Lo and MacKinlay (1990) explain that despite negative autocorrelation in individual stock returns, weekly portfolio returns are positively autocorrelated and are the result of important cross-autocorrelations.

Stock Returns and Business Cycle: In any Finance model, one main driver for stock returns should be the underlying economic condition. Nevertheless, the link between the two is not that strong in the data. For example, Prof. Shiller wrote a paper in 1984 entitled, “Do Stock Prices Move Too Much to be Justified by Subsequent Change in Dividend?” In this paper, he made the observation that the stock market is too volatility (e.g., 20% per year) compared to the volatility in the fundamental: dividends or earnings.

If you plot the time-series of realized stock returns against the business cycle, you do find a link between the two. In particular, depressed expected business conditions are associated with high expected excess returns. This observation gives rise to predictive regressions using a set of variables that are related to business conditions, including default spreads, term premiums, and dividend-price ratio. By far, the best predictor for stock market returns is the dividend-price ratio. We will re-visit the default spread and term premium as we cover the fixed-income market.

Dividend-Price Ratio as a Predictor: Let’s run this regression at the annual frequency:

$$R_{t+1} = a + b \left(\frac{D}{P} \right)_t + \epsilon_{t+1},$$

where D/P is the dividend-price ratio (aggregate dividend divided by the value-weighted CRSP index). The general finding that is the coefficient b of this predictive regression is positive and statistically significant.

The key to this regression is $1/P$. The aggregate price level is usually depressed during poor business condition (e.g. recessions). Going forward, the stock return is expected to be high. Hence the positive regression coefficient. Using D/P is just a way to scale the overall time trend of stock price increase. Using aggregate earnings, one can use replace D/P by E/P , although the earnings number is more noisy and biases the regression coefficient downward.

One important observation of this predictive regression is that the power of predictability is very weak even for the best predictor. At an annual frequency, the R-squared of this

predictive regression is around 5%, indicating that 95% of the future variance remains unpredictable.

Market Timing as a Trading Strategy: This kind of result explains why market timing is not a very popular trading strategy among market participants. Of course, there is nothing wrong with taking a long position on the market if your objective is to be compensated from such a risk exposure. But if you are in and out of the market with the belief that you can predict the market, then you should be very careful. The empirical evidence tells you that you are going to be exposed to quite a bit of uncertainty. Moreover, this is a very special kind of uncertainty — the market risk, the most dangerous kind.

This is why David Swensen wrote in his book, “Market timing, according to Charles Ellis, represents a losing strategy: There is no evidence of any large institutions having anything like consistent ability to get in when the market is low and get out when the market is high. Attempts to switch between stocks and bonds, or between stocks and cash, in anticipation of market moves have been unsuccessful much more often than they have been successful. Serious investors avoid timing markets.”

Predicting Bond Returns: Let’s start with Fama and Bliss (1987). Consider the following simple strategy: At time t , purchase a discount bond, $P_{t,n}$, with n years to mature. One year later at time $t + 1$, sell the same bond, $P_{t+1,n-1}$, which has $n - 1$ years to maturity. The log-return of this “long-bond” strategy is

$$h_{t+1}^{(n)} = \ln P_{t+1,n-1} - \ln P_{t,n}.$$

Alternatively, one can purchase a one-year discount bond and wait for it to mature one year later. The log-return on this “short-bond” strategy is known at time t to be $r_{t,1}$, the continuously compounded one-year yield. The difference between the two, $h_{t+1}^{(n)} - r_{t,1}$, indicates how attractive long-term bonds are relative to the short-term bond.

Just like what we’ve learned from the equity market, a lot of effort has been put into predicting the excess return of long term bonds. The conditioning information used in this type of predictive regressions, however, is more straightforward. Fama and Bliss (1987) use the n -year forward-spot spread to predict the excess return of a n -year bond,

$$h_{t+1}^{(n)} - r_{t,1} = a + b \left(f_t^{(n)} - r_{t,1} \right) + \epsilon_{t+1}^{(n)},$$

where

$$f_t^{(n)} = \ln P_{t,n-1} - \ln P_{t,n},$$

is the time- t forward rate from $n - 1$ to n .

According to the expectations hypothesis, long yields are the average of future expected short yields, indicating that excess returns should not be predictable. Fama and Bliss (1987) find that the yield spread can in fact predict long-bond returns with R-squared about 18%. Campbell and Shiller (1991) provide further evidences against the expectations hypothesis. The bond predictability was further refined and extended by Cochrane and Piazzesi (2005), who find that a single factor, a single tent-shaped linear combination of forward rates, predicts excess returns on one- to five-year maturity bonds with R-squared up to 44%. The return-forecasting factor is countercyclical and forecasts stock returns.

Predicting Currency Returns: In this domain, the currency carry strategy provides the strongest evidence of return predictability, although this predictability requires a cross-section of currencies. Focusing on one currency, there is no well-established empirical results on predictability, although some of the largest directional trading profits in macro hedge funds were from the currency market.

