GFLEC
GLOBAL FINANCIAL LITERACY
EXCELLENCE CENTER

Basic Probability & Risk vs. Return

Lecture 9

Risk, uncertainty, and probability

There is a considerable amount of **uncertainty** in finance and investing.

For example, while planning for retirement, it's impossible to know exactly what returns stocks and bonds will yield in the future.

In the first half of this lecture, we discuss **probability** to help understand decision-making under uncertainty.

Probability

The probability, P(E), of an event E occurring can be calculated as:

$$P(E) = \frac{|Outcomes(E)|}{|Outcomes|}$$

Where |Outcomes(E)| is the number of outcomes in which E occurs, and |Outcomes| is the total number of outcomes.

This is best illustrated through examples.

Urns

Ex. An urn contains four blue balls and five red balls. What is the probability that a ball randomly chosen from the urn is blue?

Ans.

There are nine possible outcomes – one for each ball in the urn.

Four of these outcomes involve picking a blue ball.

Thus the probability of picking a blue ball is 4/9 = 0.444 = 44.4%.

Dice

Ex. What is the probability that when two dice are rolled, the sum of the numbers on the two dice is 7?

Ans. There are 6 possible numbers for the first die and 6 possible numbers for the second. When both are rolled, any possible number from the first may be paired with any number from the second, and so there are a total of 6*6 = 36 outcomes.

Of these six outcomes, the following pairs add to seven:

{(1,6), (2,5), (3,4), (4,3), (5,2), (6,1)}

Where the first number refers to the value on the first die and the second to the value on the second.

Thus, there are 6 such outcomes, and the probability of rolling a sum of 7 is 6/36 1/6 = 0.167 = 16.7%.

Populations

Ex. The following table presents hypothetical employment numbers for a population, in thousands:

Gender	Age	Employed	Unemployed	Total
Male	Under 18	100	25	125
	18 or older	700	35	735
Female	Under 18	110	20	130
	18 or older	850	65	915
	Total:	1,760	145	1905

What is the probability that a randomly selected male is unemployed?

What is the probability that a randomly selected male over 18 is unemployed?

What is the probability that a randomly selected unemployed person is a female?

What is the probability that a randomly selected person is under 18?

Populations

Gender	Age	Employed	Unemployed	Total
Male	Under 18	100	25	125
	18 or older	700	35	735
Female	Under 18	110	20	130
	18 or older	850	65	915
	Total:	1,760	145	1905

Ans.

The probability that a randomly selected male is unemployed is:

$$\frac{25 + 35}{125 + 735} = \frac{60}{860} = 7.0\%$$

The probability that a randomly selected male over 18 is unemployed is:

$$\frac{35}{735} = 4.8\%$$

Populations

Gender	Age	Employed	Unemployed	Total
Male	Under 18	100	25	125
	18 or older	700	35	735
Female	Under 18	110	20	130
	18 or older	850	65	915
	Total:	1,760	145	1905

Ans. (continued)

The probability that a randomly selected unemployed person is female is:

$$\frac{20+65}{145} = \frac{85}{145} = 58.6\%$$

And the probability that a randomly selected person is under 18 is:

$$\frac{125+130}{1905} = \frac{255}{1905} = 13.4\%$$

The Law of Large Numbers

Law of large numbers

Probabilities are important because of the **law of large numbers**:

- The law of large numbers states that, **with a large number of trials, the proportion of time that an event occurs is unlikely to be far from the probability of that event occurring**.
- And the larger the number of trials, the less likely the proportion will deviate materially from the probability.

Ex. A flipped coin has a 50% probability of landing heads. Therefore, when the coin is flipped multiple times, on average the coin will come up heads half the time. The **law of large numbers** states that, with a large number of trials, the proportion of heads is unlikely to deviate from 0.50.

So, although flipping 7 heads in 10 coin flips (for a proportion of 70% heads) is not rare, if the coin is instead flipped 1,000 times, it would be very rare to observe the same proportion of 70% heads with 700 heads.

Law of large numbers

The following chart demonstrates the law of large numbers. Each of the four lines depicts the proportion of heads observed during separate runs of 1,000 coin flips.

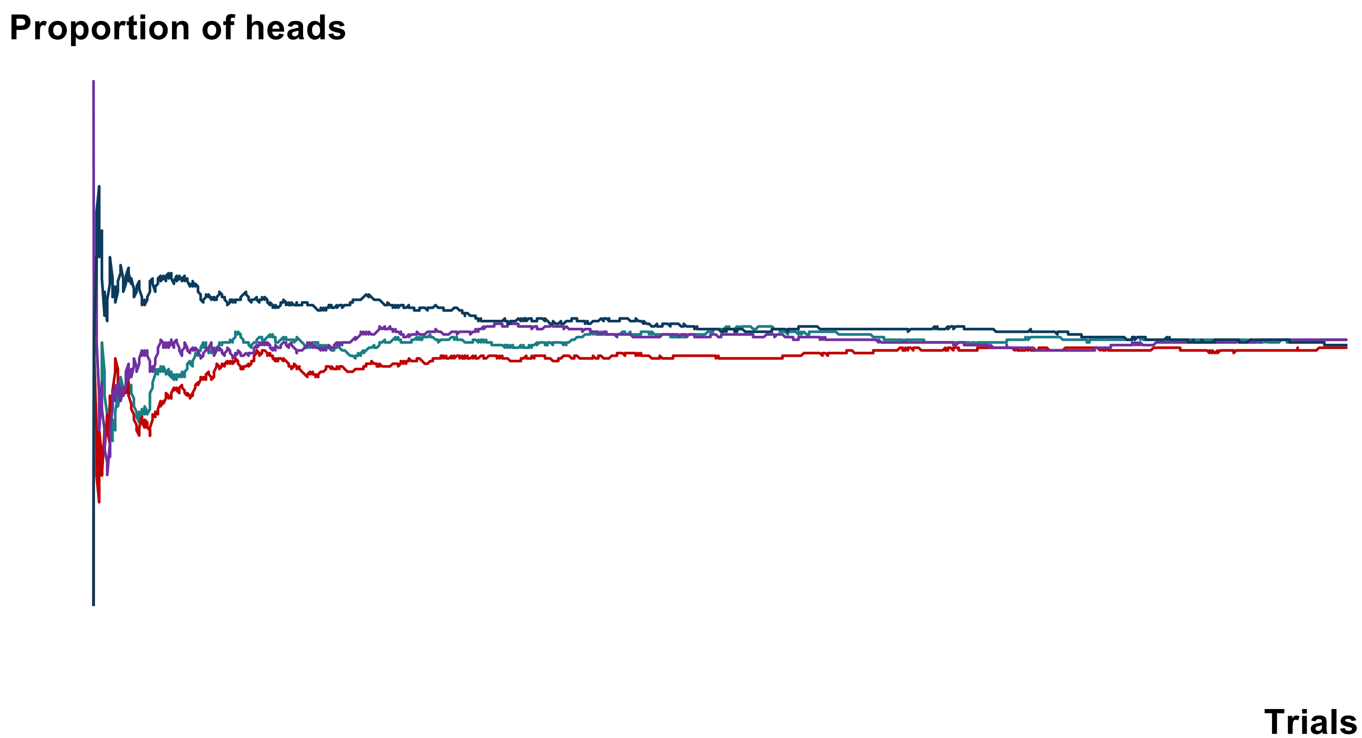

In each case, the proportion of heads converges towards 50% as the number of trials approaches 1,000.

The gambler’s fallacy

A common misinterpretation of the law of large numbers occurs when an unlikely string of events occurs and an observer believes that the event must be less likely in the subsequent outcome or outcomes to compensate and to push the average occurrences to the long-run mean predicted by the law of large numbers.

For example, if six coin straight coin flips come up heads, some might think that a tails must be more likely in the next few coin flips to push the number of heads to 50%.

