Shanghai Advanced Institute of Finance | Jun Pan
Shanghai Jiao Tong University | junpan@saif.sjtu.edu.cn

Empirical Asset Pricing, Fall 2019 | November 20, 2019

Classes 9 & 10: Options

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 Options, an Overview

- **Why Options?** The development of options as an exchange-traded product was an important landmark in the practice of Finance. It offers investors an alternate way to buy and sell the risk inherent in the underlying stock. In the language developed later, it offers non-linear exposures to the underlying stock or index. This non-linearity cuts the entire distribution of stock returns into various pieces.

  After the 2008 crisis, people sneered at the Wall Street practices such as tranching and repackaging. I think this is very unfortunate. A small fraction of Wall Street clearly mis-used and abused derivatives and contributed to the financial crisis in 2008. Looking back into the history of financial innovation, this was not the first time, nor will it be the last.

  Finance is about optimal allocation of risk: match the right kind of risk to the right kind of investors and distribute the right kind of investment to the right kind of firms or entrepreneurs. If we think of this distributional effort as a network of pipelines, then financial markets on equity, bond, foreign exchange, and commodity offer the basic infrastructure. The limited flexibility of these markets gave rise to derivatives.

  Options are a very good example. When we invest in the stock market, we have to take the whole package: the entire distribution of the stock. In our earlier classes, we talked about how we can minimize our exposure to idiosyncratic risk by forming portfolios and how we can take out the market risk by long/short strategies. The motivation behind options is the same. What if we are interested in hedging out not the entire market risk, but only a specific portion of the market risk, say the left tail? The long/short strategy will not help us do that: you are either all in or all out. But buying a put option on the S&P 500 index will achieve this goal for you. Now the question is how much are you willing to pay for this product? This is option pricing.

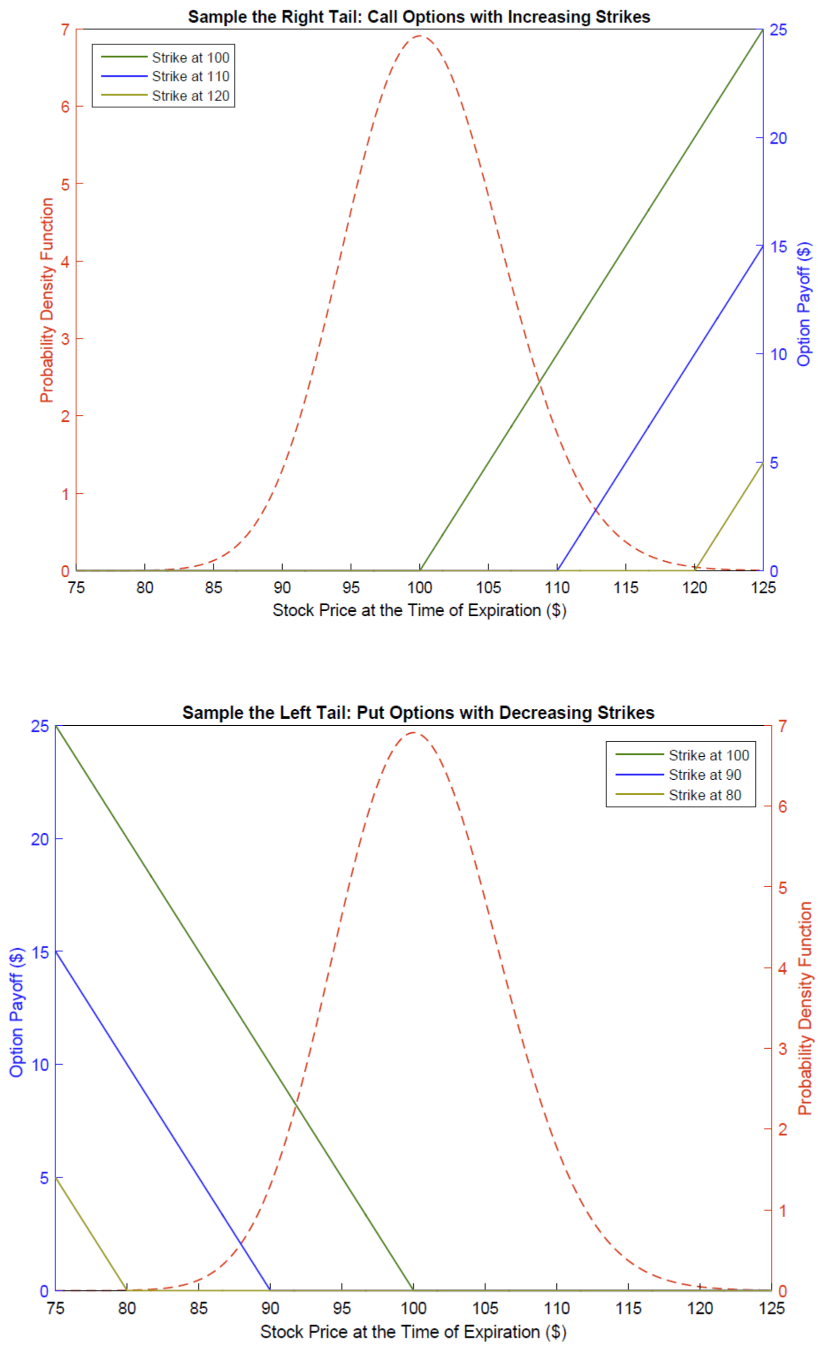

Figure 1: The distribution of a stock plotted against the payoff function of call and put options with varying strike prices.

Let me expand on this example further. As illustrated in Figure 1, moving the strike price of a call option from left to right with increasing strike prices, we are making the call option more and more out of the money. At the same time, this call option becomes more and more sensitive to the right tail of the distribution. Likewise, moving the strike price of a put option from right to left with decreasing strike prices, we are making the put option more and more out of the money. At the same time, this put option becomes more and more sensitive to the left tail of the distribution. Effectively, the market's valuations of such OTM call and put options provide us information about the right and left tails. As we learned early, the left and right tails are not abstract concepts. They are made of extreme financial events: crises show up on the left tail and rallies add to the right tail.

This is as if we are given a high definition camera with a super strong zooming ability. We can point our camera to the right tail and zoom into that area using an OTM call option. Likewise, an OTM put option allows us to zoom into the left tail. If you a photographer, you would be overjoyed to own such a high-definition camera. Likewise, if you are in the business of risk, you would naturally be drawn to these new financial instruments.

- **History:** These new instruments called options first showed up as an exchange-traded product in April 1973, exactly one month before the publication of the Black-Scholes paper. On the first day of trading, 911 contracts of calls were traded on 16 underlying stocks. One option contract is on 100 underlying shares.

  By 1975, the Black-Scholes model was adopted for pricing options. This is an excerpt from an interview with Prof. Merton: *Within months they all adopted our model. All the students we produced at MIT, I couldn't keep them in-house; they were getting hired by Wall Street. Texas Instruments created a specialized calculator with the formula in it for people in the pits. Scholes asked if we could get royalties. They said, "No." Then he asked if we could get a free one, and they said, "No."*

  It was not until 1977, four years after the trading of call options, when trading in put options begins. In 1983, the first index option (OEX) begins trading and a few months later SPX, options on the S&P 500 index, was launched. My PhD thesis was on option pricing and when I first started to work on the CBOE data in 1997, OEX, options on the S&P 100 index, still had a large market presence. By now, it has only a tiny market share. In 1993, CBOE started to publish the VIX index, which was effectively the Black-Scholes implied volatility for an at-the-money one-month to expiration SPX. In 2004, CBOE launches futures on VIX and later options on VIX.

- **Trading Volume and Market Size:** To gauge the activity of a market, the most frequently used measure is trading volume. For the U.S. equity market, the exchange-listed stocks are traded on 11 stock exchanges ("lit" markets) and about 45 alternative trading systems ("dark pools"). According to summary data from BATS, for the month of September 2015, the average daily trading in the stock market is 7.92 billion shares and \$321 billion (dollar volume). For the same month, the overall daily trading volume in the options market is about 16.94 million contracts and \$6.30 billion (dollar volume). As you can see, in terms of trading volume, the options market is small compared with its underlying stock market.

  In comparing the trading volumes in the stock and options markets, one interesting observation is that, after the 2008 crisis, the trading in the stock market has been badly hurt. For example, the average NYSE group trading volume peaked around 2.6 billion shares per day in 2008 and has decreased quite dramatically to a level near 1.0 billion shares per day in 2013 and 2014. This is not an NYSE specific problem. The overall stock market trading peaked in 2009 around 9.76 billion shares per day and bottomed to 6.19 billion shares in 2013. By contrast, the trading volume in options did not suffer this dramatic reduction. The average daily trading volume was around 14 million contracts in 2008, increased to 18 million contracts in 2011, and held up steady at around 16 million contracts in 2013.

  In terms of size, the U.S. equity market has a total market value of \$26 trillion by end-2014. At the end of September 2015, the open interest for equity and ETF options is 292 million contracts, and 23.7 million contracts for index options. Given that the average premium is around \$200 for equity and ETF options and \$1,575 for index options, this open interest amounts to \$95.7 billion in total market value. Again, the options market is small compared to its underlying stock market.

- **Leverage in Options:** Although the options market is small compared to the underlying stock market, the risk in this market is anything but small. Because of the non-linearity, the leverage inherent in options could be large. Given an investment of the same dollar amount, the profit and loss in options could be many times larger than those in the underlying stocks.

  For example, let's consider a one-month at-the-money put option. Using the Black-Scholes pricing formula, Figure 2 plots the returns to this option as a function of the underlying stock returns, assuming the stock return volatility is 20% per year. As we can see from the plot, for a 10% drop in the underlying stock price, the option yields a return over 300%. So the inherent leverage in options amplifies a dollar's

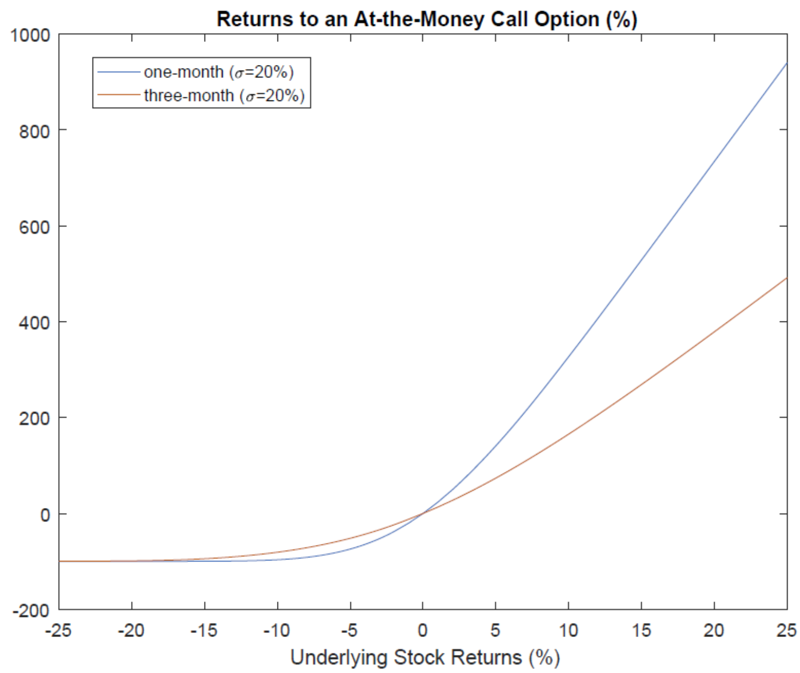

Figure 2: The return of an at-the-money put option plotted against the underlying stock return

investment in the underlying stock to 10 dollars in options. Likewise, a 10% increase in the underlying stock price translates to a near -100% drop in the put option. This amplification effect shows up in call options as well, except that the profit and loss of a call option is in the same direction as the underlying stock. Because of these amplification effects, the beta of options on the S&P 500 index can be easily around 20 or -20. Searching through the thousands of stocks listed on the three major U.S. exchanges, you will not be able to find one single stock with this kind of beta. This is what a very simple, almost innocent, non-linearity in the payoff function does to the transformation of risk.