3 Time-Varying Return Volatility

Just when Prof. Fama and his PhD/MBA students were busy working on the cross section of expected stock returns, another area of Finance was taking shape. In this area, tools developed in Econometrics and Statistics are applied to financial time series such as the time-series of stock returns. Given how difficult it is to estimate the first moment (expected returns), much of the attention was devoted to estimating the second moment, stock return volatility. The most visible figure in this area is Prof. Rob Engle, who was awarded a Nobel Prize in 2003 for “methods of analyzing economic time series with time-varying volatility (ARCH).” The ARCH paper was published in 1982 when Prof. Engle was an Economics professor at UCSD. The more famous GARCH extension was later published in 1986 by his PhD student Prof. Bollerslev.

3.1 Market Risk Measurement

Need for Better Risk Management Tools: In the early 1990s, there were two developments that made volatility models attractive and relevant. First, the need of a better option pricing model becomes quite obvious after the 1987 stock market crash. The Black-Scholes model builds a very good foundation, but it lacks flexibility in handling the richer reality. In the Black-Scholes model, stock returns are normally distributed with a constant volatility

σ . Any casual inspection of the data would inform us that volatility is not a constant. So having a better volatility model would be a first step toward a better option pricing model. Second, and even more pressing was the need for better risk management tools. The mortgage-backed security was developed in the 1980s, and the over-the-counter derivatives market started to take off by the late 1980s. By the early 1990s, the increasing activity in securitization and the increasing complexity in the fixed-income products have made the trading books of many investment banks too complex and diverse for the chief executives to understand the overall risk of their firms.

It was an industry wide phenomenon. For example, in “Money and Power,” Cohan wrote about the difficult year of 1994 at Goldman after the two amazingly profitable years in fixed-income in 1992 and 1993. In the book, he quoted Henry Paulson, “What came out of the 1994 debacle was best practices in terms of risk management, The quality of the people, and the processes that were put in place – anything from the liquidity management to the way we evaluated risk and really the independence of that function – changed the direction of the firm.”

JPMorgan’s RiskMetrics: Among all Wall Street firms, JPMorgan’s effort was by far the most visible and influential. The 2009 New York Times article titled “Risk Mismanagement” gave a very good account of the events.

JPMorgan’s chairman at the time VaR took off was a man named Dennis Weatherstone. Weatherstone, who died in 2008 at the age of 77, was a working-class Englishman who acquired the bearing of a patrician during his long career at the bank. He was soft-spoken, polite, self-effacing. At the point at which he took over JPMorgan, it had moved from being purely a commercial bank into one of these new hybrids. Within the bank, Weatherstone had long been known as an expert on risk, especially when he was running the foreign-exchange trading desk. But as chairman, he quickly realized that he understood far less about the firm’s overall risk than he needed to. Did the risk in JPMorgan’s stock portfolio cancel out the risk being taken by its bond portfolio – or did it heighten those risks? How could you compare different kinds of derivative risks? What happened to the portfolio when volatility increased or interest rates rose? How did currency fluctuations affect the fixed-income instruments? Weatherstone had no idea what the answers were. He needed a way to compare the risks of those various assets and to understand what his companywide risk was.

What later became RiskMetrics was an internal effort developed within JPMorgan in 1992 in response to the CEO’s question. Quoting the New York Times article again,

By the early 1990s, VaR had become such a fixture at JPMorgan that Weatherstone instituted what became known as the 415 report because it was handed out every day at 4:15, just after the market closed. It allowed him to see what every desk's estimated profit and loss was, as compared to its risk, and how it all added up for the entire firm. True, it didn't take into account Taleb's fat tails, but nobody really expected it to do that. Weatherstone had been a trader himself; he understood both the limits and the value of VaR. It told him things he hadn't known before. He could use it to help him make judgments about whether the firm should take on additional risk or pull back. And that's what he did.

Global Risk Factors: In 1994, JPMorgan started to make RiskMetrics publicly available. It published its technical document outlining its risk measurement methodologies. It also made available two sets of volatility and correlation data used in the computation of market risk. In Spring 1996, I was hired as a research assistant to Prof. Darrell Duffie to work on risk management and Value-at-Risk. I spent a lot of time reading the technical document of RiskMetrics and I also downloaded the volatility and correlation dataset everyday just to play with them and track the movements. It was through having to deal with the datasets that the immensity of the global financial markets became real to me.

For example, the 1996 RiskMetrics data files covered over 480 financial time series that were important for the trading books of most investment banks. This includes

- Equity indices across the world.
- Foreign exchange rates.
- The term structure of interest rates across the world:
 - money market rates (1m, 3m, 6m, and 12m) for the short end.
 - government bond zero rates (2y, 3y, 4y, 5y, 7y, 9y, 10y, 15y, 20y, and 30y) for the longer end.
- The term structure of swap rates across the world (2y, 3y, 4y, 5y, 7y, and 10y).
- Commodities: spot and futures of varying maturities.

Variance-Covariance Matrix: In order to measure the firm-wide risk exposure, one need to first map each position to these risk factors, and then calculate the volatility of the overall portfolio or portfolios by asset class: interest rates, equity, currency, and commodities. One of the key building block of this calculation is the variance-covariance matrix of the risk factors. For the 480 risk factors used by JPMorgan in 1996, this involves calculating the volatility for each of the 480 risk factors, and then calculate the pair-wise correlations between the 480 risk factors. In the rest of the class, we will be busy doing these calculations.