This belief is incorrect and is known as the **gambler’s fallacy**. In fact, the event is just as likely in the subsequent outcome(s).

This is not a contradiction of the law of large numbers, as demonstrated by the following example…

The gambler’s fallacy

Ex. A coin is flipped 10 times and 8 of the flips are heads.

What is the probability that the next flip is a heads?

The coin is then flipped another 990 times, for a total of 1,000 flips. Of these 990 additional flips, how many are expected to be heads?

Given the 8 initial heads and the expected number of heads in the additional 990 coin flips, what is the total expected number of heads over the 1,000 coin flips, and what is the expected proportion of heads over these 1,000 coin flips?

Does this contradict the law of large numbers?

The gambler's fallacy

Ans.

The probability that the next flip is heads is 50% (it is **not** influenced by the previous string of heads). They are **independent** of one another.

Of the next 990 flips, 50%, or 495, are expected to be heads.

Thus, of the 1,000 total coin flips, 495+8 = 503 are expected to be heads.

This corresponds to a proportion of 0.503, which is close to the 0.50 probability of flipping a heads predicted by the law of large numbers.

Note that, after starting at a proportion of 0.80 heads after the first 10 flips, it was not necessary for the subsequent flips to have a lower probability of heads for the proportion to converge towards 0.50.

Expected Value

Expected value

When values are assigned to each outcome of a random variable, the **expected value** of the random variable is calculated as:

$$E[X] = p_1X_1 + p_2X_2 + \cdots + p_NX_N$$

Where *X* is a random variable with *N* possible outcomes, *E[X]* is the expected value of *X*, X_i is the value assigned to the i^{th} outcome of *X* and p_i is the probability of the i^{th} outcome occuring.

This is an abstract description of expected value and the concept is better illustrated by example…

Expected value

Ex 1. If the value 3 is assigned to a heads in a coin flip, and 1 is assigned to a tails, what is the expected value of this coin flip?

Ans. $E[X] = (1/2)*3 + (1/2)*1 = 1.5 + 0.5 = 2$

Ex 2. If the value on each face of a die is assigned to that face of the die, what is the expected value of a die roll?

Ans. $E[X] = (1/6)*1 + (1/6)*2 + (1/6)*3 + (1/6)*4 + (1/6)*5 + (1/6)*6 = 3.5$

Expected value

The importance of expected value is that, as a corollary of the **law of large numbers**, with a large number of trials, the **average value of the outcomes will tend toward the expected value**.

A popular application of expected value is to analyze bets:

Ex. If a coin comes up heads, a bettor wins $3. If it comes up tails, the bettor wins nothing. What is the expected value of this bet? Would you pay $1 to make this bet? Would you pay $2?

Ans. The expected value is (1/2)*$3 + (1/2)*0 = $1.50.

Thus, on average, a bettor can expect to receive $1.50 per bet on such a bet.

If the cost is $1, the bettor would make $0.50 on average, and should take the bet. If the cost is $2, the bettor will lose about $0.50 per bet in the long run, and so should avoid making a habit of such bets.

Expected value

The following chart demonstrates expected value. Each of the four lines depicts the average winnings for the bet analyzed in the previous example during separate runs of 1,000 coin flips.

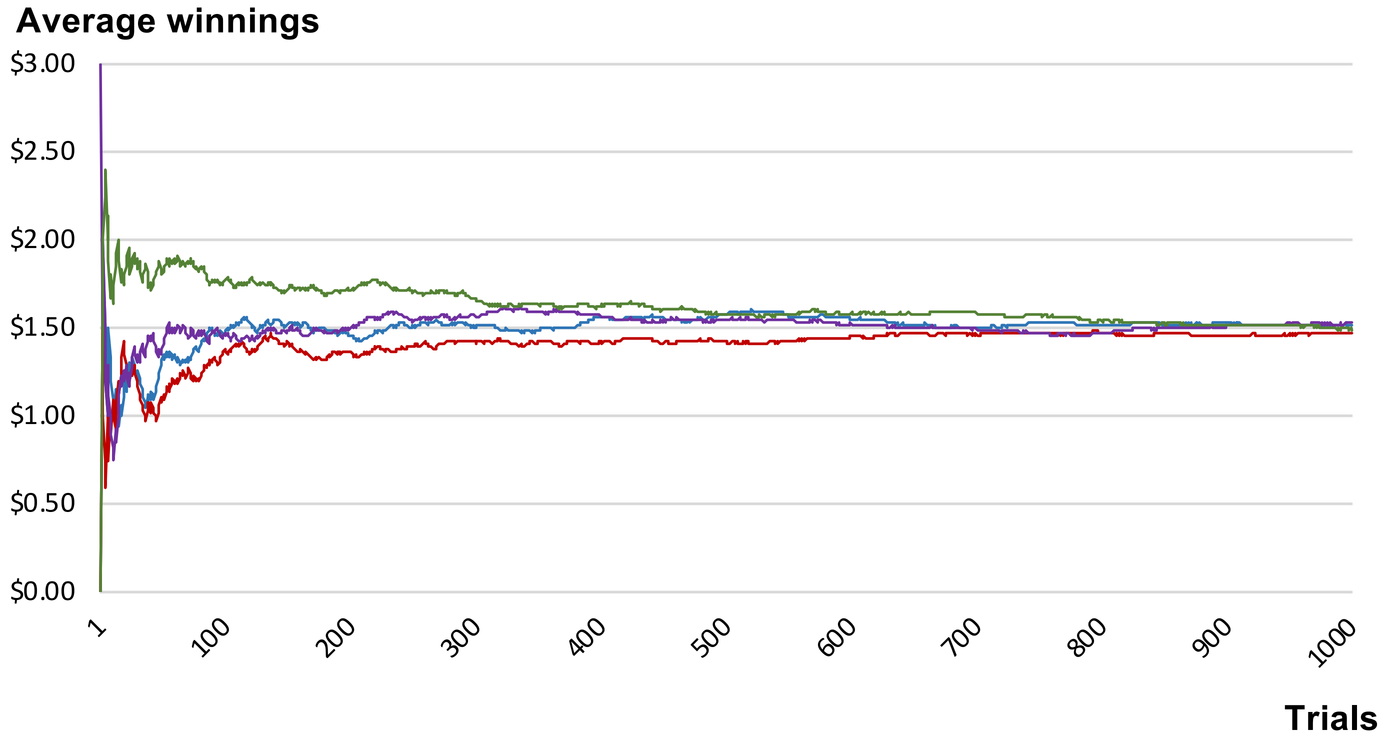

In each case, the average winnings converges to $1.50 as the number of trials approaches 1,000.

Lotteries

Lotteries are well-suited for evaluation by expected value. Lotteries tend to cost little and give bettors an enormous potential payout, but also a small probability of winning.

Ex. A lottery costs $1 to enter. The jackpot is $50,000.

To play, the bettor picks a sequence of 5 digits, each 1-9. If that number matches a randomly chosen string of 5 digits, that bettor wins the jackpot. If not, the bettor receives nothing.

What is the expected profit of the lottery for the bettor?

If 10,000,000 people play, how much money will the lottery holder win or lose, on average?

Lotteries

Ans.

There are 9*9*9*9*9 = 9^5 = 59,049 possible combinations of strings of 5 digits. Only one of these will match the bettors.

Thus, the probability of the bettor winning the lottery is 1/59,049, and the expected value of the lottery is (1/59,059)*$50,000, or about 85 cents.

Since the lottery costs $1, the expected profit to the bettor is -$0.15, and so the bettor will lose about 15 cents per play in the long-run.

The opposite outcome applies to the lottery holder, who will earn about $0.15 per bet, or about $1,500,000 if 10,000,000 bettors play.

Note that with such a large number of participants, it's unlikely that the lottery holder will lose money because of the law of large numbers.