- **Types of Options:** Broadly speaking, there are three types of exchange-traded options: equity, ETF, and index options. Equity options are American-style call and put options on individual stocks. One contract is on 100 underlying shares and the option settles by physical delivery. This CBOE link gives the exact specifications of equity options. Using the September 2015 numbers as an example, the average daily trading volume for equity options is around 7.57 million contracts and $1.58 billion per day in dollar trading volume.

  On any given day, there are thousands of stocks with options traded. Larger stocks usually have higher options trading volume. In September 2015, options on AAPL are by far the most active options traded. Other popular stocks include FB, BAC, NFLX, and BABA, although there is quite a bit of variation over time in terms which stock options show up among the actively traded. If you are curious, this OCC link provides monthly summaries of all equity and ETF option trading volume by exchange.

ETF options are American-style call and put options on ETFs. Again, one contract is on 100 underlying shares and the option settles by physical delivery of the underlying ETF. This CBOE link gives the exact specifications of ETF options. Since the mid-2000s, the growth in ETF options is an important development in the options market. For September 2015, the average daily trading volume for ETF options is around 7.33 million contracts per day, on par with the trading activity for equity options. The dollar trading volume in ETF options averages to \$1.50 billion per day, similar in magnitude to equity options.

Among the popular ETFs are SPY, EEM, IWM, and QQQ, which command relatively high option trading volume. By far, the most actively traded ETF option is SPY (options on SPDR). For September 2015, SPY options are traded on 12 options exchanges with an average daily volume of 3.3 million contracts.

Index options are European-style call and put options on stock indices. Except for mini products, one contract is on 100 underlying index. Instead of physical delivery, the settlement of index options is done by cash. This CBOE link gives the specifications of SPX, the most important index options. For September 2015, the average daily trading volume of SPX is about 1.14 million contract and \$2.78 billion in dollar trading volume. Recall that the overall dollar trading volume in the options market is about \$6.30 billion. This implies that over 30% of the options dollar trading volume comes from SPX. It is therefore not surprising that all options exchanges would like to get involved with this product. So far, CBOE is able to maintain the exclusive license in this product.

You might also notice that both SPX and SPY are trading on the S&P 500 index. There are, however, a few differences between these two products. SPX is on the index itself while SPY is on the ETF SPDR, which is about 1/10 of the index. As a result, per contract, SPX is larger in size than SPY. Recall that average daily trading volume in SPY from 12 exchanges adds up to 3.3 million contracts. This translates to a daily trading volume around \$804 million, a large number for ETF and equity options but small compared with the daily dollar volume of \$2.78 billion for SPX. Finally, while SPY is an American-style option, SPX is European-style; SPY is physical settlement while SPX is cash settlement.

Regardless of their differences, SPX and SPY share the same underlying. Therefore there must be market participants who actively trade between these two contracts to profit from any temporary mis-pricing between the two. As a result, the pricing of these two contracts should be very much aligned with one another, taking into account of the difference in their exercise style. For those who are interested, it might be a good exercise to go to the CBOE's website to get quotes for near-the-money near-the-term SPX and SPY call and put options, back out the Black-Scholes implied volatilities from these contracts and see if there are any significant pricing differences (above and beyond the quoted bid and ask spreads).

- **Options Exchanges:** As we see earlier, over 30% of the option dollar trading volume comes from SPX: call and put options on the S&P 500 index. Not surprisingly, CBOE fought really heard to keep its exclusive rights to SPX. In 2012, after 6 years of litigation, CBOE won the battle and was able to retain its exclusive licenses on options on the S&P 500 index. As a result, CBOE remains its dominance in index options with over 98% of the market share. In addition to SPX, options on VIX have also grown in popularity, which is also traded exclusively on CBOE.

  In other areas, however, CBOE has not been able to retain its market power. Until the late 1990s, CBOE was the main exchange for options trading. By the early 2000s, however, CBOE was losing its market share in equity options to new option exchanges like ISE. For equity options in September 2015, CBOE accounts for 16.32% of the trading volume, PHLX has a market share of 17.50%, and the rest are shared by BATS (14.37%), ARCA (11.81%), ISE (10.53%), AMEX (8.89%) and others. Trading in ETF options took off around the mid-2000 and have been spread over many options exchanges in a way similar to equity options: CBOE (15.93%), ISE (15.55%), PHLX (15.02%), BATS (10.48%), ARC (10.15%), AMEX (10.14%), and others.

  You might have noticed the fragmentation of the options market. Indeed, equity and ETF options are traded in 12 different options exchanges. This phenomenon of market fragmentation is not option specific. For example, US stocks regularly trade on 11 exchanges. In addition to these exchanges which are called "lit" markets, a non-trivial amount (20% to 30% in 2015) of stock trading is done in alternative trading systems such as "dark pool."

- **Market Participants:** One advantage of options being traded on exchanges is its accessibility. Investors of all types come to the market to trade. Another advantage is its transparency. Information on transaction prices and volumes is readily available to investors. On any given day, you can see how many put options are bought on the

S&P 500 index or on AAPL versus how many call options are bought. The same thing cannot be said about the over-the-counter (OTC) derivatives market. While pricing information on OTC derivatives can be obtained from Bloomberg or Datastream, the real-time transaction information is very much protected by dealers as proprietary information. In my personal view, if the trading information in products such as CDS, CDO, synthetic CDO, and CDO2 were available to the public back in 2005, more people would have paid attention to this market.

Like most markets, there are designated market makers in the options market. Their presence in the market is to facilitate trading and provide liquidity. They make money by quoting bid and ask prices: buy at the bid and sell at the ask. The bid-ask spread (ask price - bid price) is the source of their profit. In the options market, the percentage bid-ask spread is much larger than that of the underlying stock, reflecting the leverage risk inherent in options. It is also a reflection of the relative illiquidity in options. When there are buying and selling imbalances, market makers might have to keep an inventory, which exposes them to market risk. This risk exposure is further exaggerated if this imbalance is caused by some private information the market maker is not aware of (information asymmetry). The inventory cost and the cost of information asymmetry are two important drivers for the bid and ask spread in financial prices. In the options market, it is typical for a market maker to minimize his exposure to the underlying stock by delta hedging.

Coming back to the topic of SPX and SPY, two options products with very similar underlying risk. You might notice that there is a substantial difference in their bid/ask spreads. In particular, the average bid/ask spreads (as a percentage of the option price) are much higher for SPX than SPY. The average percentage bid/ask spread for SPX is about 9% while that for SPY is about 1%. If the market risk is similar, then where does this difference in trading cost arise?

Investors who trade against the market makers can be summarized into four groups: customers from full service brokerage firms (e.g., hedge funds), customers from discount brokerage firms (e.g., retail investors), and firm proprietary traders. Using CBOE data from 1990 through 2001, we see that customer from full serve brokerage firms are the most active participants in the options market while firm proprietary traders concentrate their trading mostly on index options as a hedging vehicle. Of course, these are older data and the options market has exploded after 2001.

Another way to look at the market participants is through their trading activities against the market makers. Some investors come to the options market to buy options to open a new position, while other buy options to close an existing position. Some sell options to new a new position while others sell options to close an existing position. In doing so, their trading motives are very different.

2 The Black-Scholes Option Pricing Model

- **The Model:** Let S_t be the stock price at time t. For simplicity, let's first assume that this stock pays no dividend. Later we will add dividend back. We model the dynamics of the stock price by the following model (geometric Brownian motion):

$$dS_t = \mu\, S_t\, dt + \sigma\, S_t\, dB_t\,. \tag{1}$$

  This equation does not look very appealing at the moment, but you will come to appreciate or even like it later. Under this model, the expected stock return is μ and its volatility is σ, both numbers are in annualized terms. So if you like, μ is about 12% and σ is about 20%. Moreover, under this model, stock returns (to be more precise, log-returns) are normally distributed. Let me use the rest of this section to explain why it is so.

  Let S_T be the stock price at time T. Implicitly we are planning ahead for the time T, when the option expires. Standing here at time 0 and holding a European-style option, all we care about is the final payoff:

$$\text{Payoff of a call option struck at K} = (S_T - K)\,\mathbf{1}_{S_T>K} \tag{2}$$

  where $\mathbf{1}_{S_T>K} = 1$ if $S_T > K$ and zero otherwise. Let's focus on call options for now. Once we know how to deal with call options, the put/call parity will get us to put options very easily.

  Option pricing bolts down to calculating the present value of the payoff in equation (2). How should this calculation be done? What is the discount rate to use in order to bring the random cash flow to today? Let's keep this question hanging for a while.

- **Brownian Motion:** Since it is the first time we are working with Brownian motions, let me summarize the following three important properties of Brownian motions and relate them to Finance:

  - *Independence of increments:* For all $0 = t_0 < t_1 < \ldots < t_m$, the increments are independent: $B(t_1) - B(t_0)$, $B(t_2) - B(t_1)$, $\ldots$, $B(t_m) - B(t_{m-1})$. Translating to Finance: stock returns are independently distributed. No predictability and zero auto-correlation $\rho = 0$.

- *Stationary normal increments:* $B_t - B_s$ is normally distributed with zero mean and variance $t-s$. Translating to Finance: stock returns are normally distributed. Over a fixed horizon of T, return volatility is scaled by $\sqrt{T}$.
- *Continuity of paths:* $B(t)$, $t \geq 0$ are continuous functions of t. Translating to Finance: stock prices move in a continuous fashion. There are no jumps or discontinuities.

- **The Model in R_T:** Let's perform this very important transformation:

$$S_T = S_0 \, e^{R_T} \, .$$

Another way to look at it is by,

$$R_T = \ln(S_T) - \ln(S_0) \, ,$$

which tells us that R_T is the log-return of the stock over the horizon T. Now I am going to do one magic and you just have to trust me on this. Next semester when you take 450, you will learn the mechanics behind it, which is call the Ito's Lemma.

$$dR_t = \left(\mu - \frac{1}{2}\sigma^2\right) \, dt + \sigma \, dB_t$$

Comparing with equation (1), the dynamics of R_t is simpler. It does not have those μS_t and σS_t terms. Instead, we have $\mu - \sigma^2/2$ as its drift and σ as its diffusion coefficient. The extra term of $\sigma^2/2$ is often call the Ito's term.

