Investments

Masters in Finance

NOVA SCHOOL OF BUSINESS & ECONOMICS

Spring 2025, Martijn Boons

Book chapter 5

1. Multi-period returns

Multi-period returns: definition and example

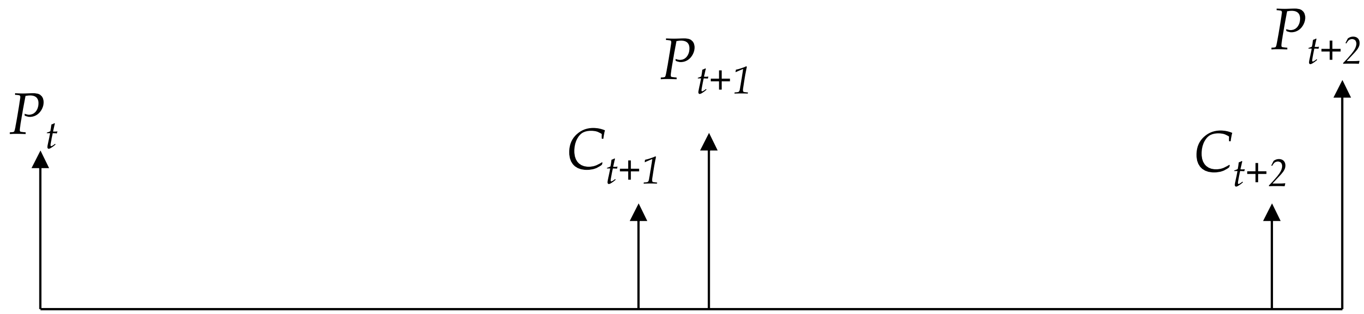

- Return from holding asset from *t* to *t*+2, assuming we <u>reinvest</u> capital gain and intermediate income back into the asset:

$$r_{t,t+2} = \frac{C_{t+1} + P_{t+1}}{P_t} \times \frac{C_{t+2} + P_{t+2}}{P_{t+1}} - 1 = \left(1 + r_{t,t+1}\right) \times \left(1 + r_{t+1,t+2}\right) - 1$$

- Example: you earn 10% in year 1 and 20% in year 2;

  2-year return = $(1 + 0.1) \times (1 + 0.2) - 1 = 1.32 - 1 = 32\%$

  ≠ sum of returns (30% = return without reinvesting)

- This is what most funds do in the real world, but other reinvestment strategies are possible
  - Not reinvesting intermediate income at all (i.e., consume it)
  - Reinvest intermediate income at the risk-free rate (i.e., put it on a bank account)

Expected rates of return in multi-period setting

NOVA SCHOOL OF BUSINESS & ECONOMICS

- Suppose that risky returns are identically and independently distributed (IID) over time, with
  - μ the per period expected return = the average return over possible states *s:* $\mu = \sum_s p(s) r\,(s)$
  - IID is a reasonable assumption given market efficiency.

- Total expected returns over *n* periods:

$$E\left[r_{0,n}\right] = E\left[\left(1 + r_{0,1}\right)\left(1 + r_{1,2}\right) \ldots \left(1 + r_{n-1,n}\right)\right] - 1 =$$
$$E\left[\left(1 + r_{0,1}\right)\right] E\left[\left(1 + r_{1,2}\right)\right] \ldots E\left[\left(1 + r_{n-1,n}\right)\right] - 1 =$$
$$(1 + \mu)^n - 1$$

Since E(XY) = E(X)E(Y)+Cov(X,Y), but Cov(X,Y)=0 when X,Y are independent

Volatility and time

- Suppose that risky returns are identically and independently distributed (IID) over time, with
  - σ^2 the per-period variance = an average over possible states *s:* $\sigma^2 = \sum_s p(s)\,[r(s) - \mu]^2$
  - σ the per-period standard deviation, expressed in same percentage units as returns

- If returns are relatively small, then over *n* periods we have

$$r_{0,n} = \left(1 + r_{0,1}\right)\left(1 + r_{1,2}\right) \ldots \left(1 + r_{n-1,n}\right) - 1 \approx r_{0,1} + r_{1,2} + \cdots + r_{n-1,n}$$

such that

Since Var(X+Y)=Var(X)+Var(Y)+2Cov(X,Y) and Cov(X,Y)=0

$$Var\left(r_{0,n}\right) \approx Var\left(r_{0,1} + r_{1,2} + \cdots + r_{n-1,n}\right) = Var\left(r_{0,1}\right) + Var\left(r_{1,2}\right) + \cdots + Var\left(r_{n-1,n}\right) = n \times \sigma^2$$