Let me quote a few paragraphs from the 2012 annual report of Goldman Sachs so as to give you an update on the risk management effort on Wall Street since 1996.

We also rely on technology to manage risk effectively. While judgment remains paramount, the speed, comprehensiveness and accuracy of information can materially enhance or hinder effective risk decision making. We mark to market approximately 6 million positions every day. And, we rely on our systems to run stress scenarios across multiple products and regions. In a single day, our systems use roughly 1 million computing hours for risk management calculations.

When calculating VaR, we use historical simulations with full valuation of approximately 70,000 market factors. VaR is calculated at a position level based on simultaneously shocking the relevant market risk factors for that position. We sample from 5 years of historical data to generate the scenarios for our VaR calculation. The historical data is weighted so that the relative importance of the data reduces over time. This gives greater importance to more recent observations and reflects current asset volatilities, which improves the accuracy of our estimates of potential loss. As a result, even if our inventory positions were unchanged, our VaR would increase with increasing market volatility and vice versa.

As you can see, roughly 6 million positions are mapped into 70,000 market factors in Goldman's risk management system. If I understand their statement correctly, this implies a variance-covariance matrix of 70,000 by 70,000.

Back in the mid-1990s, all three of my Chinese classmates at the NYU Physics department went to work at Citibank after graduation. I was the only exception, who went on to get another PhD in Finance. So much for the future of Physics, which was better off without us. During one of my visits back to New York, I visited them at Citibank, thinking how exciting it was for them to be working in the *real* world with a real paycheck. And I was very surprised to see how bored they all looked. One of them worked in the risk management group and his job was to calculate the variance-covariance matrix everyday. He looked miserable. I guess this is not the most exciting job if you have to do it everyday. Many years later, I got an email from Mr. Variance-Covariance, who has become a managing director at Citibank. A happy ending, by Wall Street standard.

3.2 Estimating Volatility using Financial Time Series

In general, volatility is very easy to estimate. Unlike in the case of expected returns, volatility can be measured with better precision using higher frequency data. The convention in this field is to use daily data. For stock market returns, having one month of daily data could get

you a pretty accurate estimate. We will continue to use the time series of aggregate stock returns as an example.

I should mention that in this field, log returns are being used more often than percentage returns:

$$R_t = \ln S_t - \ln S_{t-1},$$

where S_t could be date- t stock price, currency rate, interest rate, or commodity futures price. At the daily frequency, the magnitude of returns are generally very small. Using the handy Taylor expansion, we know that for small x ,

$$\ln(1 + x) \approx x.$$

Repeating this for log-returns, we have

$$R_t = \ln S_t - \ln S_{t-1} = \ln \left(\frac{S_t}{S_{t-1}} \right) = \ln \left(1 + \frac{S_t - S_{t-1}}{S_{t-1}} \right) \approx \frac{S_t - S_{t-1}}{S_{t-1}}.$$

In other words, working with log-returns or percentage-returns does not make too much of a difference when returns are small in magnitude.

Also notice that our attention is no longer on the first moment. Getting the volatility right is our main task. We calculate the variance by,

$$\text{var}(R_t) = E(R_t - \mu)^2 = E(R_t^2) - \mu^2.$$

The volatility estimate is

$$\text{std}(R_t) = \sqrt{\text{var}(R_t)} = \sqrt{E(R_t^2) - \mu^2} = \sqrt{E(R_t^2)} \times \sqrt{1 - \frac{\mu^2}{E(R_t^2)}}.$$

At the daily frequency, μ is around a few basis points for the US equity market, while the daily volatility is around 100 basis points (i.e., 1%). As a result, $\mu^2/E(R_t^2)$ is a really really small number and $\sqrt{1 - \mu^2/E(R_t^2)} \approx 1 - \frac{1}{2}\mu^2/E(R_t^2)$ is very close to one. So it does not make a big difference whether or not we subtract μ from the realized returns such as R_t in estimating the volatility. You will notice that most of the time, people drop μ for simplicity:

$$\text{std}(R_t) \approx \sqrt{E(R_t^2)}.$$

Of course, this relationship between daily vol and daily μ exists mostly true for assets in the risky category. For fixed income product, this might not be true. If you would like to be safe, one way to approach the data is to first demean the time-series data: $R_t - \mu$, and

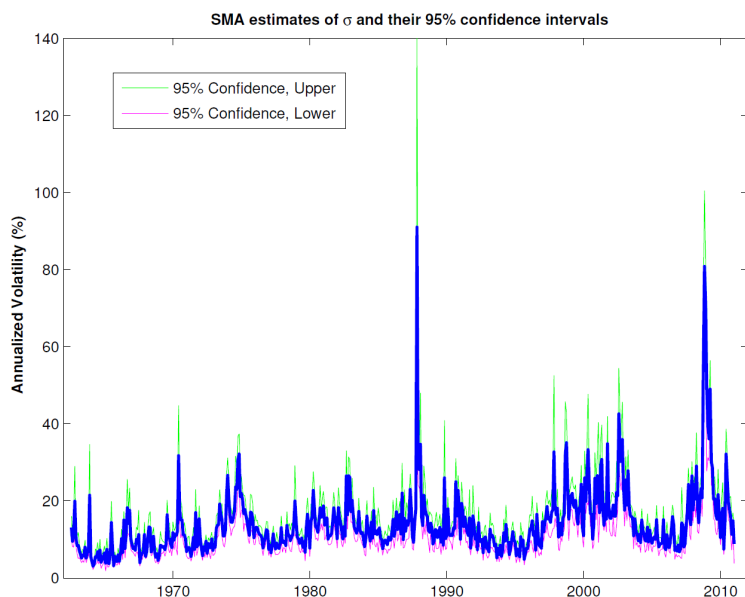


Figure 1: Time-Series of Stock Volatility using SMA.

then apply the volatility models.