With this dynamics for R_t, we can now fix the time horizon T and write out R_T:

$$R_T = \int_0^T dR_t = \left(\mu - \frac{1}{2}\sigma^2\right) T + \sigma\sqrt{T}\,\epsilon_T \, , \tag{3}$$

where ϵ_T is a standard normal random variable (zero mean, variance equals to 1). You will agree with me that $\int_0^T dt = T$. Let me explain why $\int_0^T dB_t = B_T - B_0$ is $\sqrt{T}\,\epsilon_T$: it comes from the second property, stationary normal increments, of the Brownian motion.

When it comes to valuation under the Black-Scholes model, the math will be done at the level of equation (3). As you can see, it is not that scary, isn't it? This model tells us that the log-return of a stock over a fixed horizon of T is normally distributed with mean $(\mu - \sigma^2/2)\,T$ and standard deviation of $\sigma\sqrt{T}$. Other than the Ito's term, $\sigma^2/2$, everything looks quite familiar. No?

- **The Ito's Term:** Now let me explain why we have this Ito's term. In the continuous-time model of equation (1), the stock price grows at the instantaneous rate of $\mu\, dt$:

$$E(S_T) = S_0\, e^{\mu T}\,,$$

or equivalently, with a continuously compounded discount rate μ:

$$S_0 = e^{-\mu T}\, E(S_T)\,.$$

Now let's do the same calculation with our model for log-return in Equation (3),

$$E(S_T) = S_0\, E\left(e^{R_T}\right)\,.$$

When it comes to calculating expectation of a convex function involving a normally distributed random variable x, this is a useful formula for you to have

$$E(e^x) = e^{E(x)+\mathrm{v}ar(x)/2}\,.$$

Let me emphasize, this works only when x is normally distributed. Applying this formula to the above calculation, we have

$$E(S_T) = S_0\, E\left(e^{R_T}\right) = S_0\, e^{E(R_T)+\mathrm{v}ar(R_T)/2} = S_0\, e^{\left(\mu-\sigma^2/2\right)T+\sigma^2\, T/2} = S_0\, e^{\mu T}\,,$$

which is exactly what we wanted in the first place.

To summarize, the transformation from S_T to $\ln(S_T) - \ln(S_0)$ introduces some concavity, because $\ln(x)$ is a concave function. This is why $-\sigma^2/2$ shows up in R_T. The transformation from R_T to e^{R_T} introduces some convexity, because e^x is a convex function, and $\sigma^2/2$ gets added back during the transformation. So everything works out.

In essence, Mr. Ito is busy because we are doing concave/convex transformations on random variables. If there is no random variable involved, then Mr. Ito will not be this busy. For example, let's make x a number by setting $\mathrm{v}ar(x) = 0$. What do we have for $E(e^x) = e^{E(x)+\mathrm{v}ar(x)/2}$? We have $E(e^x) = e^x$ and nothing else. The Ito's term disappeared.

- **Risk-Neutral Pricing:** Now let's come back to the present value calculation. As discussed earlier, the payoff of a call option at time T is as in Equation (2). It is a random payoff, depending on the realization of S_T. It is a non-linear random payoff with a kink at the strike price K: the payoff is zero if S_T falls below K and is $S_T - K$ if S_T rallies above K and the option is exercised at time T. So what is the present value of this random non-linear payoff? Which discount rate should we use?

Risk-neutral pricing is the answer to that question. Although it has "risk-neutral" in its name, it is anything but risk-neutral. Let me first tell you the approach of the risk-neutral pricing. Recall that after some hard work, we have

$$R_T = \left(\mu - \frac{1}{2}\sigma^2\right) T + \sigma\sqrt{T}\,\epsilon_T\,.$$

I am going to call this model the actual dynamics and label it by "P." Then I am going to introduce a different model, called risk-neutral dynamics and label it by "Q."

$$\text{Actual Dynamics ("P"):} \qquad R_T = \left(\mu - \frac{1}{2}\sigma^2\right) T + \sigma\sqrt{T}\,\epsilon_T \tag{4}$$

$$\text{Risk-Neutral Dynamics ("Q"):} \quad R_T = \left(r - \frac{1}{2}\sigma^2\right) T + \sigma\sqrt{T}\,\epsilon_T^{\mathrm{Q}} \tag{5}$$

By writing down the model in the Q-dynamics, I am bending the reality by forcing the stock return to grow at the riskfree rate r. And then I am going to do my present value calculation under this bent reality: 1) the expectation of the future cash flow is done under the Q-measure and 2) this expectation is discounted back to today using the riskfree rate r. And somehow, two wrongs make one right, the calculation works out. You just have to trust me on this. This pricing framework is widely adopted on Wall Street in fixed income, credit, and options.

- **Pricing a Stock:** Before applying this risk-neutral pricing framework on options, let's first try it on something easier: the linear random payoff of S_T. We know what the answer should be: the present value should be S_0. We've already done it under the P-dynamics: $S_0 = e^{-\mu T} E(S_T)$. It works out and using μ as the discount rate makes perfect sense ... because this is how the dynamics is written.

  Now let's do it under the Q-dynamics:

  $$e^{-rT}\,E^Q(S_T) = e^{-rT}\,S_0\,e^{rT} = S_0\,.$$

  So it also works! Just to emphasize that risk-neutral pricing has nothing to do with investors being risk-neutral, let's bring in a risk-neutral investor to price the same stock. He takes the P-dynamics (because it is the reality) and discounts the cash flow with riskfree rate r (because he is risk neutral):

$$e^{-rT}\, E^{\mathrm{P}}(S_T) = e^{-rT}\, S_0\, e^{\mu T} = S_0\, e^{(\mu-r)T}$$

So he is paying more than S_0 for the same cash flow. Why? Because he is risk-neutral. Recall that if S_T is the market portfolio, then $\mu - r$ is the market risk premium. Risk-averse investors demand a premium for holding the systematic risk in the market portfolio. That gives rise to the positive risk premium in $\mu - r$. A risk-neutral investor, however, is not sensitive to risk. As such, he is willing to pay more for the stock.

This exercise might seem trivial mathematically, but it is very useful in clearing our thoughts. In particular, I would like to emphasize that risk-neutral pricing does not mean pricing using a risk-neutral investor. In a way, this name "risk-neutral pricing" is unfortunate and confusing.

- **Pricing the Option:** We are now ready to price the option. Let C_0 be the present value of a European-style call option on S_T with strike price K:

$$C_0 = e^{-rT} E^{Q}\left((S_T - K)\,\mathbf{1}_{S_T>K}\right) = e^{-rT} E^{Q}\left(S_T \mathbf{1}_{S_T>K}\right) - e^{-rT} K E^{Q}\left(\mathbf{1}_{S_T>K}\right)$$

Now let me cheat a little by going directly to the solution,

$$C_0 = S_0\, N(d_1) - e^{-rT} K N(d_2)\,,$$

where $N(d)$ is the cumulative distribution function of a standard normal x:

$$N(d) = \mathrm{Prob}\,(x \le d) = \int_{-\infty}^{d} \frac{1}{\sqrt{2\pi}}\, e^{-\frac{x^2}{2}}\, dx\,.$$

In Matlab, $N(d)$ is normcdf(d). The two critical values d_1 and d_2 are,

$$d_1 = \frac{\ln\left(S_0/K\right) + \left(r + \sigma^2/2\right) T}{\sigma\sqrt{T}}\,; \quad d_2 = \frac{\ln\left(S_0/K\right) + \left(r - \sigma^2/2\right) T}{\sigma\sqrt{T}}$$

Comparing where we are now with the solution, we see some internal logic. In particular, it is obvious that

$$N\left(d_2\right) = E^{Q}(\mathbf{1}_{S_T>K}) = \mathrm{Prob}^{Q}\left(S_T > K\right)$$

and

$$N\left(d_1\right) = e^{-rT} E^{Q}\left(\frac{S_T}{S_0}\,\mathbf{1}_{S_T>K}\right)\,.$$

- **Understanding $N(d_2)$ in the Black-Scholes formula:** The part associated with $N(d_2)$ is actually pretty easy. It calculates the probability that the call option is in the money under the Q-measure. So let's work it out:

  $$\text{Prob}^Q\left(S_T > K\right) = \text{Prob}^Q\left(S_0\, e^{R_T} > K\right) = \text{Prob}^Q\left(e^{R_T} > K/S_0\right) = \text{Prob}^Q\left(R_T > \ln(K/S_0)\right)\,,$$

  where, in the last step, I took a log on both side of the inequality, which is OK because $\ln(x)$ is a monotonically increasing function in x.

  Now let's use the Q-dynamics of R_T in Equation (5) to get,

  $$\text{Prob}^Q\left(R_T > \ln(K/S_0)\right) = \text{Prob}^Q\left(\left(r - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}\,\epsilon_T^{\text{Q}} > \ln(K/S_0)\right)\,.$$

  Moving things left and right, we get

  $$\text{Prob}^Q\left(\epsilon_T^Q > \frac{-\ln(S_0/K) - \left(r - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}\right)\,,$$

  or equivalently,

  $$\text{Prob}^Q\left(-\,\epsilon_T^Q < \frac{\ln(S_0/K) + \left(r - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}\right)\,,$$

  which is really $N(d_2)$, knowing that ϵ_T^Q is standard normally distributed.

- **Understanding $N(d_1)$ in the Black-Scholes formula:** The part associated with $N(d_1)$ is more subtle. Recall that

  $$N\left(d_1\right) = e^{-rT}E^Q\left(\frac{S_T}{S_0}\,\mathbf{1}_{S_T>K}\right)$$

  So $N(d_1)$ involves a calculation that takes into account that we are calculating the expectation of S_T when S_T is greater than K (the option expires in the money). So it is not a simple probability calculation such as $N(d_2)$. Here, it involves an interaction term. As a result $N(d_1)$ should always be larger than $N(d_2)$. This is true because $d_1 = d_2 + \sigma\sqrt{T}$. As you will see later, this difference between d_1 and d_2 is really where the option value of an option comes from. In other words, $\sigma\sqrt{T}$ is the best summary of the option value.

  Given the amount of math we have been doing up to this point, I have a feeling that most of you are not willing to go further. For those of you who are interested, you can do the math to prove that $N(d_1)$ is in fact $e^{-rT}E^Q\left(\frac{S_T}{S_0}\,\mathbf{1}_{S_T>K}\right)$.

For those who are not willing to go through the math, let me offer this observation. Under the Q-dynamics, the drift in R_T is $(r-\sigma^2/2)T$ and the volatility is $\sigma\sqrt{T}$. That's how we get the expression of d_2 (and our previous calculation just proved this point). Comparing d_1 and d_2 this way, we notice that suppose we bend the reality further by making the drift in R_T to be $(r+\sigma^2/2)T$ and keep the same volatility. Then, under this strange dynamics, let's call it QQ, we have $N(d_1) = \text{Prob}^{QQ}(S_T > K)$. Intuitively, because of the interaction term, the valuation is higher. One simple way to express this higher valuation is by allowing R_T to grow faster than its Q-measure, with a drift of $(r+\sigma^2/2)T$. Under this probability measure, the probability of S_T is greater than K (the option expires in the money) becomes $N(d_1)$. I'll stop here.