- This means that standard deviation increases with the square root of time:

$$\sigma_{Annual} = \sqrt{12}\sigma_{Monthly}$$

Estimates of risk and return

- True means and variances are unobservable, because the states s are unknown, so we estimate them.
- Historical arithmetic average estimated using a sample of *N* returns:

$$\bar{r} = \frac{1}{N}\sum_{n=1}^{N} r_{n-1,n}$$

- Arithmetic average $\bar{r}$ is our first estimate of true expected return μ
  - ➢ We will later find forward-looking estimates using asset pricing models
  - ➢ Similarly, $\hat{\sigma}^2 = \frac{1}{N}\sum_{n=1}^{N}(r_{n-1,n} - \bar{r})^2$ estimates the variance

- Historical average return is a poor proxy for expected future return
  - Poorly estimated and sensitive to outliers

  → Disclaimer: past performance is no guarantee for future performance!

- Expected returns compensate for risk, which is estimated more precisely based on historical data (and especially using high-frequency data).

Geometric average return

- The geometric average tells us, for a sequence of returns, what is the constant per-period return that would have generated the same end-of-period wealth

$$r_g = \left(\prod_{n=1}^{N} \left(1 + r_{n-1,n}\right) \right)^{1/N} - 1 = \left(1 + r_{0,N}\right)^{1/N} - 1$$

where $r_{0,N}$ is the total return from time 0 to time N.

- Geometric mean ≈ Arithmetic mean – ½*Variance

Variance reduction return

- There is an important lesson in understanding the difference between arithmetic and geometric returns.
  - Which two consecutive returns do you prefer?
    A. 90%,-50% or
    B. 10%,-10%

  - Lower volatility of B more than makes up for the large difference in (arithmetic) average returns (A. 20% vs. B. 0%).

- Lowering return volatility is advantageous for long-term investors.

  - We will see later that diversification reduces portfolio volatility without affecting arithmetic average returns, thus leading to higher geometric average returns.

Tesla

- Calculate average returns and risk for an investment in Tesla stock and compare these numbers to the S&P500

- What do these numbers tell you about the average return and risk to expect from investing in Tesla going forward?

  - Average return:
    - Most probably little:
      - Historically, few stocks have repeated such a 15-year track record.
      - Price may be high due to (i) overpricing, (ii) low discount rates or (iii) high expected growth and all three would imply exceptional returns are unlikely to repeat in the future.

  - Risk:
    - Quite a bit: 60% per year isn't unheard of looking at other individual stocks.

2. Historical performance

Historical average returns (1928-2013, US)

1928-2013	Geometric Average	Arithmetic Average
Stocks	9.6%	11.5%
T-bills	3.5%	3.6%
Long-term Government Bonds	4.9%	5.2%
Inflation	3.1%	3.1%

Source: Ibbotson, CRSP

- Example series in Excel.

- If we invest $1 now in each of these asset classes, by how much can we expect it to grow after 30 years <u>if future returns are like the past</u>?
  - Use geometric averages
  - E.g., (1+0.096)^30 = 15.6 for investment in stocks

- Equity premium = expected excess stock return minus risk-free (T-bills) =
  - 7.9% (typically reported using arithmetic averages)
  - This is a real return!

Risk

1928-2013	Standard Deviation
Stocks	20.0%
T-bills	3.1%
Long-term Government Bonds	7.9%
Inflation	4.1%

Source: Ibbotson, CRSP

- Risk is uncertainty about the future
  - While equity premium is positive on average, investors know that stocks can strongly underperform in any given year.
  - By how much?