SMA: The simple moving average model fixes a window, say one month, and use the daily returns within this window to calculate the sample standard deviation. The window is then moved forward by one step, say one day, and the whole calculation gets repeated again.

In Figure 1, I use a moving window of one month and move the window one month at a time to plot a monthly time-series of SMA volatility estimates. In this space, volatility is usually quoted at an annualized level. So I multiply the volatility estimates by $\sqrt{252}$ (assuming 252 business days per year). This annualized volatility corresponds to the volatility coefficient σ in the Black-Scholes model. So we also compare these volatility numbers with the option-implied volatility.

Just so you are convinced that these volatility estimates can be estimated with a pretty good precision using only one month of daily data, I also plotted the 95% confidence intervals. If you compared Figure 1 against Figure 2, you can see the marked difference in estimation precision. Using one month of daily data to estimate the average return, what you get is very much noise.

Another observation I would like you to make is the variation of market volatility over time. Its pattern is very different from that of market returns. Volatility is persistent: a day of high volatility is usually followed by another day of high volatility. Volatility also tends to spike up once in a while. If you plot these events against the NBER business cycles, you see that volatility usually spikes up during recessions. But recessions are not the only time




Figure 2: Time-Series of Monthly Average Stock Returns using SMA.

when volatility spikes up. Whenever the market is in trouble, volatility goes up. Let me name a few recent events: the 1987 stock market crash, the 1997 Asian Crisis, the 1998 LTCM crisis, the 2000-01 tech bubble/burst, the 9/11, and the 2008 financial crisis. Finally, whenever volatility is at an usually high or low level, it tends to revert back to its historical average. This pattern is called mean reversion. Using a longer sample that includes the Great Depression, the historical average of volatility is around 20%. Using the more recent sample, the average volatility is around 15%.

The plot has not been updated for the last few years. Nowadays, the best place to get stock market volatility (without having to do any calculation) is the CBOE VIX index. For example, over the one-week period from August 17 to 24, 2015, the VIX shoots up from 13.02% to 40.74%, because of the concern over the Chinese stock market.

Comparing the volatility estimate with the VIX index, which is effectively the option implied volatility, you notice that the option implied volatility seems to be consistently higher than the volatility estimate. We will visit this issue again when we cover the options market.

EWMA: The exponentially weighted moving average model is an improvement over the SMA model. Instead of applying equal weights to all observations with a fixed window, EWMA applies an exponential weighting schedule. It chooses a decay factor λ , which is a




Figure 3: The Exponential Weighting Scheme in EWMA.

number between zero and one, and performs the volatility estimate by:

$$\sqrt{(1 - \lambda) \sum_{n=0}^N \lambda^n (R_{t-n})^2}.$$

Let's first put aside the term $1 - \lambda$ and focus on the terms within the summation. We put a weight of 1 for today ($n = 0$), λ for yesterday, λ^2 for the day before yesterday, and so on. With this weighting schedule, as we move further back into the history and away from today t , the contribution of $(R_{t-n})^2$ decreases according to the exponential schedule. Hence the name. Figure 3 gives a graphical presentation of this exponential weighting scheme.

Going back to one of the paragraphs I quoted from Goldman's annual report: "We sample from 5 years of historical data to generate the scenarios for our VaR calculation. The historical data is weighted so that the relative importance of the data reduces over time. This gives greater importance to more recent observations and reflects current asset volatilities, which improves the accuracy of our estimates of potential loss." So effectively, the length of the window N is set at five years and a decay factor is selected to put more weight to the more recent events. In Goldman's report, the value of the decay factor was not reported. In RiskMetrics, λ was fixed at 0.94 for all time series. Also, the choice of the window size is not important because the decay factor λ effectively selects the window size for you. Figure 3 gives a nice graphical presentation on how the window size is determined by the decay factor: a strong decay factor ($\lambda = 0.8$) implies a smaller window while a mild decay

factor ($\lambda = 0.97$) implies a larger window. A window of 5 years is definitely not necessary: the return happened 5 years ago has a weight of $\lambda^{5 \times 252}$, which is too small to matter.

Now let's come back to the term $1 - \lambda$. Notice that

$$1 + \lambda + \lambda^2 + \lambda^3 + \dots = 1/(1 - \lambda).$$

So $(1 - \lambda)$ is there because of normalization. In the same way, the normalization factor in the SMA model is $1/N$. In the SMA model, if I increase the window size N , then each observation carries a smaller weight: a smaller $1/N$. Likewise, if I change λ from 0.94 to 0.97 in EWMA, the effective window size increases (see Figure 3). As a result, each observation carries a smaller weight: a smaller $1 - \lambda$.