- **Add Dividend Yield:** We are going to apply the Black-Scholes model to SPX. So it is important that we can handle stocks paying dividend with a constant dividend yield, which, for the S&P 500 index, is a good enough approximate. Let q be the dividend yield. Again, let S_T be the time-T stock price, ex dividend. Then, the stock dynamics becomes,

$$dS_t = (\mu - q)\; S_t\, dt + \sigma\, S_t\, dB_t\,.$$

And the dynamics for R_T changes to

$$\text{Actual Dynamics ("P"):} \qquad R_T = \left(\mu - q - \frac{1}{2}\sigma^2\right) T + \sigma\sqrt{T}\,\epsilon_T$$

$$\text{Risk-Neutral Dynamics ("Q"):} \quad R_T = \left(r - q - \frac{1}{2}\sigma^2\right) T + \sigma\sqrt{T}\,\epsilon_T^{\mathrm{Q}}$$

And the Black-Scholes pricing formula becomes

$$C_0 = e^{-qT}\, S_0\, N(d_1) - e^{-rT}\, K\, N(d_2)\,,$$

where $N(d)$ is the cumulative distribution function of a standard normal and

$$d_1 = \frac{\ln{(S_0/K)} + (r - q + \sigma^2/2)\,T}{\sigma\sqrt{T}}\,; \quad d_2 = \frac{\ln{(S_0/K)} + (r - q - \sigma^2/2)\,T}{\sigma\sqrt{T}}$$

- **Arbitrage Pricing and Dynamic Replication:** In Finance, when it comes to valuation, there are just two approaches: equilibrium pricing and arbitrage pricing. We've touched on equilibrium pricing in the CAPM, where mean-variance investors optimize their utility functions and the equity and bond markets clear. What we've been doing so far in this class falls squarely into the category of arbitrage pricing. The essence of arbitrage pricing is replication: replicate a stream of random payoffs with existing securities whose market values are known to us. The present value of this cash flow equals to the cost of the replication.

The best example in our current setting is the put/call parity. As I am sure that you've learned in 15.415 (or 15.401), the time-T payoff of buying a European-style call and selling a European-style put (with the same strike price K) is the same as taking a long position in the underlying stock and borrowing K from the bond market. The present value of the underlying stock is $e^{-qT}S_0$, where, as usual, we use ex dividend stock price. The present value of the bond-borrowing portion is $e^{-rT}K$, with r being the riskfree rate, continuously compounded. So the replication cost is $e^{-qT}S_0 - e^{-rT}K$. The present value of buying a call and selling a put is, by definition, $C_0 - P_0$. As a result, $C_0 - P_0 = e^{-qT}S_0 - e^{-rT}K$.

As you can see, in getting this relation, we do not have to use any model, just simple logic. In practice, this put/call ration holds pretty well in the market. There are investors actively arbitrage between the options and the cash (i.e., the S&P 500 index or the S&P 500 index futures via "E-mini") markets. Even if the Black-Scholes model fails (which it does), this relation still holds. Arbitraging using put/call parity is very similar to arbitraging between the futures and cash markets (arbitraging between Chicago and New York).

When it comes to pricing call and put options, however, we do need to use a model. So far, we've used the Black-Scholes model. It turns out that even with a stock and a bond, we can still replicate the non-linear payoff of an option. This is the important insight of Prof. Black, Merton, and Scholes: dynamic replication. You need to continuously rebalance your hedging portfolio, doing delta hedging at a super high frequency. I am sure that you've got a heavy dosage of that in your 15.415. So I am not going to spend time on dynamic replication or delta hedging.

Recall that the third property of a Brownian motion is continuity of paths. This implies that stock prices move in a continuous fashion. There is no jumps or discontinuities. This is why models like geometric Brownian motions are called diffusion models. As you can see, the property of dynamic replication falls apart as soon as we move away from the Brownian motion by adding random jumps to the model. This is just one example. If we add another streams of random shocks to volatility, making it a stochastic process (instead of a number $\sigma = 20\%$), then this replication also falls apart.

As such, the Black-Scholes formula is very much confined to the model itself. We will see that the Black-Scholes model does not hold very well in the market. We will then extend in two dimensions: adding jumps to the model to allow crashes; relaxing σ from a number to a stochastic process and build a stochastic volatility model.

- **Why so many equations?** Since Fall 2015, because of the MFin students, I made a conscious effort in being as rigorous as possible and giving you as much detail as possible. While using the Black-Scholes model as a black box is fine for most people, I feel that most of you deserve to know a little bit better. In past years, 15.450 was taught along with 15.433. So I made the comfortable choice of letting the professor in 15.450 carry more of the math burden. Now that 15.450 has been moved to the Spring semester, I feel that I've lost my excuse. And Prof. Wang kept asking me to push you more. So this is my effort in pushing you.

  If you've seen this before, don't presume that you know everything. Honestly, I started to work in this area as soon as I entered the PhD program at Stanford GSB 20 years ago. But I've only developed these intuitions over the years. So take your time to digest the materials and make them your own.

3 Using the Black-Scholes Formula

- **Pricing ATM Options:** By definition, an at-the-money option has the strike price of $K = S_0\, e^{(r-q)T}$. Going back to d_1 and d_2, we notice that by setting the strike price at this level, $d_1 = \frac{1}{2}\sigma\sqrt{T}$ and $d_2 = -\frac{1}{2}\sigma\sqrt{T}$. Effectively, by having an option with this strike price, we take away the moneyness component of the option and focus exclusively on the option value. Also notice that at this strike price, $e^{-qT}\, S_0 = e^{-rT}\, K$, which implies that, via the put/call parity, $C_0 = P_0$ for this pair of at-the-money call and put options. For the case of $\sigma = 20\%$ and $T = 1/12$, we have $d_1 = \sigma\sqrt{T}/2 = 0.0289$. Figure 3 plots the respective $N(d_1)$ and $N(d_2)$ for the case.

  Applying the Black-Scholes formula (assuming $q = 0$), we have

$$C_0 = P_0 = S_0\,(N(d_1) - N(d_2)) = S_0\left[N\left(\frac{1}{2}\sigma\sqrt{T}\right) - N\left(-\frac{1}{2}\sigma\sqrt{T}\right)\right] .$$

  Using the fact that $N(d)$ is the cdf of a standard normal:

$$N(d) = \int_{-\infty}^{d} \frac{1}{\sqrt{2\pi}}\, e^{-\frac{x^2}{2}}\, dx\,,$$

  we can further simplify the pricing formula,

$$\frac{C_0}{S_0} = \frac{P_0}{S_0} = \frac{1}{\sqrt{2\pi}} \int_{-\frac{1}{2}\sigma\sqrt{T}}^{\frac{1}{2}\sigma\sqrt{T}} e^{-\frac{x^2}{2}}\, dx\,.$$

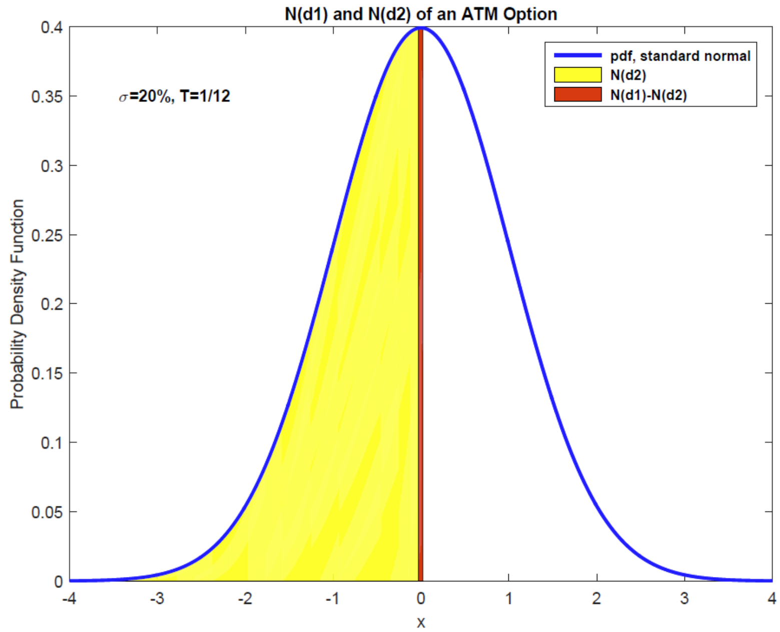

Figure 3: The $N(d_1)$ and $N(d_2)$ for an one at-the-money call or put option with one-month to expiration. The underlying stock volatility is 20%.

Now let's use a Taylor expansion that is very useful in Finance: $e^x \approx 1 + x$, for small x. Applying this to the integrand,

$$e^{-\frac{x^2}{2}} = 1 - \frac{x^2}{2}.$$

Let's replace the integrand with this approximate:

$$\frac{C_0}{S_0} = \frac{P_0}{S_0} \approx \frac{1}{\sqrt{2\pi}} \int_{-\frac{1}{2}\sigma\sqrt{T}}^{\frac{1}{2}\sigma\sqrt{T}} \left(1 - \frac{x^2}{2}\right) dx = \frac{1}{\sqrt{2\pi}} \left(\sigma\sqrt{T} - \frac{1}{24}\left(\sigma\sqrt{T}\right)^3\right) \approx \frac{1}{\sqrt{2\pi}} \sigma\sqrt{T},$$

where I dropped the cubic term to make our approximation even simpler. But you can see, if you include the next order of approximation, the net effect will make the option price lower. This approximation works well for small $\sigma\sqrt{T}$. For a typical one-month option on the S&P 500 index, $\sigma = 0.20$ and T=1/12, we have $\sigma\sqrt{T}$ being around 0.0577. As a comparison, the higher order term $\left(\sigma\sqrt{T}\right)^3/24$ is 8×10^{-6}. So this level of $\sigma\sqrt{T}$, our approximation works really well.