The Normal distribution

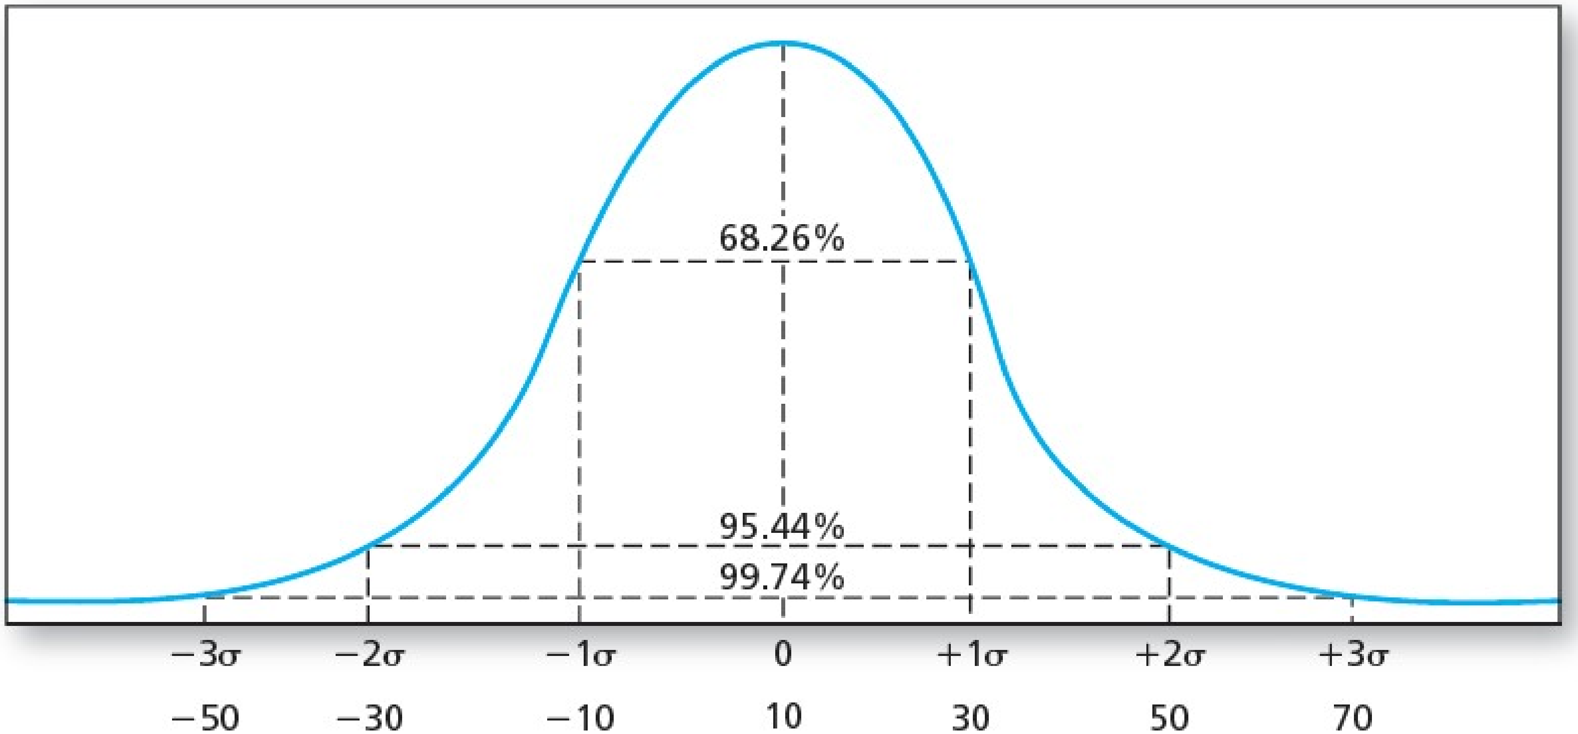

Assuming stock returns are normally distributed with mean 10% and standard deviation 20%, we would expect a return below x in y% of the observations:

x	z-score = $\frac{x-mean}{st.dev.}$	y
10%	0.0	50.00%
0%	-0.5	30.85%
-10%	-1.0	15.87%
-30%	-2.0	2.28%
-50%	-3.0	0.13%

Sharpe ratio

- Sharpe ratio measures the tradeoff between expected return and risk
  - How much expected excess return can be obtained per unit of standard deviation

$$SR = \frac{\mu - r_f}{\sigma}$$

- To estimate the Sharpe ratio, replace with sample moments

$$\widehat{SR} = \frac{\bar{r} - r_f}{\hat{\sigma}}$$

- In our sample, in annual terms, $\widehat{SR} = 0.4 = \frac{11.5\% - 3.6\%}{20\%}$.