SMA and EWMA: The difference between these two volatility estimates becomes most visible immediately after a large price movement. Figure 4 uses the famous Black Wednesday of 1992 as an example to illustrate this point. This technical document by RiskMetrics was first written around 1994. If it were written today, then the 2008 crisis would be plotted here as an example.

After a large price movement, up or down, the response of the EWMA estimate is very fast, because it carries a higher weight for the most recent event. If the market calms down after the large price movement, then the EWMA estimate will soon come down to a lower level. The behavior of the SMA estimate is the opposite. Its response is typically sluggish and it carries that piece of information for the duration of its window size. For this reason, the EWMA is the preferred volatility estimate when it comes to monitoring market volatility at the daily frequency.

Black Wednesday 1992: As a side, the 1992 sterling crisis was an important event in the global currency market. The followings are excerpts from Steven Drobny's book on "Inside the House of Money: Top Hedge Fund Traders on Profiting in the Global Markets."

The United Kingdom joined the European Exchange Rate Mechanism (ERM) in 1990 at a central parity rate of 2.95 deutsche marks to the pound. To comply with the ERM rule, the UK government was required to keep the pound in a trading band within 6 percent of the parity rate. In September 1992, as the sterling/mark exchange rate approached the lower end of the trading band, traders increasingly sold pounds against deutsche marks, forcing the Bank of England to intervene and buy an unlimited amount of pounds in accordance with ERM rules. Finally, on the evening of September 16, 1992, Great Britain humbly announced that it would no longer defend the trading band and withdrew the pound from the ERM system. The pound fell approximately 15 percent against the deutsche mark over the next

Chart 5.2
Log price changes in GBP/DEM and VaR estimates (1.65σ)




Figure 4: The Black Wednesday 1992 and the Volatility Estimates of SMA and EWMA.

few weeks, providing a windfall for speculators and a loss to the UK Treasury (i.e., British taxpayers) estimated to be in excess of £3 billion.

It was reported at the time that Soros Fund Management made between \$1-2 billion by shorting the pound, earning George Soros the moniker “the man who broke the Bank of England.” But he was certainly not alone in betting against the pound. In fact, the term global macro first entered the general public’s vocabulary on Black Wednesday.

Going back to our class on Predictability and Market Efficiency, there are few things to be learned. First, you predict the market by following the information, which, in this case, includes the ERM rule, the economic condition at UK, the government’s ability and political resolve to defend its currency. Second, the “arbitrage” is risky. In order for George Soros to make \$1 billion with a 15% drop in sterling, his short position at the time had to be over \$6 billion. This is the style of global macro: large and risky directional bets. Of course, they don’t always make money and we’ve seen a few times when Soros lost by the same order of magnitude. Third, most of the global macro opportunities (or losses) in currencies and emerging markets happened because of some frictions outside of the financial markets: currency pegging, government intervention, central bank and policy errors, etc. As the governments and central banks become smarter in their interaction with the markets, such outside returns may be slowly going away.

Computing EWA Recursively: Today is day $t - 1$. Let σ_t be the EWMA volatility estimate using all the information available on day $t - 1$ for the purpose of forecasting the volatility on day t . Notice the dating convention: the time- t estimate is observed on day $t - 1$. In my personal opinion, we should date σ_t by $t - 1$, not t . But this is the convention in this area. So let’s go with convention.

Moving one day forward, it’s now day t . After the day is over, we observe the realized return R_t . We now need to update our EWMA volatility estimator σ_{t+1} using the newly arrived information (i.e. R_t):

$$\sigma_{t+1}^2 = \lambda \sigma_t^2 + (1 - \lambda) R_t^2. \quad (1)$$

A good exercise for you would be to start right from the beginning,

$$\sigma_2^2 = \lambda \sigma_1^2 + (1 - \lambda) R_1^2$$

and then apply the recursive formula a few times to convince yourself that this recursive

approach does get you the exponential weighting scheme of EWMA:

$$\begin{aligned}\sigma_3^2 &= \lambda \sigma_2^2 + (1 - \lambda) R_2^2 = \lambda^2 \sigma_1^2 + (1 - \lambda) (\lambda R_1^2 + R_2^2) \\ \sigma_4^2 &= \lambda \sigma_3^2 + (1 - \lambda) R_3^2 = \lambda^3 \sigma_1^2 + (1 - \lambda) (\lambda^2 R_1^2 + \lambda R_2^2 + R_3^2) \\ &\dots \\ \sigma_t^2 &= \lambda^{t-1} \sigma_1^2 + (1 - \lambda) (\lambda^{t-2} R_1^2 + \lambda^{t-3} R_2^2 + \dots + R_{t-1}^2)\end{aligned}$$

For those of you who like things to be precise: as $t \rightarrow \infty$, we are back to the exact formulation of the EWMA. And whatever σ_1 we started with does not make a difference.

If you are an Excel user, you will appreciate the convenience of this recursive formula. If you care about saving CPU time, you will also appreciate the convenience of this recursive formula. When we update the information on day t to calculate σ_{t+1} , all of the past information has been neatly summarized by σ_t . The new information waiting for us to be included is the realization of R_t . We weight the new information R_t^2 by $1 - \lambda$ and “decay” the old information σ_t^2 by λ . Adding these two pieces together, we get the updated variance estimate. It would be difficult not to appreciate the elegance of this recursive approach. No?