As shown in Figure 4, as $\sigma\sqrt{T}$ becomes large, this approximation becomes imprecise. Moreover, the approximation is bias upward compared with the Black-Scholes pricing. This makes sense because the next higher order term is negative. It also makes sense because C_0/S_0 cannot grow linearly with $\sigma\sqrt{T}$ forever. The call option price is bounded

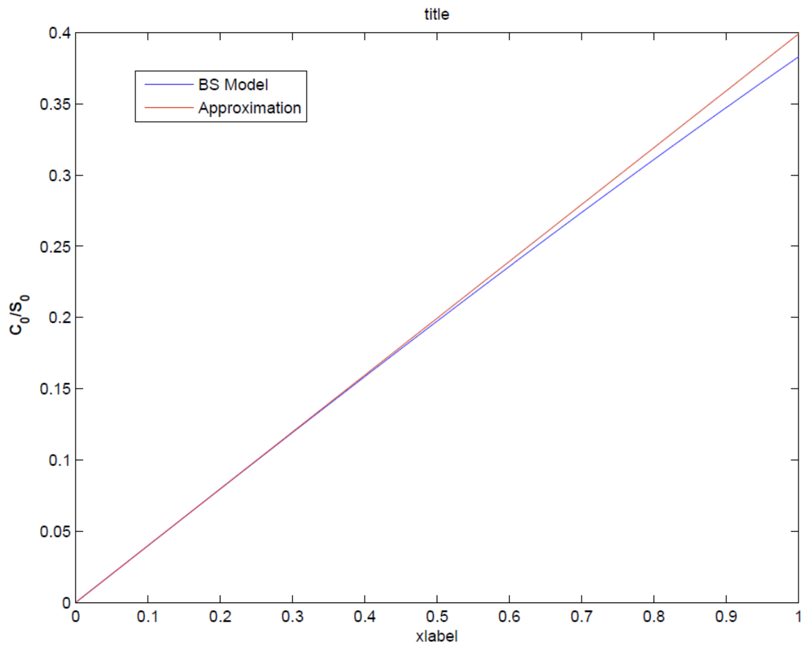

Figure 4: The ratio of an at-the-money call or put option price to the underlying stock price, C_0/S_0 or P_0/S_0, as a function of $\sigma\sqrt{T}$. The approximation of $C_0/S_0 = P_0/S_0 \approx \sigma\sqrt{T}/\sqrt{2\pi}$ is in red and the Black-Scholes pricing is in blue.

from above by the underlying stock price: C_0/S_0 cannot be bigger than 1. At some point, this ratio has to taper off.

What kind of options will give us $\sigma\sqrt{T}$ that is too large for this approximation to work? Options on volatile stocks with long time to expire. For example, for an option with $\sigma = 100\%$ and 1 year to expiration, $\sigma\sqrt{T} = 1$. As you can see from Figure 4, our approximation is no longer very good.

- **ATM Options as Financial Vehicles on $\sigma\sqrt{T}$:** In spending time to analyze the at-the-money options, we learned an important lesson. In fact, it is the cleanest way to understand what options are really about. By buying a call option, we get a positive exposure to the underlying stock; by buying a put option, we get a negative exposure. Neither of these exposures is unique to options. There are other ways we can get this kind of exposure. And the exposure can be easily hedged out by stocks. But what's unique about options is the volatility exposure. In the Black-Scholes model, volatility is a constant. So you might not appreciate the significance of this volatility exposure. As soon as we allow volatility to move around, which is true in reality, then you find in options a vehicle that is unique in offering exposures to $\sigma\sqrt{T}$. Nothing in the stock market can offer this kind of exposure.

Recall that dynamic replication makes options a redundant security within the Black-Scholes model. At that point, you might be wondering to yourself that: if it is redundant, then what is the point? Well, in reality, with random shocks to volatility and fat-tails in stock returns, options are not at all redundant. That is why, as beautiful and revolutionary as the dynamic replication theory is, I do not want us to spend too much time on it.

Going back to our discussions regarding $N(d_1)$ and $N(d_2)$, the example of ATM options further clarifies what really matters in d_1 and d_2. It's the fact that d_1 is always larger than d_2, by the amount of $\sigma\sqrt{T}$ in the Black-Scholes model. If you trace back to the calculation of d_1, you notice that it comes from $E^Q\left(S_T\mathbf{1}_{S_T>K}\right)$, the positive interaction between S_T and $\mathbf{1}_{S_T>K}$. Within the Black-Scholes setting, we have the exact formulation of this option value. As we later move away from the Black-Scholes model, $N(d_1)$ and $N(d_2)$ will be replaced by other formulas. That is why I have been emphasizing calculations like $E^Q\left(\mathbf{1}_{S_T>K}\right)$ and $E^Q\left(S_T\mathbf{1}_{S_T>K}\right)$ for call options. These calculations are the main building blocks of a call option, whose values might be different in different models. Likewise, for put options, calculations like $E^Q\left(\mathbf{1}_{S_T<K}\right)$ and $E^Q\left(S_T\mathbf{1}_{S_T<K}\right)$ are the main building blocks.

- **The Black-Scholes Option Implied Volatility:** Once we understand that options are unique financial vehicles for volatility, then volatility will be the first thing we would like to learn from options. Indeed, the Black-Scholes option implied volatility is such a concept.

  For a call option with strike price K and time to expiration T, we can calculate its Black-Scholes price by plugging the model parameters. We obtain the underlying stock price S_0 from the stock market, the riskfree rate r from the Treasury or LIBOR market. If this option is on the S&P 500 index, we can assume a flow of dividend payment in the form of a dividend yield q. We can approximate q with its historical average, say 2%. Now the only parameter left for us to move around is σ. Of course, we can go to the underlying stock market to measure the volatility. But let's not do that. Let's instead back out the volatility σ^I so that the model price for this option agrees with the market price of this option. This is the Black-Scholes implied volatility.

  In doing this exercise, we are not assuming the Black-Scholes model is correct. We are only using the model as a tool for us to transform the option price from the dollar space to the volatility space. Why is this useful? Because options with different strike prices and times to expiration will differ quite a lot in their market value. A deep in-the-money option might be worth hundreds of dollars, while a deep out-of-the-money option on the same underlying might be worth just a few dollars. A short-dated options is worth much less than a long-dated options. Since all of these options are on the same underlying, you would like to be able to compare their pricing. But comparing these options in the dollar space is not at all intuitive. By contrast, all of these options share the same underlying. Hence the same σ. So comparing these options in the volatility space is much more intuitive and productive. In fact, in OTC markets, options are typically quoted not in dollar but in the Black-Scholes implied volatility. This is analogous to the adoption of yields in the bond market. So Black-Scholes implied vols in options and yields in bonds.

4 Bring the Black-Scholes Model to the Data

The Black-Scholes option pricing model, along with the arbitrage-free risk-neutral pricing framework, is something of a revolution in Finance. It managed to attract many mathematicians, physicists, and even engineers to Finance. But if the progression stopped right at the level of modeling and pricing, it would have been rather boring: you take the pricing formula, plug in the numbers, and get the price. So things would have been pretty mechanical. Real life is always more interesting than financial models. In this class, let's bring the model to the data and enjoy the discovery process.

- **ATM Options and Time-Varying Volatility:** Volatility plays a central role in option pricing. In the Black-Scholes model, volatility σ is a constant. If you take this assumption literally, then the Black-Scholes implied vol σ_t^I should be a constant over time. In practice, this is not at all true. As we learned in our earlier class on time-varying volatility, using either SMA or EWMA models, the volatility measured from the underlying stock market moves over time. Recall this plot, Figure 5, in Classes 8 & 9, where the option-implied volatility is plotted against the volatility measured directly from the underlying stock market. In both cases, stock return volatility varies over time.

  One interesting observation offered by Figure 5 is that the option-implied volatility is usually higher than the actual realized volatility in the stock market. In other words, within the Black-Scholes model, the options are more expensive than what can be justified by the underlying stock market volatility. If you believe in the Black-Scholes model, then selling volatility (via selling near-the-money options, calls or puts) will be a very profitable trading strategy.

  Figure 6 plots the time-series of VIX (option-implied volatility using SPX) against the time-series of the S&P 500 index level. As you can see, the random shocks to

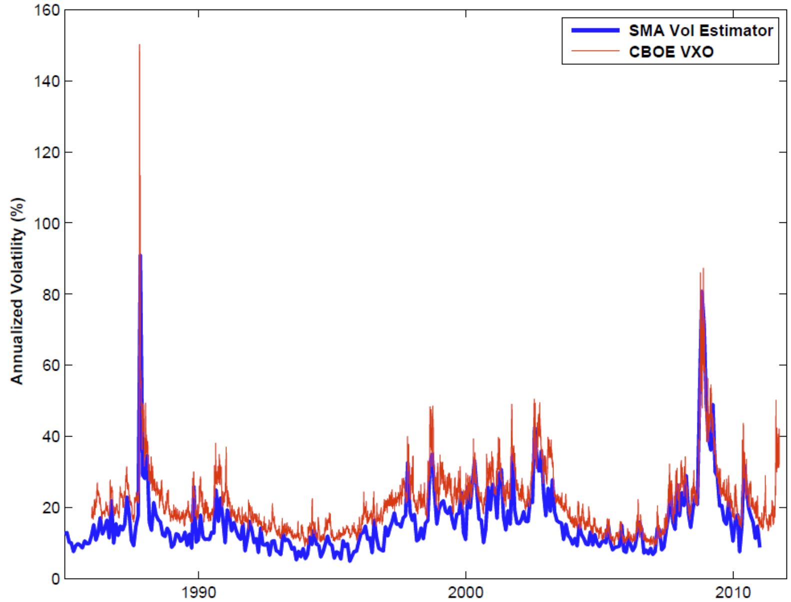

Figure 5: Time-Varying Volatility of the S&P 500 Index. The red line is the option-implied volatility using SPX traded on CBOE. The blue line is measured directly from the underlying stock market using daily returns of the S&P index.

VIX, especially those sudden increases in VIX are often accompanied by sudden and large drops in the index level. Of course, this observation is outside of the Black-Scholes model, where σ is a constant. But this plot gives us the intuition as to what could go wrong with selling volatility: you lose money when the markets are in crisis. Basically, by selling volatility on the overall market (e.g., SPX), your capital is at risk exactly when capital is scarce. In the language of the CAPM, you have a positive beta exposure.

But this positive beta exposure is more subtle than the simple linear co-movement captured by beta. As highlighted by the shaded areas, volatility typically spikes up when there are large crises. Just to name a few: the October 1987 stock market crash, the January 1991 Iraq war, the September 1993 Sterling crisis, the 1997 Asian crisis, the 1998 LTCM crisis, 9/11, 2002 Internet bubble burst, 2005 downgrade of GM and Ford, 2007 pre-crisis, March 2008 Bear Stearns, September 2008 Lehman, the European and Greek crises in 2010 and 2011, and the August 2015 Chinese spillover. In other words, what captured by Figure 6 is co-movement in extreme events, like the crisis beta in Assignment 1 (risk exposure conditioning on large negative stock returns). Also, as shown in Figure 6, not all crises have the same impact. For example, the downgrade of GM and Ford was a big event for the credit market, but not too scary for equity and index options.

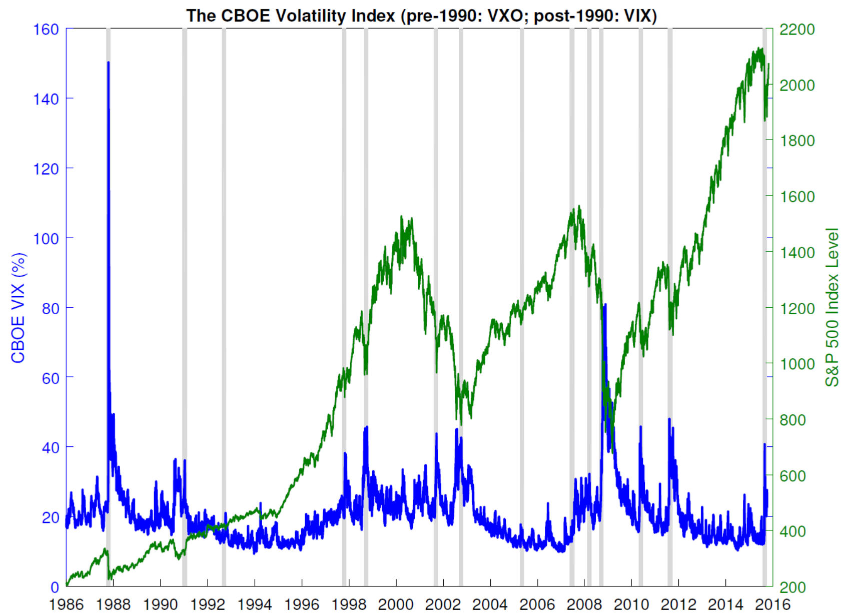

Figure 6: Time-Series of the CBOE VIX Index Plotted against the Time-Series of the S&P 500 Index Level. Prior to 1990, the old VIX (VXO) is used. Post 1990, the news VIX index is used.