- Investment managers talk a lot about SR of their portfolio and anything higher than 1 should likely be taken with a grain of salt

The equity premium puzzle (1/3)

- The equity premium is one of the most important numbers in financial economics
  - At what age can you retire? How much will you have? How much should you save as part of your salary?
  - Podcast: All else equal – How risky should your retirement savings be?
  - For reasonable levels of risk aversion and given that you are relatively young, popular life-cycle investment models imply you should invest >>50% in equities
  - The puzzle: this is not at all what we see in reality…

- Problem 1: Historical average realized returns are not necessarily a good estimate of expected future returns

  - Historical averages better for diversified portfolios than individual assets, but even for the market portfolio:
  - large uncertainty. With 100+ years of data
    - standard error on mean estimate is 2% (≈20%/$\sqrt{2024-1928}$)
    - Two-standard error (95%) confidence interval is

      [7.9-4%,7.9+4%]=[3.9%,11.9%]
    - What if there was a structural break? Markets and investors' portfolios surely more diversified now than in 1930.

The equity premium puzzle (2/3)

- Problem 2: Survivorship bias
  - US economy was historically very successful, and the US equity market followed suit:

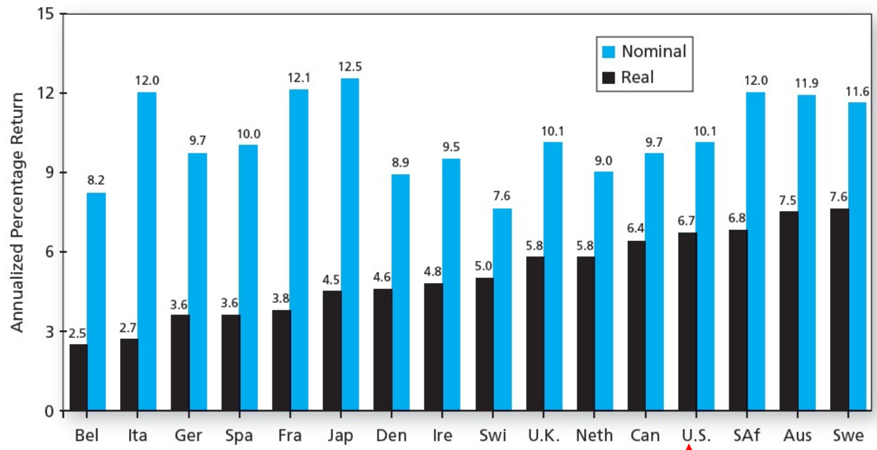

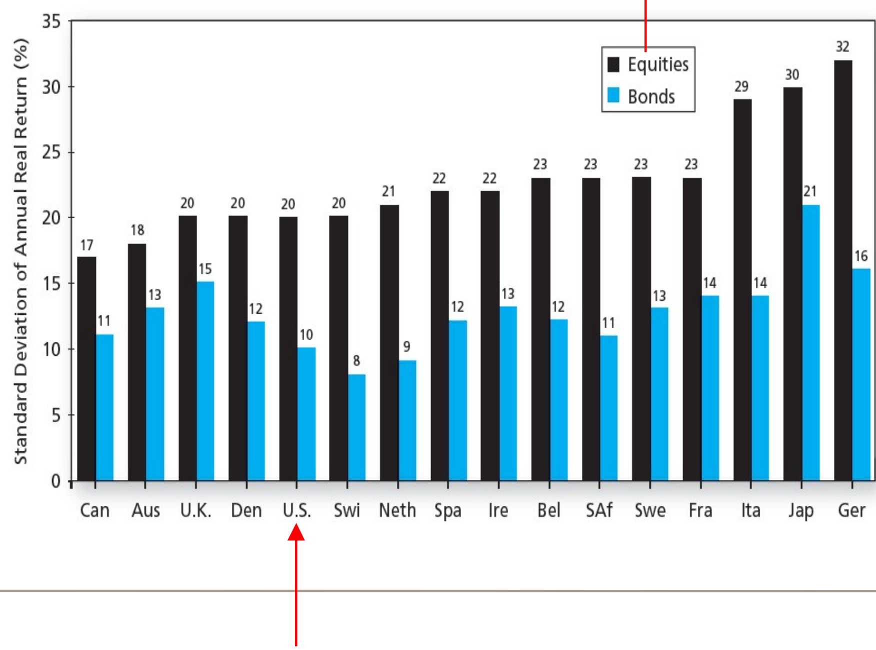

The equity premium puzzle (3/3)

- Problem 3: Alternative approaches to measure the equity premium based on fundamental stock valuation yield numbers around 4%
  - Recall: $k = D/P + g$ → Expected return = dividend yield + expected growth rate of dividends
  - Claus and Thomas (2001) use analyst forecasts of dividends (and growth rates)
  - Fama and French (2002) use alternative measures of expected dividend growth icw current D/P
  - 4% is closer to the median across countries
- Intuition (see) : Success of US economy has led to
  - discount rate shocks that were negative on average
  - crashes that did not happen, but were rationally expected ex ante
  - Both lead to an upward bias in realized returns.