Expected return

The expected return on an investment (either business or financial) can be computed using the method of expected value.

Consider an investment that returns 8% in one year half of the time and 4% the other half of the time. If $100 is invested, the $100 will grow to become $108 half the time and $104 the other half of the time. Thus, the expected value of the investment after one year is:

$$\left(\frac{1}{2}\right) * \$108 + \left(\frac{1}{2}\right) * \$104 = \$106$$

This implies an expected return of:

$$E[R] = \frac{\$106}{\$100} - 1 = 6\%$$

Note that this expected return can be calculated directly using the returns:

$$E[R] = \left(\frac{1}{2}\right) * 8\% + \left(\frac{1}{2}\right) * 4\% = 6\%$$

Risk vs. Return

Risks in investing

There are two main sources of risk in investing that we will consider in this lecture.

- **Default risk** is the risk that a borrower will not repay a loan or a bond issuer will not make the bond payments.
- **Financial risk** is the risk that the price of an asset will fall and an investor will be forced to sell it at a loss.
- T-bills and bank saving accounts are exposed to neither default risk or financial risk. T-bills are considered **risk-free** because the U.S. government has never, and is not expected to, default on its obligations. Bank accounts are insured by the FDIC.
- Corporate bonds are exposed to default risk and stocks are exposed to financial risk.
- When the payouts from a financial asset are uncertain, that asset is risky. And the more uncertain the returns, the riskier the asset is.

Risk Aversion

Risk aversion

Consider the following options. Would you prefer if:

- I give you a guaranteed $10,000.
- I flip a coin. If it comes up tails, I give you $20,000. But if it comes up heads, I give you $0.

Typically, people prefer the first option of the two above.

If we play the second option enough times, I will give you $10,000 *on average*, but there is uncertainty. So we say the second option is **riskier**.

This tendency to prefer a certain payout to an uncertain payout with the same average payout is known as **risk aversion.**

The risk premium

Which of *these* options would you prefer:

- I gave you a guaranteed $10,000.
- I flip a coin. If it comes up heads, I still give you $0. But if it comes up tails, <u>I give you $30,000</u>.

Even someone who chose the first option for the last choice may be induced to choose the second option with these new choices:

- The second option is still uncertain (risky).
- But now, the second option yields a higher average payout than the first option. I will give you $15,000, on average.
- This higher average payout may be enough to compensate you for the risk.
- This premium needed to induce a person to choose a risky option over a riskless option is known as the **risk premium**.

Risk premium

Ex 1. A risk-averse investor has the option to either (a) invest in a security that will return either 8% or 6%, each with a 50% probability, or (b) invest in a security that will return either 12% or 2%, each with a 50% probability. Given that this investor is risk-averse, is it clear which option she will choose?

Ans.

The expected return of the first option is 0.5*8% + 0.5*6% = 7%. The expected return of the second option is 0.5*12% + 0.5*2% = 7%.

Both options yield the same expected return, but the range of returns is much wider in option (b), which is therefore a riskier option.

Thus, the risk-averse investor will invest in the less risky option, which is option (a).

Risk premium

Ex 2. The same investor is then given a third option to (c) invest in a security that will return either 16% or 2%, each with a probability of 50%. Is it still clear which option she will choose?

Ans.

The expected return of the new option is 0.5*16% + 0.5*2% = 9%.

This option is riskier than option (a), but also yields more.

If the additional 2% return is sufficient to compensate her for the additional risk, the investor will choose option (c). If not, she will choose option (a). This depends on how risk-averse she is, so her choice is not clear.

Risk and the discount rate

If the cash flows for an investment are uncertain, most investors will require a higher rate of return (because of the **risk premium**) and discount the cash flows using a **higher discount rate**.

- Investors have many different options for assets to invest in, from riskless savings accounts to risky stocks and bonds.
- If stocks, which have highly volatile returns, offered the same average return as a savings account, most investors would prefer a safe investment in a savings account.
- Similarly, if bonds, which sometimes default, offered the same interest rate as a savings account, no investor would buy bonds.
- Therefore, the prices of stocks and bonds must be low enough such that they offer high enough returns to compensate investors for these risks.
- Investors price the stocks and bonds with a higher discount rate to achieve their desired higher expected return.

Risk and the discount rate

If the cash flows for an investment are uncertain, an investor will require a higher rate of return (because of the **risk premium**) and discount the cash flows using a **higher discount rate**.

Ex 1. An investor has the option to invest in:

(a) a project that generates a certain \$50 a year for three years and costs \$100

(b) a project that generates either \$30 or \$90, each with 50% probability, each year for three years and also costs \$100.

The investor requires a 5% return for riskless investments like the first. If the second project is also discounted by 5%, which project will the investor choose?

If the investor is risk-averse, will he be satisfied with a 5% return on a risky project? If the investor instead requires a 12% expected return to compensate for the uncertainty in the cash flows of the second choice, which project will he choose? What if he requires an 18% expected return?

Risk and the discount rate

Ans.

The net present value of the first project is:

$$NPV_1 = \frac{\$50}{1.05} + \frac{\$50}{1.05^2} + \frac{\$50}{1.05^3} - \$100 = \$36.16$$

The expected annual cash flows for the second project are:

$$0.5 * \$30 + 0.5 * \$90 = \$60$$

If the investor is not risk-averse (and is **risk-neutral**), he will discount the second project at the same 5% rate used for the first and value the project at:

$$NPV_2 = \frac{\$60}{1.05} + \frac{\$60}{1.05^2} + \frac{\$60}{1.05^3} - \$100 = \$63.39$$

At a discount rate of 5% for both projects, the investor will prefer the second. However, most investors are risk-averse, and so would require a risk premium for the uncertain cash flows offered by project (b)…

Risk and the discount rate

Ans. (continued)

If the investor instead requires a 12% return on project (b), he will value it at:

$$NPV_2 = \frac{\$60}{1.12} + \frac{\$60}{1.12^2} + \frac{\$60}{1.12^3} - \$100 = \$44.11$$

The investor will still prefer the second project, even when it is discounted at a higher rate of 12%. This means that the additional \$10 per year is sufficient compensation for the added risk.

If the investor instead requires an 18% return for the riskier project, its value falls to:

$$NPV_2 = \frac{\$60}{1.18} + \frac{\$60}{1.18^2} + \frac{\$60}{1.18^3} - \$100 = \$30.46$$

In this case, he will prefer the first project. Even though the second offers a higher *expected* return, it is not sufficient to compensate him for the added risk.

The Credit Spread

Default risk and the credit spread

For bonds, cash flows are uncertain because the issuer may **default**, and investors require a higher return to compensate them for this risk.

- Because an issuer may default, the investor may receive less than the face value of the bond.
- The expected value of the bond is therefore less than the face value.
- Bonds are valued based on these expected cash flows, and so a decreased expected cash flow leads to a decreased price.
- The decreased price leads to a higher interest rate (*remember that bond prices are inversely related to interest rates*).
- This higher implied interest rate over the rate on an equivalent bond with little or no default risk is known as a **credit spread**.

Default risk and the credit spread

Ex.

A zero-coupon bond that makes a single principal repayment in one year with no default risk is discounted at 5% & is priced at:

$$P = \frac{\$100}{1.05} = \$95.24$$

If the bond has a default probability of 5% and the investor gets nothing if the bond defaults, the expected value of the principal repayment is:

0.05*$0 + 0.95*$100 = $95

If this is still discounted at 5%, the new price of the bond is:

$$P = \frac{\$95}{1.05} = \$90.48$$

Default risk and the credit spread

Ex. (continued)

The implied yield on this bond is 10.52%:

$$\$90.48 = \frac{\$100}{1+r} \rightarrow r = \frac{\$100}{\$90.48} - 1 = 10.52\%$$

In this case, the credit spread is 10.52% - 5% = 5.52%.