The comovement in Figure 6 gives rise to a negative correlation between the S&P 500 index returns and changes in VIX, which ranges between -50% to -90%. Figure 7 is an old plot from Classes 8 & 9, which uses the EWMA model to estimate the correlation between the two. As you can see from the plot, the correlation has experienced a regime change. During the early sample period, the correlation hovers around -50%, while in more recent period, the correlation has become more severe, hovering around -80%.

All of these observations have direct impact on how options should be priced in practice: the Black-Scholes model needs to allow σ to vary over time. The time variation of σ should not be modeled in a deterministic fashion. As shown in Figure 6, the time series of σ_t is affected by uncertain, random shocks. So just like the stock price S_t follows a stochastic process (e.g., geometric Brownian motion), σ_t itself should follow a stochastic process with its own random shocks. Moreover, the random shocks in σ_t should be negatively correlated with the random shocks in S_t to match the empirical evidence in Figure 6. There is a class of diffusion models called stochastic volatility models developed exactly for this purpose.

These models are similar to the discrete-time models like EWMA or GARCH, which

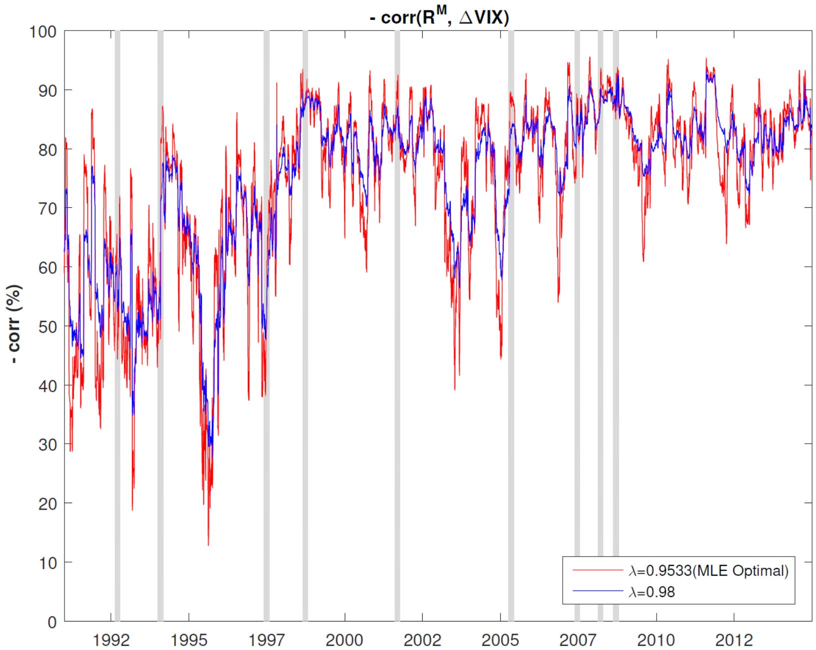

Figure 7: The Time-Series of EWMA Estimates for the Correlations between the S&P 500 Index Returns and Daily Changes in the VIX Index.

also allow volatility to be time-varying. But one distinct feature of stochastic volatility models is that it has its own random shocks. In EWMA or GARCH, the time-varying volatility comes from the random shocks in the stock market. We will come back to the stochastic volatility model later in the class, which are very useful in pricing options of different times to expiration, linking the pricing of long-dated options to that of short-dated options.

- **OTM Options and Tail Events:** In developing our intuition for the Black-Scholes model, we've focused mostly on the ATM options, which are important vehicles for volatility exposure. Now let's look at the pricing of the out-of-the-money options.

  Recall the risk-neutral pricing of a call option,

$$C_0 = E^Q\left(e^{-rT}(S_T - K)\mathbf{1}_{S_T>K}\right) = \boxed{e^{-rT}\, E^Q\left(S_T \mathbf{1}_{S_T>K}\right)} - \boxed{e^{-rT}\, K\, E^Q\left(\mathbf{1}_{S_T>K}\right)},$$

  where the pricing bolts down to calculations involving $E^Q(\mathbf{1}_{S_T>K})$ and $E^Q(S_T\mathbf{1}_{S_T>K})$. For $K > S_0 e^{rT}$, the call option is out of the money. In fact, the larger the strike price K, the more out of the money the option is, and the smaller $E^Q(\mathbf{1}_{S_T>K})$. So if we focus on OTM calls, we zoom into the right tail.

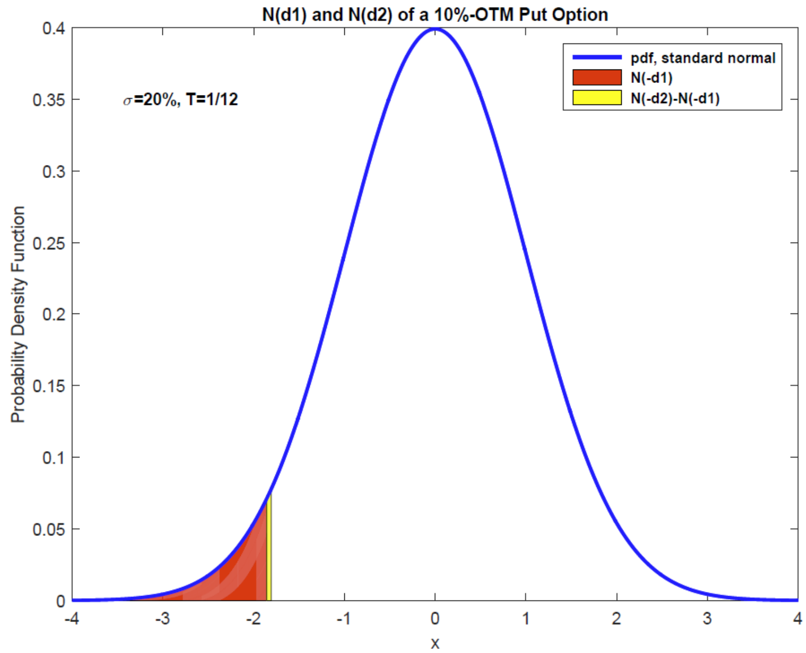

Figure 8: The Distribution of a Standard Normal with $N(-d_1)$ and $N(-d_2)$.

.

Likewise, the risk-neutral pricing of a put option is,

$$P_0 = E^Q\left(e^{-rT}(K - S_T)\mathbf{1}_{S_T<K}\right) = \boxed{e^{-rT}\, K\, E^Q\left(\mathbf{1}_{S_T<K}\right)} - \boxed{e^{-rT}\, E^Q\left(S_T \mathbf{1}_{S_T<K}\right)},$$

where the pricing bolts down to calculations involving $E^Q(\mathbf{1}_{S_T<K})$ and $E^Q(S_T\mathbf{1}_{S_T<K})$. For $K < S_0 e^{rT}$, the put option is out of the money. In fact, the smaller the strike price K, the more out of the money the option is, and the smaller $E^Q(\mathbf{1}_{S_T<K})$. So if we focus on OTM puts, we zoom into the left tail.

Within the Black-Scholes model, the above calculations can be taken to the next level using the probability distribution of a standard normal:

$$P_0 = \boxed{e^{-rT}\, K\, E^Q\left(\mathbf{1}_{S_T<K}\right)} - \boxed{e^{-rT}\, E^Q\left(S_T\mathbf{1}_{S_T<K}\right)} = \boxed{e^{-rT}\, K\, N(-d_2)} - \boxed{S_0\, N(-d_1)},$$

where I've changed the color coding so that this equation matches with Figure 8. More specifically, for a 10% OTM put striking at $K = S_0\, e^{r\,T} \times 90\%$, we can re-write the above pricing into:

$$\frac{P_0}{S_0} = \frac{e^{-rT}K}{S_0}\boxed{N(-d_2)} - \boxed{N(-d_1)} = 0.90 \times \boxed{N(-d_2)} - \boxed{N(-d_1)}$$

Figure 8 gives us a graphical presentation of what matters when it comes to pricing such OTM options: the left tail in red and the slice in yellow. The areas in red and yellow are mapped directly to the CDF of a standard normal (hence $N(-d_1)$ and $N(-d_2)$ - $N(-d_1)$) because we are working under the Black-Scholes model. But the intuitive goes further. For any distribution (even it is not normal), what matters for the pricing of this OTM put option is the left-tail distribution. If this left tail is fat because of many financial crises, then the pricing of OTM put options should reflect these tail events. In Assignment 3, you will have a chance to work with a model with crash and see the link between fat tails and option prices.

As mentioned a few times, the actual distribution of stock market returns is not normally distributed. This is especially true for returns at higher frequencies (e.g., daily returns). As such, the Black-Scholes model fails to capture the fat tails in the data. As we will see, this becomes a rather important issue when it comes to pricing options. Conversely, by looking at how these OTM options are priced, we learn about investors' assessment and attitude toward these tail events.

- **Option Implied Smirks:** After the 1987 stock market crash, one very robust pattern arose from the index options (SPX) market called volatility smiles or smirks.

  Consider the nearest term options, say one month to expiration (T=1/12). Let's vary the strike price of these options. Typically, for options with one month to expiration, you can find tradings of OTM puts and calls that are up to 10% out of the money. It is generally the case that OTM options are more actively traded than in-the-money options. This makes sense. If you are using options for speculations, you would prefer options that are cheaper (and are liquid) so that you can get more action for each dollar invested in options. If you are using options for hedging, it is likely that you are hedging out tails events. So either way, the OTM puts and calls are referred instruments than ITM options.

  Between OTM calls and puts of SPX, it is generally the case that OTM puts are more actively traded and the level of OTM-ness can reach up to 20%. For the S&P 500 index, a typical annual volatility is 20%, implying a monthly volatility of 5.77%. So for a 10% OTM put option, it takes a drop of 1.733-sigma (10%/5.77%) move in the S&P 500 index over a one-month period for this option to come back to the money.

  Using the market prices of all the available SPX puts and calls, we can back out the Black-Scholes implied volatility σ^I for each one of them. If investors are pricing the options according to the Black-Scholes model, then we should see σ^I being exactly the same for all of these options, regardless of the moneyness of the options. What we see in practice, however, is a pattern like that in Table 1.

Table 1: Short-Dated SPX Puts with Varying Moneyness on March 2, 2006.