Auto-Correlation Coefficient: Another way to understand the recursive formula of Equation (1) is that imposes the dynamic structure of σ^2 : persistent with an auto-correlation coefficient of λ .

Recall that we regress stock return R_{t+1} on its own lag R_t to examine the stock return predictability. We find that from 1926 to 2004, the auto-correlation coefficient is positive and statistically significant. But the magnitude of the correlation is very small. Moreover, this predictability is not very robust: over the various subsamples, the auto-correlation coefficients become statistically insignificant. In other words, the random walk model with zero auto-correlation is a reasonable model for the stock returns.

When it comes to the dynamic structure of volatility, however, the auto-correlation coefficient plays a rather important role. Models such as EWMA and GARCH became popular in practice because they allow volatility to be persistent with a high auto-correlation coefficient. In estimating the auto-correlation in stock returns, we can simply run a regression. In the case of volatility, however, we need to estimate the volatility along with the coefficient λ . For this, we need a more structured estimation approach than a regression. (If you get into this area called Econometrics, you will realize that the essence is really the same. In particular, a linear regression is really the product of a maximum likelihood estimation. See Appendix A.)

Estimating the Decay Factor: Figure 3 provides a graphical connection between the decay factor λ and the sample size. A strong decay factor, say $\lambda = 0.8$, pays more attention

to the current events and underweights the far-away events more strongly. As a result, the effective sample size is smaller with a stronger decay factor (e.g., smaller λ). As you can see from Figure 5, a strong decay factor improves on the timeliness of the volatility estimate, but the smaller sample size makes the estimate noisier and less precise. On the other hand, a weaker decay factor, say $\lambda = 0.97$, improves on the smoothness and precision, but that estimate could be sluggish and slow in response to changing market conditions, as reflected in Figure 5. So there is a tradeoff.

Minimize RMSE: Let's consider two ways to pick the optimal decay factor. In the first approach, we would like to minimize the forecast error between the model's prediction and the actual realization. Recall, on day t , we form σ_{t+1} as a forecast for the volatility on day $t + 1$. So the model's forecast error is $R_{t+1}^2 - \sigma_{t+1}^2$. Summing these forecast errors over the sample period, we calculate the root mean squared error (RMSE) by,

$$\text{RMSE} = \sqrt{\frac{1}{T} \sum_{t=1}^T (R_{t+1}^2 - \sigma_{t+1}^2)^2}$$

Note that the only parameter at our disposal is λ . Everything else comes from the data. So let's find the optimal λ^* that minimizes the forecast error:

$$\lambda^* = \arg \min_{\lambda \in (0,1)} \text{RMSE} = \arg \min_{\lambda \in (0,1)} \sqrt{\frac{1}{T} \sum_{t=1}^T (R_{t+1}^2 - \sigma_{t+1}^2)^2}$$

MLE: In the second approach, let's use the maximum likelihood estimation. To be honest, using the MLE on this problem is really an overkill, but I would like to use this opportunity to introduce you to MLE. Anybody working with data should have done MLE at least once in their life.

Recall that we talk about the pdf of a normal, which is a Gaussian function. In our current setting, the volatility is time-varying. So the stock return R_{t+1} is normally distributed only when conditioning on the volatility estimate σ_{t+1} :

$$f(R_{t+1} | \sigma_{t+1}) = \frac{1}{\sqrt{2\pi}\sigma_{t+1}} e^{-\frac{R_{t+1}^2}{2\sigma_{t+1}^2}}.$$

Notice that if I wanted to be very precise, I should have replaced R_{t+1}^2 by $(R_{t+1} - \mu)^2$ and use the MLE to estimate both λ and μ . But we talked about this. Setting $\mu = 0$ here is a good compromise to make.





Figure 5: Time Series of EWMA Volatility Estimates with Varying Decay Factors.

The next step of MLE is to take log of the pdf:

$$\ln f(R_{t+1}|\sigma_{t+1}) = -\ln \sigma_{t+1} - \frac{R_{t+1}^2}{2\sigma_{t+1}^2},$$

I dropped 2π since it is a constant will not affect anything we will do later. We now add them up to get what econometricians call log-likelihood (llk):

$$\text{llk} = -\sum_{t=1}^T \left(\ln \sigma_{t+1} + \frac{R_{t+1}^2}{2\sigma_{t+1}^2} \right)$$

As you can see, the only parameter in llk is our choice of λ . It turns out that the best λ is the one that maximizes llk. In practice, we take -llk and minimize -llk instead of maximizing llk.

What we just did came straight out of Econometrics. A good textbook on this topic is the Time Series Analysis by James Hamilton. Read in particular the chapter on Generalized Method of Moments. Most of the econometrics tasks we encounter in Finance can be understood from the perspective of GMM, which was developed by Prof. Lars Hansen at University of Chicago. Prof. Hansen shared the 2013 Nobel Prize with Prof. Eugene Fama and Prof. Robert Shiller. In the Appendix, I include my old PhD-era code for estimating the standard errors of mean, std, skewness, and kurtosis. As you can see, my approach was very much influenced by the GMM approach. In the Appendix, I wrote a brief note on MLE and linear regression, which could be a nice entry point to motivate you to learn more about Econometrics.