3. Non-normal returns

Normal returns

- Investment management is easier when returns are (jointly) normal
  - Standard deviation is a sufficient measure of risk when returns are symmetric
  - The dependence of returns across securities can be summarized using only the pairwise correlation coefficients
  - Scenario analysis is easy

Non-Normal returns

- What if returns are not normally distributed?

  - Variance (k=2) is no longer a complete measure of risk as higher moments matter, with k-th moment a function of $E((r-\mu)^k)$
    - Post-war US stock market return:
      - Skewness (k=3): -0.5
        - Left-tail relatively fatter than right-tail
      - Kurtosis (k=4): 1.8
        - Both tails fatter than normal
  - Measuring (tail) dependence is tricky when dealing with many assets

- The normal assumption isn't too bad and that's why it's popular.
  - Keep higher moments in mind when looking at real-world data though!

Alternatives to standard deviation

- Value at Risk (VaR)
  - Loss corresponding to a very low percentile of the entire return distribution, such as the fifth or first percentile return
  - Commonly used in banking regulation
  - **Why is VaR redundant for normally distributed returns?**
- Expected Shortfall (ES)
  - Also called conditional tail expectation (CTE), focuses on the *expected* loss in the worst-case scenarios (left tail of the distribution)
  - More conservative measure of downside risk than VaR

VaR and ES

Example (1):
Consider a return that is distributed uniformly in the interval [-20%,+30%]

- Uniform distribution: each outcome in the interval is equally likely

- Compute the VaR(5%) and the ES(5%) for a portfolio of $10M

  - VaR(5%) = the outcome for which we estimate that things will only be worse in 5% of cases
  - ES(5%) = the average outcome in the worst 5% of cases

Example (2):
Calculate the same statistics for monthly US stock market returns since the 1960s and compare to what these statistics would look like if US stock market returns were normally distributed.

4. Time-varying distributions of returns

Stylized facts

- Across different asset classes, time periods, geographies, and so on, return distributions seem to be predictably time-varying.

  – **Difficult question: How to distinguish non-normality from time-variation in μ and σ?**

- We will discuss in this course a number of timing strategies. To set the stage, let us consider the equity premium:

- Talking to practitioners, you will realize that many of them understand that the equity premium is time-varying.
  – E.g., many fund managers advocate buying equities after a crash.
  – Why?
    - Recall the value of a stock that pays a constant dividend: P = D/k
    - Rewriting, we see that a proxy for the expected return of a stock k = D/P
    - So, when prices are low after a crisis, it means that expected returns are high.
    - This is standard “Present Value” logic: prices and discount rates are negatively correlated

The equity premium

- Regressing monthly US equity returns on lagged dividend/price from 1962 to 2014:

$$R_{MKT,t+1} = a + \underbrace{b}_{(0.33,\ t-sta\ \approx 2)} DP_t + e_{t+1} \ (R^2<0.006)$$

- Monthly excess equity market returns are predictable in-sample.
- Evidence not strong: practitioners require estimate today of conditionally expected/forward-looking return, which is hard-to-estimate:
  - Before 2014, nobody knew that 0.33 x DP_t, provides a good estimate of next month's equity premium.
    - In-sample R^2 small → out-of-sample R^2 even smaller.
      - Out-of-sample R^2: Prediction at each point in time uses only historical data available at that point in time!
    - Coefficient (b) larger in 1st and 2nd sample half, suggesting structural breaks.
- Thus, trading on time-varying returns is not easy!
- What we do learn from this simple regression: expected stock market returns are time-varying.
  - To see this even better, consider 10 year returns:

The long-term equity premium and DY

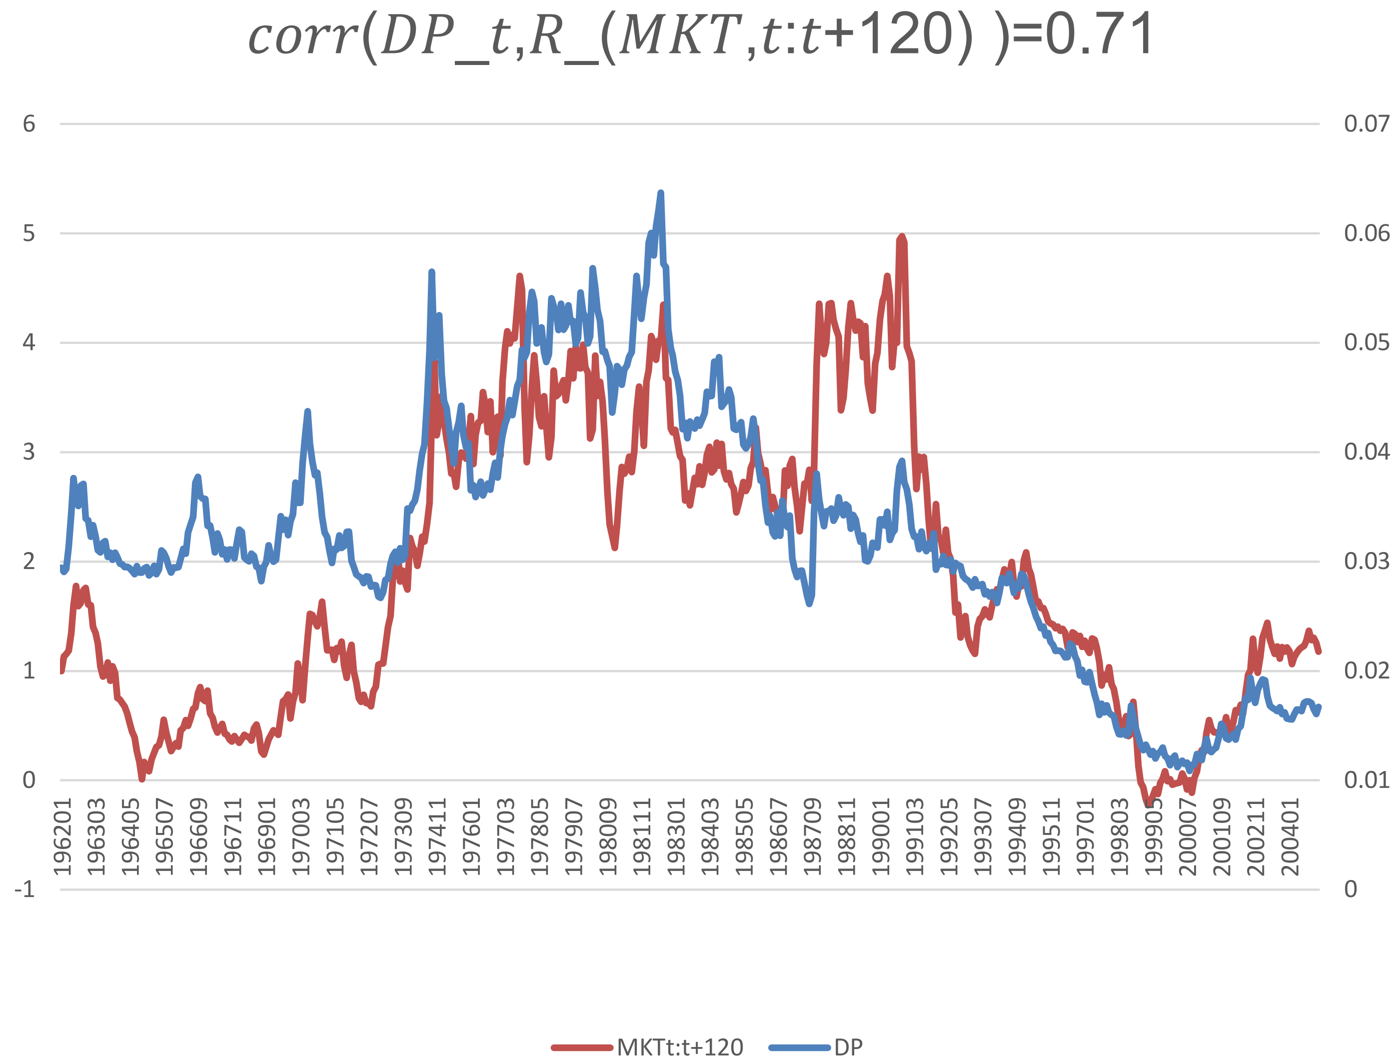

Question: What should investors do with this information?