P_0	S_0	K	OTM-ness	T	σ^I	P_0^{BS}
9.30	1287	1285	0.15%	16/365	10.06%	?
6.00	1287	1275	0.93%	16/365	10.64%	5.44
2.20	1287	1250	2.87%	16/365	12.74%	0.92
1.20	1287	1225	4.82%	16/365	15.91%	0.075
1.00	1287	1215	5.59%	16/365	17.24%	0.022
0.40	1287	1170	9.09%	16/365	22.19%	0.000013

Table 1 lists six short-dated OTM put options with exactly the same time to expiration but varying degrees of moneyness. The first option is nearest to the money, striking at $K = 1285$ when the underlying stock index is at $S_0 = 1287$. The last option is the farthest away from the money, striking at $K = 1170$. The S&P 500 index needs to drop by over 9% over the next 16 calendar days in order for this option to be in the money. Not surprisingly, options are cheaper as they are farther out of the money. But what's interesting is that their Black-Scholes implied vols exhibit this opposite pattern: the more out of the money a put option is, the higher its implied vol. In other words, even though the pricing of $0.40 (per option on one underlying share of the S&P 500 index) seems very cheap in dollars and cents, it is actually over priced. Plugging a $\sigma = 10.06\%$ to the Black-Scholes model (which a closer to the market volatility around March 2, 2006), the model price for this OTM put is $0.000013. In other words, this option is so out of the money, the Black-Scholes model (with normal distribution) deems its value to be close to zero. In practice, however, there are people who are willing to pay $0.40 for it.

Why? Don't they know about the Black-Scholes option pricing formula? If they care about tail events, then what about OTM calls which are sensitive to right tails? As we see in the data, the tail fatness shows up in both the left and the right. But the OTM calls are not over-priced. If anything, the implied vols of OTM calls are on average slightly lower than ATM options. That is why we are calling this pattern volatility smirk, which is an asymmetric smile.

- **Expected Option Returns:** Another way to look at the profit/loss involved in options is to calculation their expected returns like we do in the stock market. Table 2 was reported in a 2000 *Journal of Finance* paper by Prof. Coval and Shumway.

  As shown in Table 2, the weekly returns of buying put options are on average negative. There are quite a bit of variation in these returns. For the farther OTM put options,

Table 2: Expected Options Returns

Strike - Spot	-15 to -10	-10 to -5	-5 to 0	0 to 5	5 to 10
	Weekly SPX Put Option Returns (in %)				
mean return	-14.56	-12.78	-9.50	-7.71	-6.16
max return	475.88	359.18	307.88	228.57	174.70
min return	-84.03	-84.72	-87.72	-88.90	-85.98
mean BS β	-36.85	-37.53	-35.23	-31.11	-26.53
corrected return	-10.31	-8.45	-5.44	-4.12	-3.10

Option data from Jan. 1990 through Oct. 1995.

the return could be as positive as 475.88%, or as negative as -84.03%. This option has a beta of -36.85, which is due to the inherent leverage of these options. The CAPM-alpha of this investment is -10.31% per week. Whoever is selling this option would make a lot of money...on average. But he needs to be well capitalized when an event like 475.88% happens.

Calculations like those in Table 2 are rather imprecise because of the large variations in option returns. So we do not want to take the numbers too literally. But the qualitative result of this Table is important: when it comes to investing in options, there are large variations in option returns. Moreover, buying put options give you negative alpha. The more out of the money the put option is, the more negative the alpha becomes. For investors who are selling such put options, they are able to capture such alpha. But such trading strategies are in generally very dangerous. You need to be well capitalized to survive large crises like the 1987 stock market crash. Otherwise, you are just one crisis away from bankruptcy.

The results shown here in the return space is very much consistent with the earlier results in the implied-vol space, where OTM put options are over priced relative to near-the-money options. The level of over-pricing gets more severe as the put option becomes more out of the money and are more sensitive to market crashes. So it is not surprising that the put option returns are on average negative. Most of the times, you purchase an insurance against a market crash, but the crash does not happen and your put option expires out of the money. But once in a while, a crisis like 1987 or 2008 happens, then this put option brings you over-sized returns. Sitting on the other side of the trade are investors who sell/write you these crash insurances. Most of the times, they are able to pocket the premiums paid for the insurance without having to do anything. But once in while, they lose quite a bit of money if a crisis like 1987 or 2008 happens. As such, the risk profile of such option strategies differs quite significantly from that of a stock portfolio, where all instruments are linear. In Assignment 3, you will have a chance to see this kind of risk/return tradeoff of options in more details for yourselves.

5 When Crash Happens

- **Crash and Crash Premium:** The empirical evidence we've seen so far indicates that strategies involving selling volatility and selling crash insurance are profitable. As you will see for yourself in Assignment 3, the return distribution of such option strategies differs quite significantly from that of a stock portfolio, where all instruments are linear.

  In the presence of tail risk, options are no longer redundant and cannot be dynamically replicated. As such, two considerations involving the tail risk become important in the pricing of options. First, the likelihood and magnitude of the tail risk. Second, investor's aversion or preferences toward such tail events. The "over-pricing" of put options on the aggregate stock market (e.g., the S&P 500 index) reflects not only the probability and severity of market crashes, but also investors' aversion to such crashes — crash premium.

  In fact, as you will see in Assignment 3, the probability and severity of market crashes implicit in the volatility smirk are such that investors are pricing these OTM put options as if crashes like 1987 would happen at a much higher frequency. In other words, investors are willing to pay a higher price for such crash insurances even though they are "over-priced" relative to the actual amount of tail risk observed in the aggregate stock market. And the sellers of such crash insurances are only willing to sell them if they are being compensated with a premium, above and beyond the amount of tail risk in the data. This crash premium accounts for most of the "over-pricing" in short-dated OTM puts and ATM options.

  By contrast, this "over-pricing" is not severe for OTM calls because they are not very sensitive to the left tail. Instead, OTM calls are sensitive to the right tail. From how such options are priced relative to OTM puts, it is obvious that investors are not eager to pay the same amount of premium for insurances against the right tail. This makes perfect sense. The intuition comes straight from the CAPM. An OTM call is a positive beta security, which provides positive returns when the market is doing well. It is icing on the cake. By contrast, an OTM put pays when the market is in trouble — a friend in need is a friend indeed.

- **Bank of Volatility:** LTCM was a hedge fund initially specialized in fixed-income arbitrage. It was extremely successful in its earlier years. Success breeds imitation.

Soon, the fixed-income arbitrage space was crowded and spreads in arbitrage trades were shrinking. In early 1998, LTCM began to short large amounts of equity volatility. Betting that implied vol would eventually revert to its long-run mean of 15%, they shorted options at prices with an implied volatility of 19%. Their position is such that each percentage change in implied vol will make or lose $40 million in their option portfolio. The size of their vol position was so big that Morgan Stanley coined a nickname for the fund: the Central Bank of Volatility. For more details, you can read Roger Lowenstein's book on LTCM.

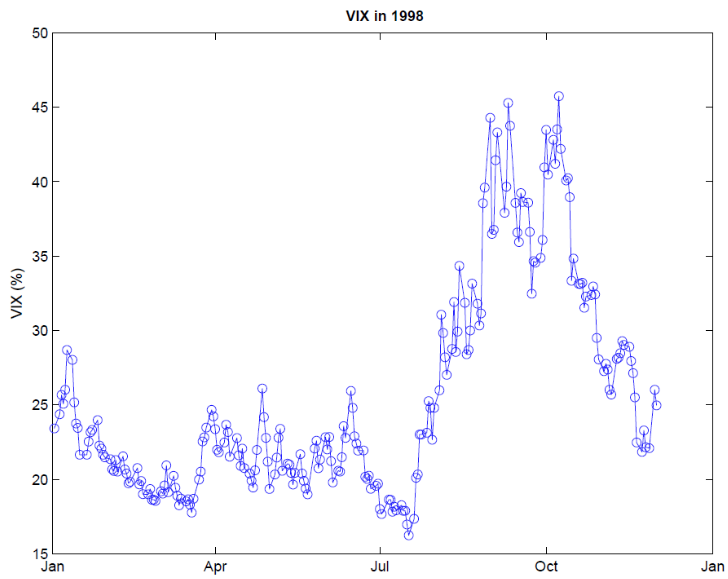

Figure 9: Time Series of CBOE VIX index in 1998.

During normal time, volatility does revert to its mean. So the idea behind the trade makes sense. Moreover, as we've seen in the data, selling volatility (via ATM options) is a profitable strategy on average because of the premium component. But the premium was not a free lunch: it exists because of the risk involved in selling volatility. As we've seen in the data, when the volatility of the aggregate market suddenly spikes up, the financial market usually is in trouble. Whenever the market is in the crisis mode, there is flight to quality: investors abandon all risky asset classes and move their capital to safe havens such as the Treasury bond market.

For the case of LTCM in 1998, it had arbitrage trades in different markets (e.g., equity, fixed-income, credit, currency, and derivatives) across different geographical locations (e.g., U.S., Japan, and European). Lowenstein's book gives more detailed descriptions of these arbitrage trades. One common characteristics of these arbitrage trades is that they locate some temporary dislocation in the market and speculate that this dislocation will die out as the market converges back to normal. In a way, these arbitrage trades betting on convergence make money because they provide liquidity to temporary market dislocations. The key risk involved in these arbitrage trades is that timing of the convergence is uncertain. Sometimes, instead of converging, the dislocation becomes even more severe before converging back to normal.

Prior to the Russian default in the summer of 1998, these arbitrage trades were not highly correlated. But after the default, most of these previously uncorrelated arbitrage trades lost money for LTCM at the same time. This certainly includes the volatility trades. As shown in Figure 9, early in the year, volatility was fluctuating around 20%. By summer 1998, however, the market became quite volatile because of the Russian default. At its peak, the VIX index was around 45%. Recall that LTCM was selling volatility when VIX was around 19% in early 1998. The position was such that each percentage change in implied vol will make or lose \$40 million. So if the volatility converges back to its long run mean of 15%, then roughly $4 \times \$40 = \$160$ can be made. But if instead of converging, the volatility increases to 45%, you can imagine the loss.

The Russian default affected not only LTCM but other hedge funds and prop trading desks who pursued the same kind of convergence trades. At a time like this, capital becomes scarce, and all leveraged investors (e.g., hedge funds or prop trading in investment banks) are desperately looking for extra source of funding. They do so by unwinding some of their arbitrage trades, further exacerbating the widening spreads. At a time like this, holding a security that pays (e.g., an existing long position in puts) could be very valuable. By contrast, a security that demands payment (e.g., an existing short position in puts) would be threatening to your survival. Therefore, being on the short side of the market volatility hurts during crises. That is why volatility is expensive (i.e., ATM options are over-priced) in the first place.

- **The 2008 crisis:** The OTM put options on the S&P 500 index is a good example for us to understand crash insurance. In writing a deep OTM put option, the investor prepares himself for the worse case scenario when the option becomes in the money. This happens when the overall market experiences a sharp decline. The probability of such events is small. But if he writes a lot of such options believing that the exposure can somehow be contained by the low probability, then he is up for a big surprise when a crisis does happen. As we learned from the recent financial crisis, some supposedly sophisticated investors wrote such OTM put options without knowing the real consequence.