ARCH and GARCH: The ARCH model, autoregressive conditional heteroskedasticity, was proposed by Professor Robert Engle in 1982. The GARCH model is a generalized version of ARCH. ARCH and GARCH are statistical models that capture the time-varying volatility:

$$\sigma_{t+1}^2 = a_0 + a_1 R_t^2 + a_2 \sigma_t^2$$

As you can see, it is very similar to the EWMA model. In fact, if we set $a_0 = 0$, $a_2 = \lambda$, and $a_1 = 1 - \lambda$, we are doing the EWMA model.

So what's the value added? This model has three parameters while the EWMA has only one. So it offers more flexibility (e.g., allows for mean reversion and better captures volatility clustering). If you are interested in estimating the GARCH model, you can use the MLE method we just discussed. Instead of estimating σ_{t+1} using EWMA, you use the GARCH model. The EWMA has only one parameter λ to estimate. The GARCH model has three parameters to estimate a_0 , a_1 , and a_2 . You will find that, just like λ , a_2 is very close to one.

In fact, a_2 captures the auto-correlation of the variance σ_t^2 : an autocorrelation coefficient that is close to one indicates a very persistent time series. Moreover, after some calculation, you notice that the long-run mean of the variance in this model is $a_0/(1 - a_1 - a_2)$. You can see how having additional parameters could provide more flexibility to the model.

The GARCH model has a pretty strong influence, and you are encouraged to dig deeper into the model if it interests you. We used to study quite a bit of GARCH at Stanford GSB. But looking back, I feel that I get most of the key intuitions by working with EWMA. Where there are too many moving parts and too many parameters, you tend to focus more on dealing with the formulas and parameters and lose track of the essence of the problem. That's why simplicity is always preferred.

EWMA for Covariance As mentioned at the beginning, our goal is to create the variance-covariance matrix for the key risk factors influencing our portfolio. Suppose that there are two risk factors affecting our portfolios. Let R_t^A and R_t^B be the realized day- t returns of these two risk factors. We estimate the covariance between A and B by

$$cov_{t+1} = \lambda cov_t + (1 - \lambda) R_t^A \times R_t^B$$

And their correlation:

$$corr_{t+1} = \frac{cov_{t+1}}{\sigma_{t+1}^A \sigma_{t+1}^B},$$

where σ_{t+1}^A and σ_{t+1}^B are the EWMA volatility estimates.

This calculation of covariance and correlation is pretty straightforward once you master the EWMA recursive formula. But let me use this opportunity to bring in volatility as a risk factor and emphasize on its importance. As recent as the early 2000s, volatility as a risk factor was not widely monitored by market participants. Of course, sophisticated investors pay attention to their exposure to volatility risk. The general intuition is that if you are short on volatility, you are going to lose during crisis. On the other hand, if you are long on volatility, you are partially hedged during these crises. Exposures to volatility risk comes certain non-linearity in one's position. The most straightforward way to be long on volatility is to buy at-the-money S&P 500 index options. As we will cover in our options class, such long positions usually are expensive. That is, you are paying a premium for such positive exposures.

Since the 2008 financial crisis, the volatility risk has got a broader audience. By now, the VIX index is reported daily in a prominent position along with the Dow, the S&P, and Nasdaq. It's often called the fear gauge. Figure 6 plots the historical VIX for the past 15 years. Going over the various events in the past, you can certainly appreciate why it is called




Figure 6: Time Series of CBOE VIX.

fear gauge.

Another important observation about the volatility risk factor is its increasing negative correlation between the market risk. Figure 7 plots the EWMA correlation estimate between daily returns of the S&P 500 index and the daily changes in VIX. As you can see, this correlation is always negative. The fact that there is a negative correlation between the stock market return and its volatility has always been documented, well before CBOE published its VIX index. As you can see, during the early sample period, the correlation was hovering around -50%. When I was a PhD student working on this topic in the late 1990s, a typical number for this correlation would be -60%. There is certainly a trend of this correlation becoming more negative in recent times. After the 2008 financial crisis, this correlation has experienced a regime switch to a more negative territory.

A negative correlation implies that whenever market drops down, the volatility goes up. Using the interpretation of VIX as a fear gauge, this means that a down market is coupled with increasing fear. The more negative correlation in recent years means a higher level of sensitivity to down markets: a market sitting at its edge, more easily spooked. As we move on to the options market, we will look at the “fear” component in VIX more closely.