Note, however, that a ***risk-averse*** investor would discount the risky cash-flow at a higher discount rate than 5%, further depressing the bond price and increasing the credit spread.

For example, if the investor discounted the expected cash flows at 15%:

$$P = \frac{\$95}{1.15} = \$82.61 \rightarrow r = \frac{\$100}{82.61} - 1 = 21.05\%$$

In this case the credit spread is 16.05%!

Default risk and returns

The following chart shows how investors require higher returns from issuers more likely to default.

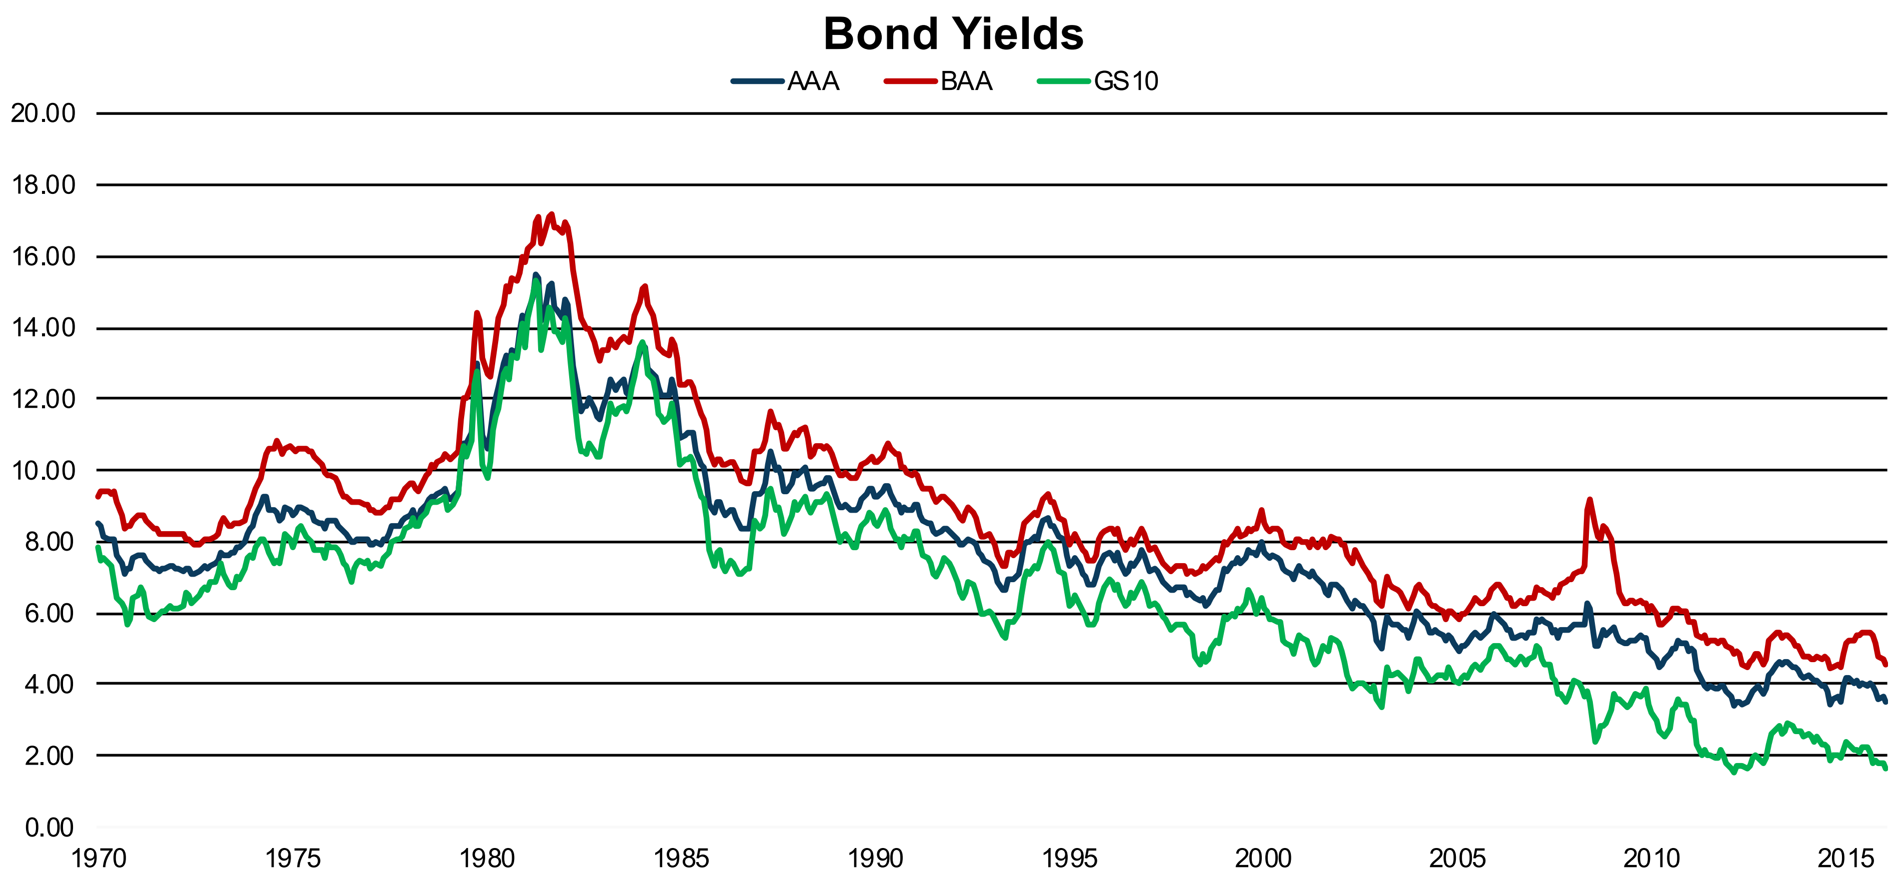

Source: Federal Reserve of St. Louis Economic Data (FRED)

Default risk and credit spreads

The credit spread for corporate issues can be found by subtracting the return on a riskless bond:

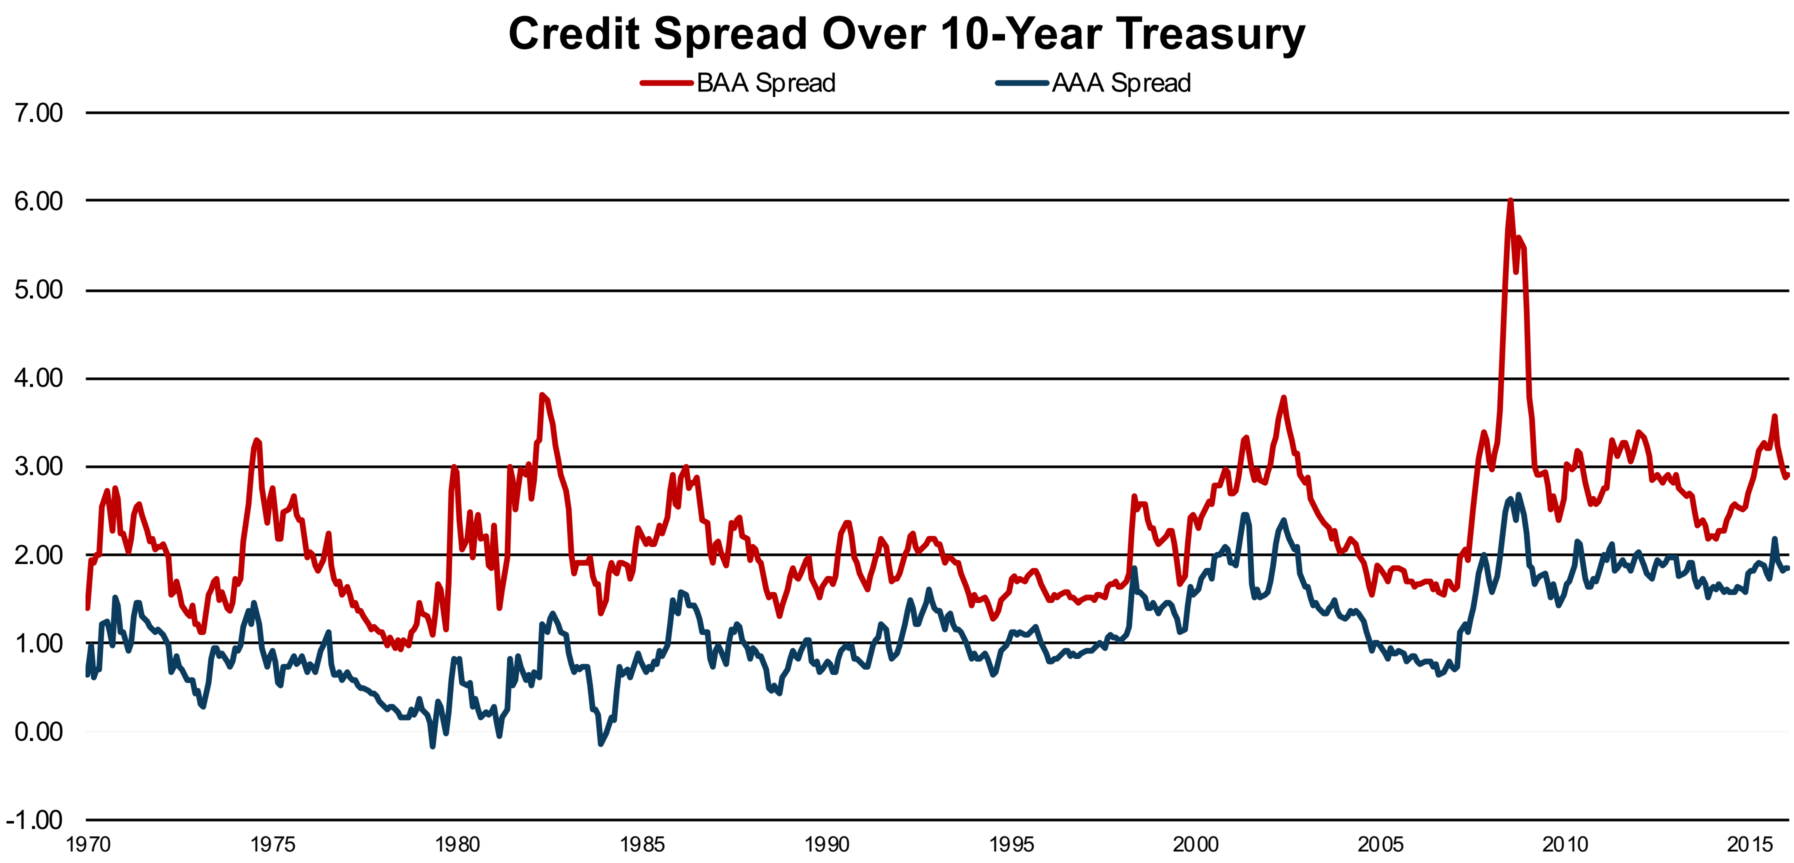

Source: Federal Reserve of St. Louis Economic Data (FRED)

Sovereign bond spreads

Credit spreads are frequently cited when evaluating the perceived financial health of countries.

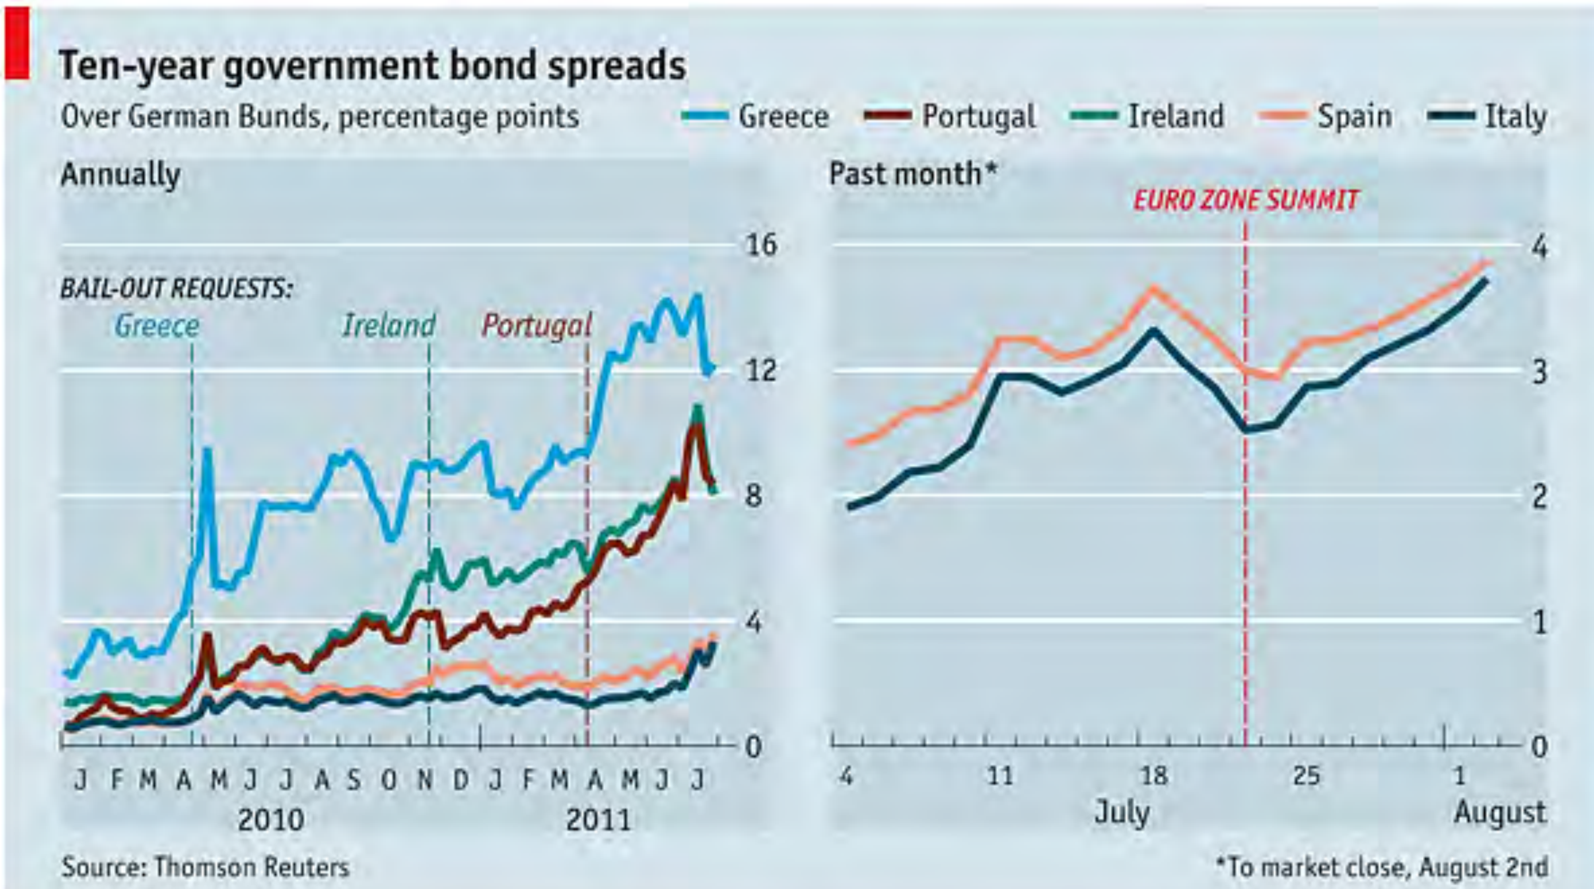

Consumer loans and the credit spread

The same credit spread principle applies to consumer loans: **the riskier the loan, the higher the interest rate**:

- If a borrower has a lower credit score (i.e. is statistically more likely to default), the bank charges a higher interest rate.
- Collateralized loans (ex. mortgage and auto loans) reduce the bank's loss even if the borrower defaults, which decreases the riskiness of the cash flows. This is why banks charge lower interest rates on collateralized loans.
- If a borrower makes a large down payment on a collateralized loan, this reduces the lender's chance of a loss, and so the lender charges a lower interest rate.
- Credit card loans are not collateralized, so lenders may lose the entire lent amount. Because such loans are so risky, lenders charge a higher interest rate.