Volatility is also time-varying

- While predicting returns is difficult, predicting volatility is much less so:
  - ➢ No autoregressive component in returns, but clearly in squared returns.
  - ➢ Financial econometrics has had great success in coming up with time-series auto-regressive type models used to filter out conditional variance ($Var_t(R_{t+1}) = \sigma_t^2$)
    - ➢ Key ingredients of models: volatility is persistent & mean-reverting.

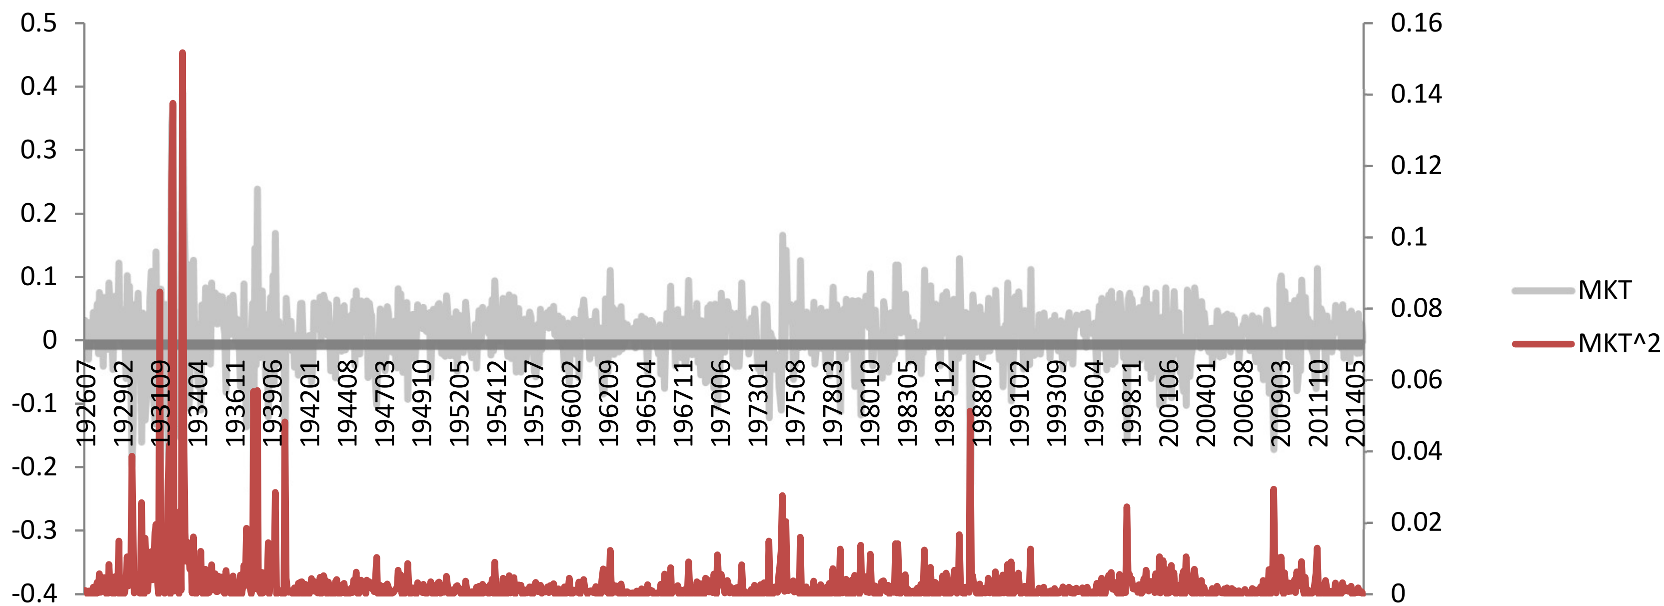

5. How financial markets work

The VIX

- The volatility index, or VIX, is the volatility implied by option prices.
  - VIX measures the expected volatility for the underlying up to expiration of the option.
  - Sophisticated methods to back out market's expected volatility from current option prices (see FT Alphaville article).
  - The VIX jumps whenever market turmoil hits and is seen by investors and popular press as an index of fear.