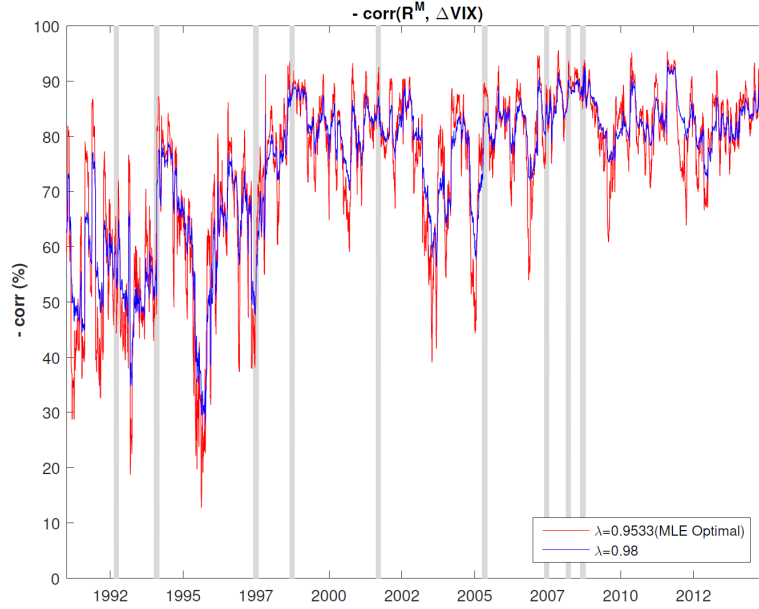


Figure 7: Time Series of CBOE VIX.

APPENDIX

A MLE and Linear Regression

Let's consider the linear regression:

$$Y_t = \alpha + \beta X_t + \epsilon_t,$$

where if we replace X with $R^M - r_f$ and Y with $R^i - r_f$, we are back with our favorite CAPM regression.

Thinking in terms of MLE, we focus on the distribution of the regression residual ϵ_t . We assume that ϵ_t is i.i.d. with normal distribution: zero mean and volatility of σ_ϵ . Now let's repeat the MLE steps for this regression:

The PDF for the Residual:

$$f(\epsilon_t) = \frac{1}{\sqrt{2\pi}\sigma_\epsilon} e^{-\frac{\epsilon_t^2}{2\sigma_\epsilon^2}}$$

Log-Likelihood:

$$\ln f(\epsilon_t) = -\ln \sigma_\epsilon - \frac{\epsilon_t^2}{2\sigma_\epsilon^2} = -\ln \sigma_\epsilon - \frac{(Y_t - \alpha - \beta X_t)^2}{2\sigma_\epsilon^2},$$

where 2π was again dropped.

Sample Log-Likelihood:

$$\text{llk} = \sum_{t=1}^T \left(-\ln \sigma_\epsilon - \frac{(Y_t - \alpha - \beta X_t)^2}{2\sigma_\epsilon^2} \right)$$

Minimize the -LLK: Find the parameter values (σ_ϵ , α , and β) that will minimize this,

$$-\text{llk} = T \times \ln \sigma_\epsilon + \frac{1}{2\sigma_\epsilon^2} \sum_{t=1}^T (Y_t - \alpha - \beta X_t)^2$$

In the EWMA case, we use the computer to minimize -llk by varying λ . In the GARCH case, we use the computer to minimize -llk by varying a_0 , a_1 and a_2 . Here, we can actually do it by hand. Let's forget σ_ϵ for now, and focus on α and β . To minimize -llk is the same as finding α and β so that

$$\frac{\partial \text{llk}}{\partial \alpha} = \frac{\sum (Y_t - \alpha - \beta X_t)}{\sigma_\epsilon^2} = 0$$

and

$$\frac{\partial \text{llk}}{\partial \beta} = \frac{\sum [X_t (Y_t - \alpha - \beta X_t)]}{\sigma_\epsilon^2} = 0.$$

Solving for the optimal α and β then reduces to solving the above two equations. The first derivatives of llk with respect to the model parameters (e.g., α and β) are also called score. If there are two parameters, then the score is a vector of two. In our current case, the score should be a vector of three because there are three parameters: α , β , and σ_ϵ . But as agreed, let's focus only on α and β and forget about σ_ϵ .

Solving for the partial derivative (score) with respect to α , we have,

$$\alpha = \frac{1}{T} \sum Y_t - \beta \frac{1}{T} \sum X_t$$

Solving for the partial derivative with respect to β , we have

$$\sum X_t Y_t - \alpha \sum X_t - \beta \sum X_t^2 = 0$$

Plugging the solution for α into the equation above, we have:

$$\sum X_t Y_t - \frac{1}{T} \sum X_t \sum Y_t = \beta \left(\sum X_t^2 - \frac{1}{T} \left(\sum X_t \right)^2 \right)$$

Let me divide both sides of the equation by T so that you can see the result more clearly,

$$\frac{1}{T} \sum X_t Y_t - \left(\frac{1}{T} \sum X_t \right) \left(\frac{1}{T} \sum Y_t \right) = \beta \left(\frac{1}{T} \sum X_t^2 - \left(\frac{1}{T} \sum X_t \right)^2 \right)$$

What we have is,

$$\text{cov}(X, Y) = \beta \text{var}(X)$$

So, as a by product of our derivation, you get to know why running a regression gets you the CAPM beta.

Also, for those of you who think more carefully, the fact that we assume ϵ is normally distributed might be bothering you. Don't be. Even if ϵ is not normally distributed, we can still do this procedure, which is called quasi-maximum likelihood estimation. The estimates might not be the most efficient, but they are still consistent. One key calculation left out is how to calculate the standard errors of α and β . For those of you who are interested in learning more, I would recommend the chapters on GMM and MLE of James Hamilton's book on Time Series Analysis.