Risk and Dispersion

Risk and dispersion

The more widely returns may fluctuate, the more risky they are said to be.

Consider the following three investments:

(A) 7% return with 100% probability
(B) 6% return with 50% probability, 8% return with 50% probability
(C) 2% return with 50% probability, 12% return with 50% probability

Each investment has an expected return of 7%.

However, while (A) yields a guaranteed 7% return, (B) and (C) may yield less. In other words (B) and (C) are risky.

And because (C) may yield much less while (B) may only yield a small amount less, (C) is considered the riskiest.

Risk and dispersion

The more widely returns may fluctuate, the more **disperse** the possible returns are said to be.

In the example above, the investment with the highest **dispersion** is (C). The chart below compares the dispersion among the three investments:

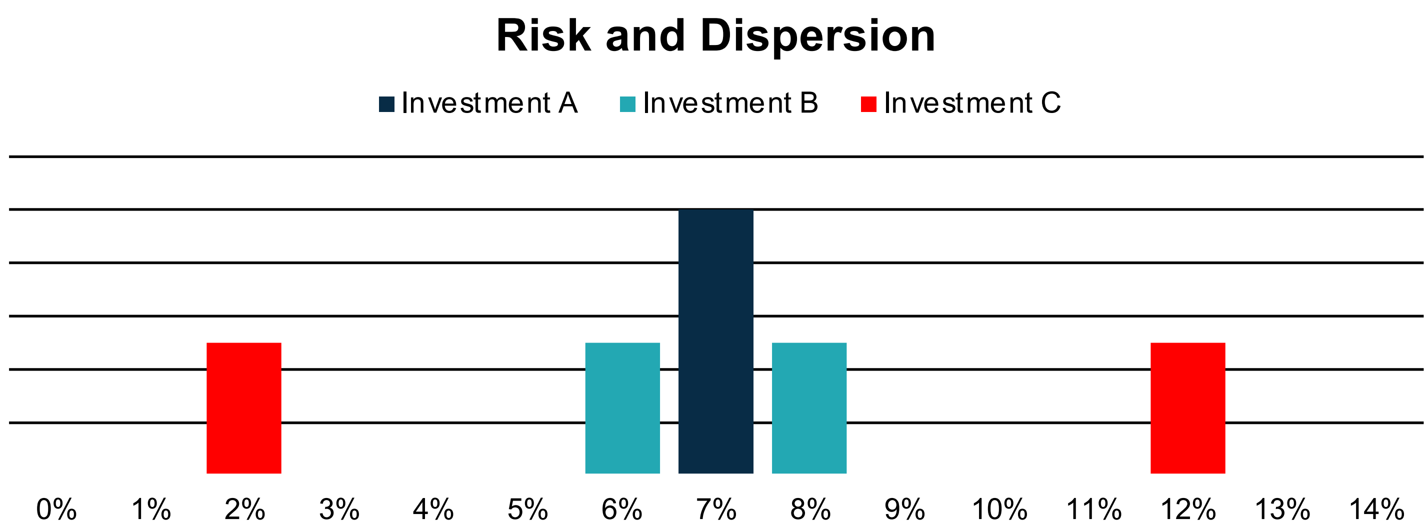

Measuring risk: the standard deviation

One formal, and popular, measure of dispersion is the **standard deviation**. A higher standard deviation implies a wider dispersion, and hence a higher risk.

Calculating the standard deviation is not necessary for this course, but for the mathematically curious, it is calculated as:

$$SD[X] = \sqrt{p_1(X_1 - \bar{X})^2 + \cdots + p_N(X_N - \bar{X})^2}$$

Where $\bar{X}$ is the expected value of X.

Measuring risk: the standard deviation

Consider again the following three investments:

(A) 7% return with 100% probability
(B) 6% return with 50% probability, 8% return with 50% probability
(C) 2% return with 50% probability, 12% return with 50% probability

Investment (A) has a standard deviation of 0% because its returns are not at all disperse and it has zero risk.

Investment (B) has a standard deviation of 1%. The possible returns on investment (B) are not very disperse and the investment is only mildly risk.

Investment (C) has a standard deviation of 5%. It has the highest standard deviation because it has the most widely dispersed returns and is the riskiest.

Risk vs. return on different assets

The following table lists the historical nominal returns and standard deviations for some common asset classes from 1926-2005:

Asset Class	Geometric Mean	Arithmetic Mean	Standard Deviation
Large Company Stocks	10.4%	12.3%	20.2%
Small Company Stocks	12.6%	17.4%	32.9%
Long-term Corporate Bonds	5.9%	6.2%	8.5%
Long-term Government Bonds	5.3%	5.5%	5.7%
U.S. Treasury Bills	3.7%	3.8%	3.1%
Inflation	3.0%	3.1%	4.3%

Source: *A Random Walk Down Wall Street* - Burton Malkiel (2007), page 185; data from Ibbotson Associates.

This table is consistent with our intuition: stocks are riskier than bonds, with small company stock being the riskiest and short-term Treasuries the least risky.

Risk vs. return

The assets also exhibit a clear relationship between risk and return:

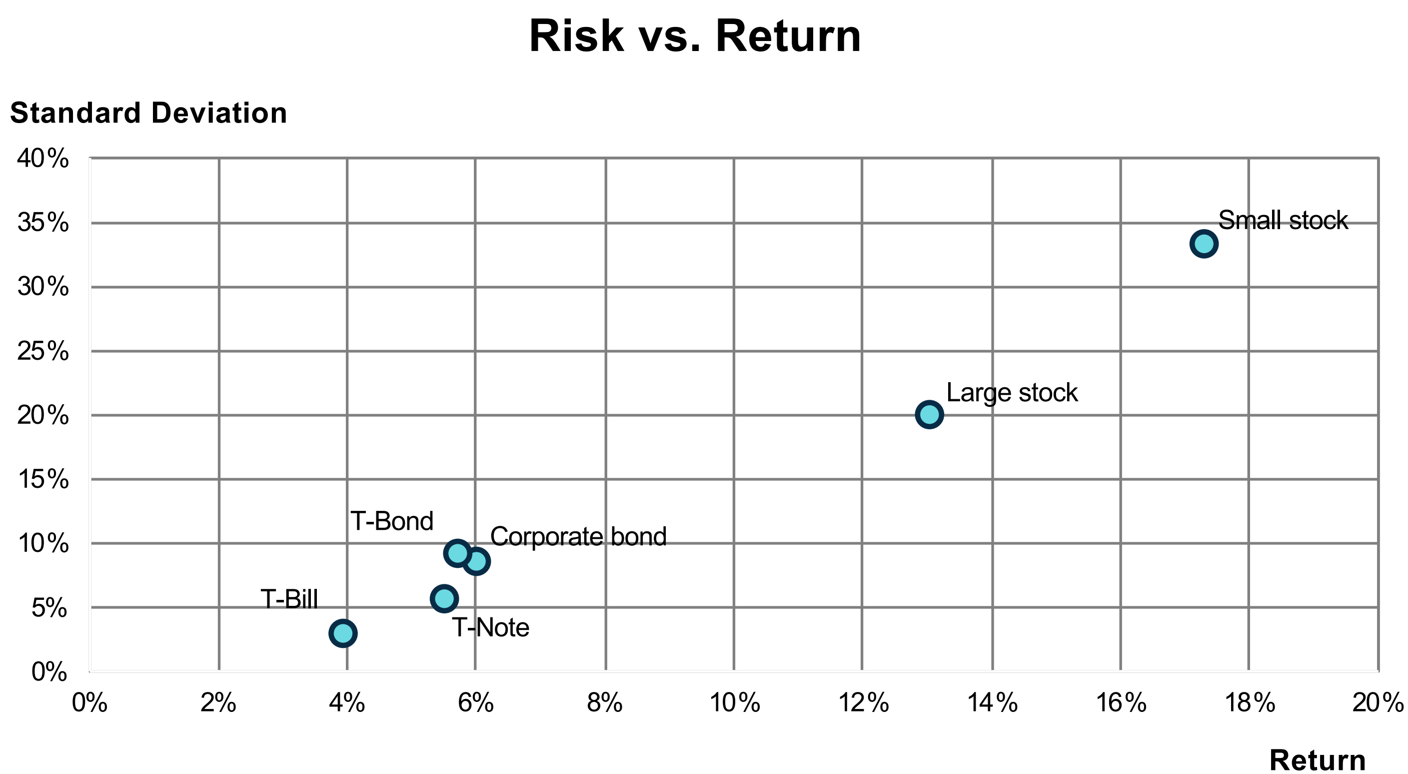

There is a trade-off between risk and return. **Assets with higher returns are also the riskiest**.

Risk on Stocks vs. Bonds

Stocks vs. bonds in the short-term

The above table only listed average returns. Although higher on average, **stock returns may fluctuate wildly from year-to-year**:

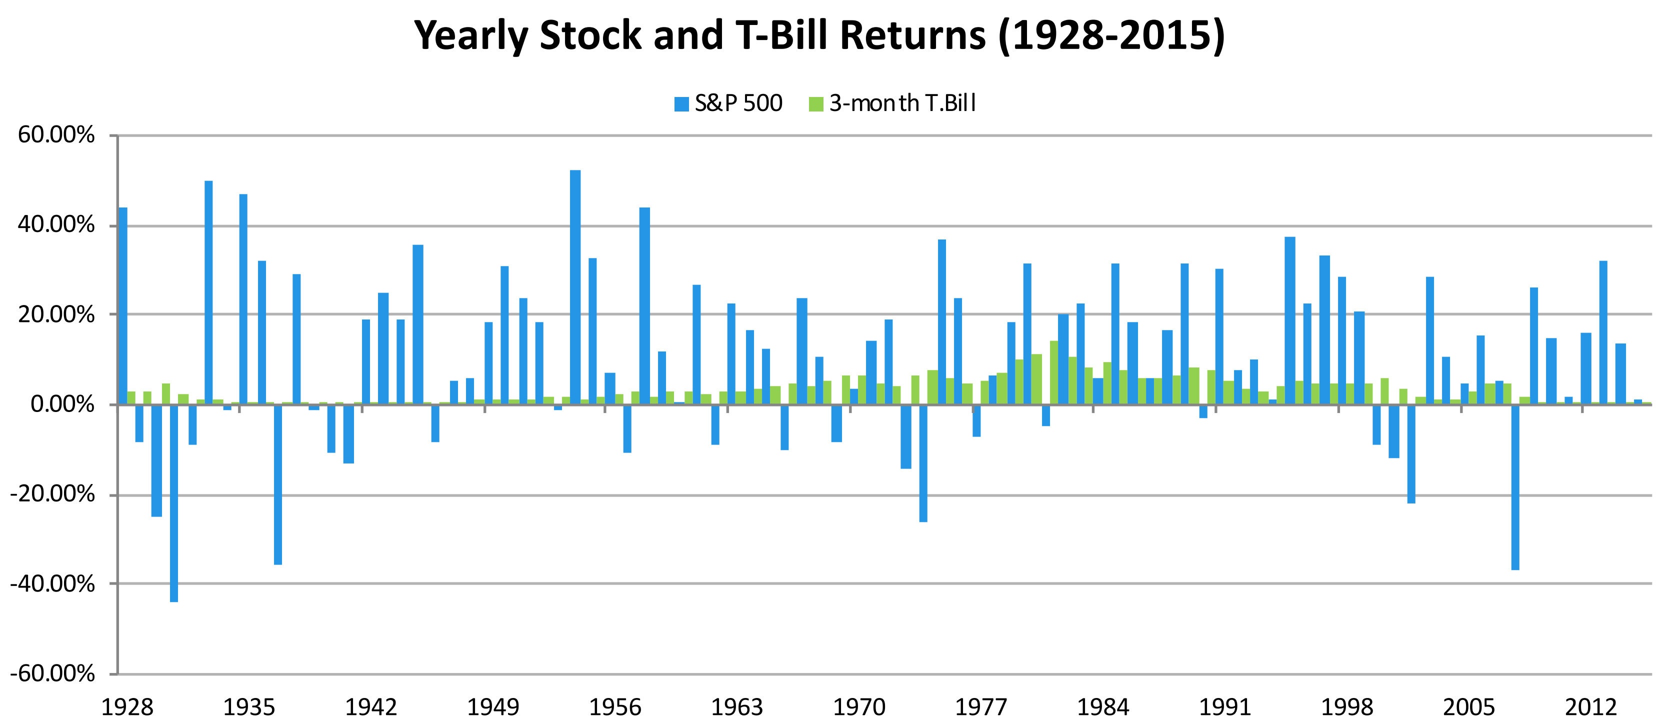

Source: Federal Reserve of St. Louis Economic Data (FRED).
Damodaran, A. "Historical returns: Stocks, T.Bonds & T.Bills with premiums." NYU Stern, 2016.

Which would you rather invest in: stocks or T-bills?

Stocks vs. bonds in the long-term

Now here is a chart of average stock and T-bill returns over 30-year periods:

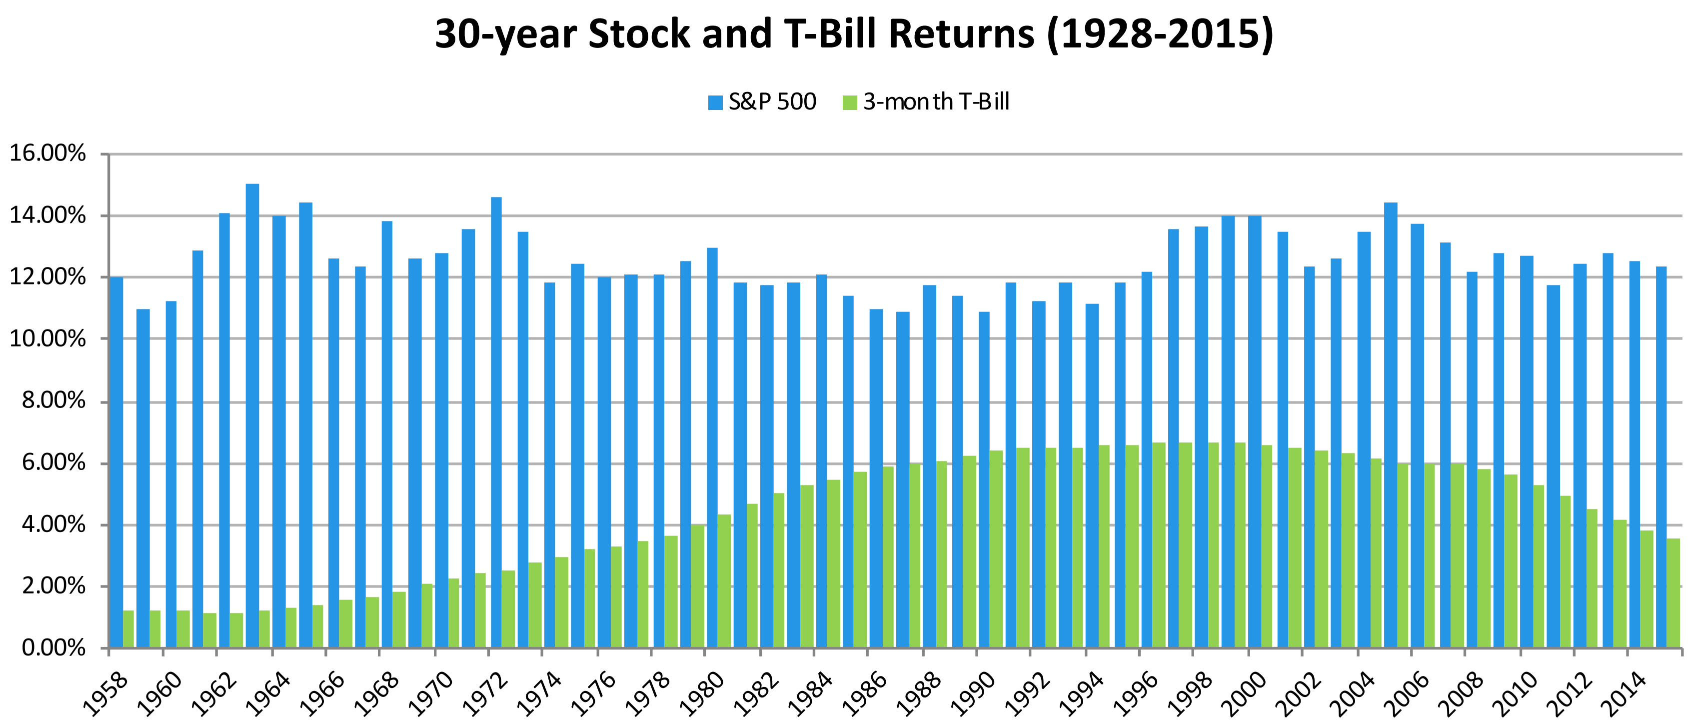

Source: Federal Reserve of St. Louis Economic Data (FRED).
Damodaran, A. "Historical returns: Stocks, T.Bonds & T.Bills with premiums." NYU Stern, 2016.

Which would you rather invest in over the long-term?

Stocks vs. bonds in the long-term

The higher long-run returns on stocks over T-bills means that each dollar invested in stocks will grow much more than a dollar invested in T-bills:

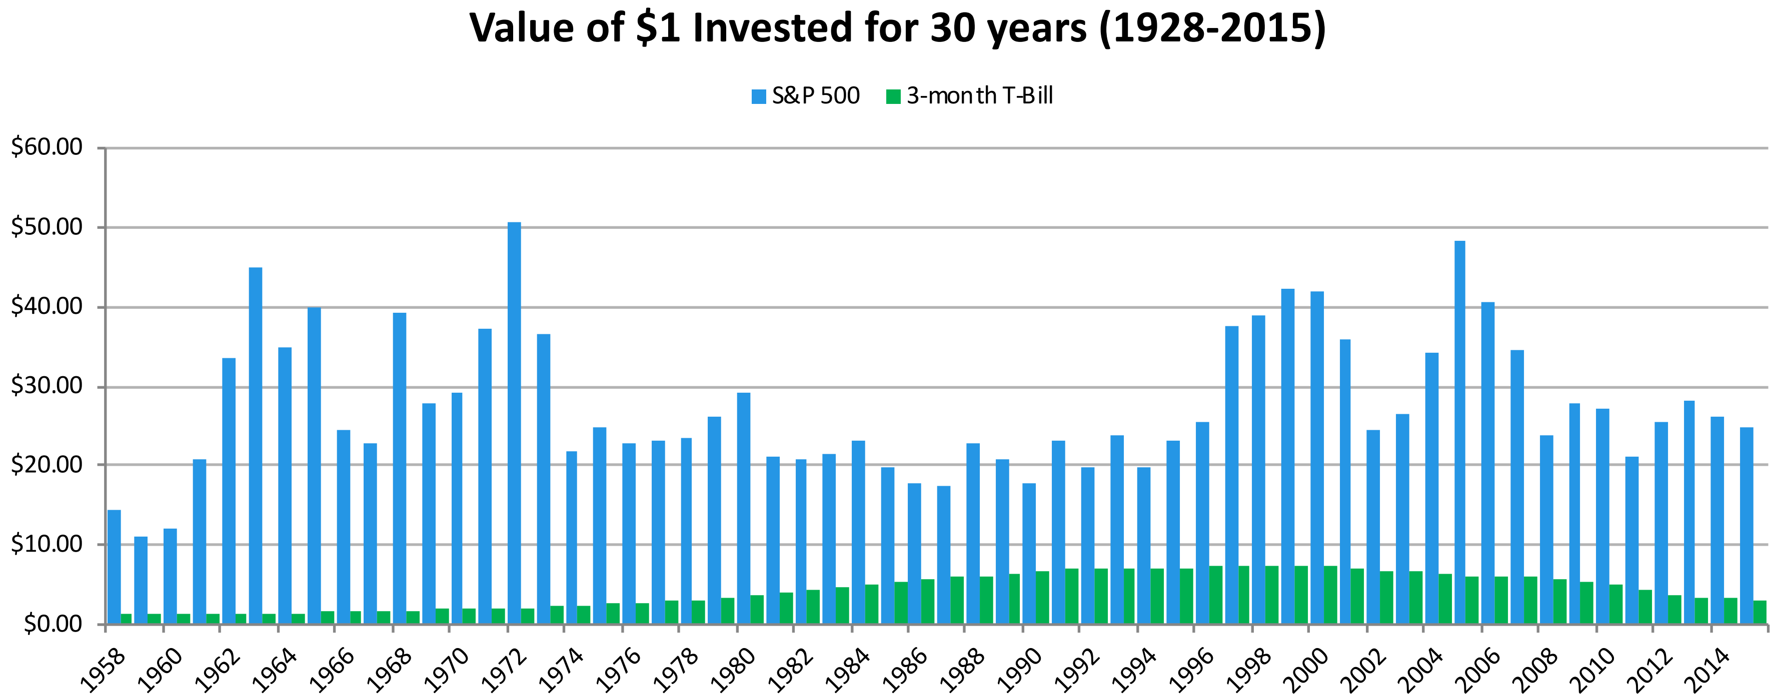

Source: Federal Reserve of St. Louis Economic Data (FRED).
Damodaran, A. "Historical returns: Stocks, T.Bonds & T.Bills with premiums." NYU Stern, 2016.

Of course, the real future **purchasing power** of a dollar invested in either stocks or T-bills will be reduced by **inflation** over such a long time period…

Risk versus return

When investing, it's important to remember the relationship between risk and return.

- Assets that are riskier tend to have higher returns. This is because **smart investors demand higher returns as compensation for higher risk**.
- The converse is also true. If an asset seems to offer a higher return, it is likely riskier. **Don't be fooled into thinking a you can get high returns without risk**.
- If you can bear it, risk may be worth the investment. Though stocks are risky from year to year and sometimes perform worse than bonds, over long stretches they have always outperformed bonds (but this may not be true in the future).
- And remember the power of interest compounding. A small difference in returns over the long run can make a big difference!

Today we learned…

- ✓ Basic probability theory
- ✓ The law of large numbers
- ✓ Expected value
- ✓ Risk vs. return
- ✓ Risk aversion
- ✓ Credit spread
- ✓ Risk on stocks vs. bonds