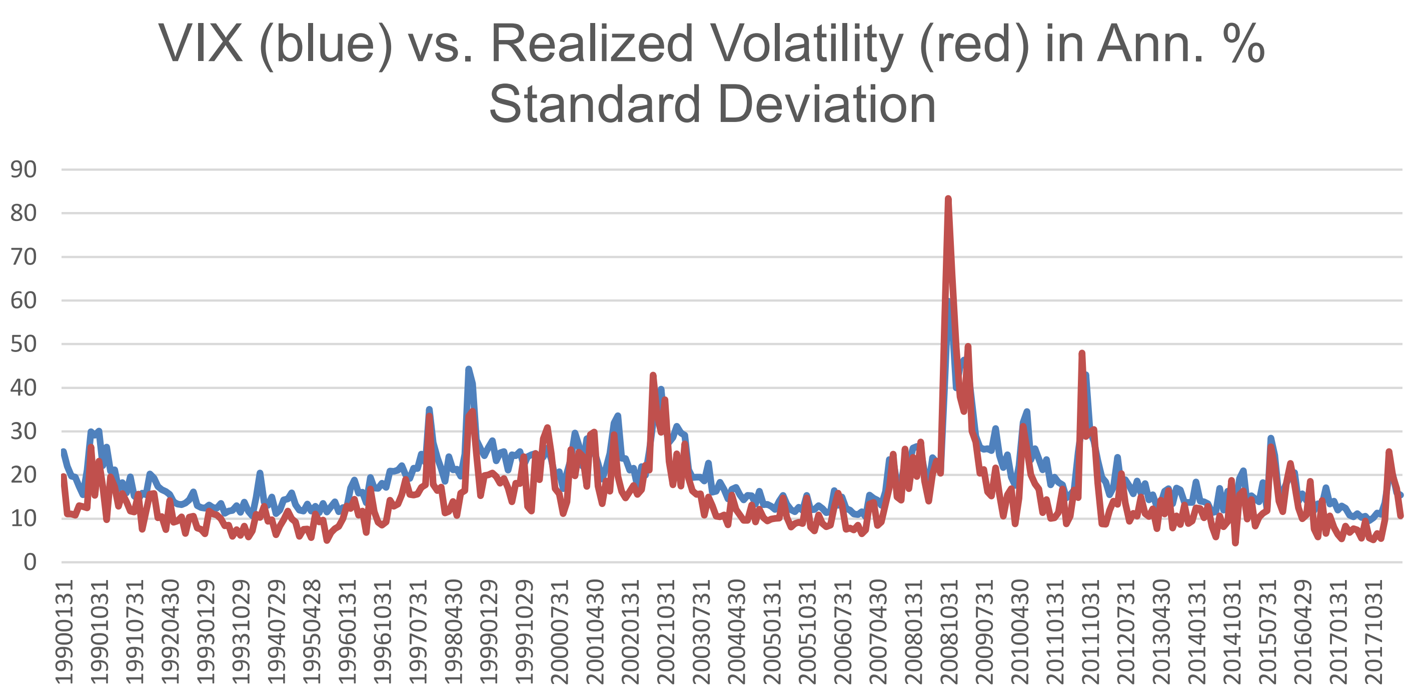

The volatility risk premium

- On average, option implied volatility is 4% above actual realized volatility. This represents a 4% **volatility risk premium**. Why!?
  - Investors are risk averse and dislike increases in volatility so much that they are willing to “overpay” for options to hedge (relative to risk-neutral investors).
    - Option prices are increasing in volatility.
  - This premium is exactly why there are other, less risk averse investors that are willing to sell options (think: hedge funds). They take on this risk and get compensated by a high expected return.
    - However, when volatility shoots up, these hedge funds incur great losses (see FT article)

What do we learn from this example?

1. If something is risky, it will carry a price or "risk premium"
   - Investors dislike volatility and are willing to pay a premium to hedge against increases in volatility.
   - Throughout the course, we take great care identifying what is risk.
     - Risk is not only volatility or poor market returns. Risk is any factor that determines bad times for investors: poor economic growth, a real estate crash, inflation, illiquidity, and so on.

2. Financial markets match investors
   - Investors that want to hedge volatility and those that are willing to provide this hedge at a certain price.
   - This matching has become increasingly efficient due faster trading platforms as well as new contracts (e.g., VIX futures).

3. Financial markets produce information
   - Market expectations of volatility can be estimated from option prices.