Gillan Tett from *Financial Times* wrote an excellent book called *Fool's Gold* with details of how investment banks developed and later competed for the market shares of the mortgage-linked CDO products. The following is a brief summary.

By 2006, Merrill, who was late into the CDO game, topped the league table in terms of underwriting CDO's, selling a total of $52 billion that year, up from $2 billion in 2001. Behind the scenes, Merrill was facing the same problem that worried Winters at J.P. Morgan: what to do with the super-senior tranch?

CDO's are the collateralized debt obligations. It pools individual debt together and slices the pool into tranches according to seniority. For a mortgage-linked CDO, the underlying pool consists of mortgages of individual homeowners. The cashflow to the pool consists of their monthly mortgage payments. The most senior tranch is the first in line to receive this cashflow. Only after the senior tranch receives its promised cashflow, the next level of tranches (often called mezzanine tranches) can claim their promised cashflow. The equity tranch is the most junior and receives the residual cashflow from the pool.

As default increases in mortgages, the cashflow to the pool decreases. The equity investors will be the first to be hit by the default. If the default rate further increases, then the mezzanine tranch will be affected. The most senior tranch will only be affected in the unlikely event that both equity and mezzanine investors are wiped out and the cashflow to the pool cannot meet the promise to the most senior tranch. Such super senior tranches are usually very safe and are Aaa rated. By contrast, the mezzanine tranches are lower rated (Baa) because of the higher default risk. And the credit quality of the equity tranch is even lower.

The pricing of such products is consistent with their credit quality: the yield on the mezzanine tranches is higher than the senior tranches to compensate for the higher credit risk. Investors, in an effort to reach for yield, prefer to buy the mezzanine and equity tranches. As a result, the investment banks underwriting the CDOs are often stuck with the super senior tranches. As the business of CDOs grew, the banks are accumulating more and more highly rated super senior tranches. Initially, Merrill solved the problem by buying insurance (credit default swaps) for its super-senior debt from AIG.

Let's take a look at what the super-senior tranch is really about. It is highly rated because of the low credit risk. Imagine the economic condition under which this credit risk affecting the super-senior tranch will actually materialize: when the default risk is so high that both mezzanine and equity investors are wiped out. A typical argument for the economics of pooling is that default risk by individual homeowners can be diversified in a pool. This is indeed true when we think about the risk affecting the equity tranch: one or two defaults in the pool would affect the cashflow to the equity tranch, but would not affect the mezzanine tranch, let alone the senior tranch. So the risk affecting the senior tranch has to be a very severe one. The default rate has to be so high that the cashflow dwindles to the extent that it would eat through the lower tranches and affect the most senior tranch. In other words, many homeowners must be affected simultaneously and default at the same time to generate this type of scenario. By then, the risk is no longer idiosyncratic but systemic. So writing an insurance on a senior tranch amounts to insuring a crisis — a deep OTM put option on the entire economy.

In late 2005, AIG told Merrill that it would no longer offer the service of writing insurance on senior tranches. By then, however, AIG has already accumulated quite a large position on such insurance. Later, AIG was taken over by the US government in a $85 billion bailout and the insurance on senior tranches was honored and made whole by AIG (and the New York Fed).

After AIG declined to insure their super senior tranch, Merrill decided to start keeping the risk on its own books. At the same time, Citigroup, another late comer, was also keen to ramp up the output of its CDO machine. Unlike the brokerages, though, Citi could not park unlimited quantities of super-senior tranches on its balance sheet. Citi decided to circumvent that rule by placing large volumes of its super-senior in an extensive network of SIVs (Special Investment Vehicle) and other off balance sheet vehicles that it created. Citi further promised to buy bank the super-senior tranch if the SIVs ever ran into problems with them.

Now let's try to understand what Merrill and Citi are actually doing by retaining the super-senior tranch. Effectively, they are holding the super-senior tranch without an insurance. If you are holding a US treasury bond, you don't have to worry about credit risk (except for when the US government defaulted). So holding a super-senior tranch without an insurance is like holding a default-free US treasury bond and selling a deep OTM option on the overall economy at the same time. Before, they were able to buy that put option from AIG to hedge out this risk. Now, they are bearing this risk themselves.

Then the crisis happened in 2007 and 2008, and the mortgage default rate increased to such an extent that it started to affect the super-senior tranches. In other words, the deep OTM put options became in the money. During the 2007-08 crisis, the pricing of these super-senior tranches became one of the biggest headaches on Wall Street.

Merrill and Citi, along with other Wall Street banks, had to take billions of dollars of writedowns.

6 Beyond the Black-Scholes Model

- **A model with market crash:** In Assignment 3, you will be working closely with a model that allows market to crash. It is a simplifies version of the model in Merton (1976).

- **A model with stochastic volatility**

Appendix

A Valuation Models in Finance

As we move on to options and fixed-income products, concepts such as present value calculation will take center stage. Looking back, you might have noticed that in our equity classes, we worked almost exclusively in the return space. We analyze the distribution of stock returns, estimate the expected return, investigate the return predictability, and study the various models of return volatility. Very rarely did we talk about valuation. For example, AAPL has a market capitalization of $642B with $112 a share right now. What kind of Finance models do we use to price this stock? Can the same model be used to price other stocks? How well does such a model work in practice?

The one exception was when we work with the book-to-market ratio in the Fama-French model. We use the book value of equity as a benchmark for the market value of equity. If investors think of buying the stock as buying the book value of the firm, then this ratio should be around one. In practice, we noticed a wide range of book-to-market ratios. For example, as of July 2015, the average book-to-market ratio is around 0.095 for stocks in decile 1 and 1.339 for those in decile 10. AAPL with its book-to-market ratio of 0.2 belongs only to decile 2. Conceptually, we can say that stocks with low book-to-market ratios are those with great growth potential. As such, investors are willing to pay multiple (in the case of 0.095, 10 times) of the book value. Quantitatively, however, why some stocks are priced at 10 times while other stocks are priced at 0.75 times? Do we have one good model to give accurate prices to this cross-section of stocks with varying book-to-market ratios?

By now, you've probably been taught various valuation models that combine cash flows with discount rates. You project the future cash flows of a firm or a project and discount them back using some discount rates estimated using a Finance model, say the CAPM. Without a question, these frameworks are useful in helping us think through the key components in a valuation project. But, quantitatively, these models do not offer the kind of precision and rigor as other models in Finance. And in practice, this seems to be true as well.

When I first read the fascinating book on the RJR Nabisco deal, "Barbarians at the Gate," my mouth was wide open as I flipped through the pages. For such a large deal, the valuation seems to be rather flexible. Over the short time span of one month and 11 days, the valuation moved from the initial $17 billion with $75 a share to $24.88 billion with $109 a share. This might be an extreme case, but other books on private equity, for example, "King of Capital," left me with the same impression: there is a lot of flexibility in valuation in this space. If you look at the venture capital space, where even the projection of future cash flow is very much up in the air, you see a similar pattern. I am not an expert in either of these areas, but it is safe to say that the level of precision required of a Finance model is relatively low in these areas, or the margin of errors allowed for such valuation models is rather high.

In writing this introduction on valuation, my objective was to compare and contrast the role of valuation in various parts of Finance. On the one end of the spectrum, you have valuations in VC and private equity. In this space, the cash flows are highly uncertain; which discount rates to use is also not clear. The role of a valuation model in such a setting is indeed very limited. If you are working in this area, spending time to perfect your Finance model is not at all your number one priority. On the other end of the spectrum, you have valuations in options and fixed income. In options, the cash flow comes from the fluctuation of the underlying stock price. In fixed income, the cash flow comes from the coupon and principal payments. In both cases, the cash flow can be modeled rather precisely and the present value calculation can be done with super high precision. In these areas, people take their valuation models rather seriously. If anything, the danger is that people take their models too literally to the extent that they are lost in their models.

I hope that you do not read this introduction as "one against another." This concern made me move this introduction to the appendix so as not to distract you from the main topic. The role of a professor is to offer knowledge and perspective. As a student, your responsibility is to absorb the useful, discard the useless and build a system for yourself. I can see how a teacher can influence his students (OK, maybe not MBAs). My cousin in Shanghai used to hate English because her English teacher was not nice to her. Isn't that crazy?

B The Motives for Option Trading

The motives behind options trading could vary from speculation to hedging. Investors with private (legal or illegal) information might choose to trade in the options market to take advantage of the inherent leverage in options. This usually happens more at the level of options on individual stocks, where option investors trade their private information about the idiosyncratic component of the stocks. I have a paper with Allen Poteshman on this topic.

As a graduate from Chicago GSB, Allen was able to get a very unique dataset from CBOE with details on option trading volumes on open buy and sell, close buy and sell from 1990 through 2001. Around the same time, I was teaching 15.433 and had to educate myself about quant investing and sorting portfolios with signals. Like some of you, after learning about this cross-sectional approach, I started to think about trading strategies. Since I spent most of my time thinking about options, the idea came quite naturally to me: would it be cool to have a signal from the options market and use it to trade in the stock market? The most obvious signal would be put/call ratio. Consider a stock with a lot of put option volume traded on it versus a stock with a lot of call option volume traded on it. One is a bearish signal on the stock and the other bullish.

My problem was that I did not have good options data with clean volume information to test this idea. Most of the publicly available data mixes open buy with close buy and open sell with close sell. As a result, the pure signal from open buy is contaminated by close buy. Likewise for the sell volume. So my test results using the publicly available data were weak and I did not want to write a paper with these weak results. This is how I located Allen and his unique dataset. I sent him an email, he sent me a disc with his data and we started to work together.

We form stock portfolios by their put/call ratios and track their performance for the next week. We find that stocks with low put/call ratio outperform stocks with high put/call ratio by 40 basis points over the next day and 1% over the next week. This predictability is stronger for smaller stocks. We also find that option volumes by customers from full service brokerage firms (e.g., hedge funds) are by far the most informative. By contrast, option volumes by firm proprietary traders do not have any predictive power. Our interpretation is that prop traders use exchange-traded options mostly for hedging needs, which is supported by the fact that prop traders are much more active on index options than equity options.

After we finished our paper in 2003 or 2004, we got a lot of interest from practitioners. We even heard from CBOE, who asked us where we got the data. We told them that the data was sitting in their mainframe and offered to help them package and sell the data (so that we can have free access). They said "No." Later they started to sell this data at a pretty high price. Several years later around 2009 or 2010, a former student of Allen got this big grant for data purchase. So I asked her to buy the very expensive CBOE data to do the same test on the more recent data. The strong predictability we found over the 1990-2001 sample no longer exists in the recent time period. It is difficult to say if our paper has any direct impact, but the market seemed to become a little more efficient. After writing this paper, Allen became more interested in the practice of Finance and left his tenured professorship to join D.E. Shaw. He also got us a coverage on the New York Times.

C Brownian Motion

During my office hours, I got a few questions about the Brownian motion and risk-neutral pricing. Let me use this appendix to explain some of the details.

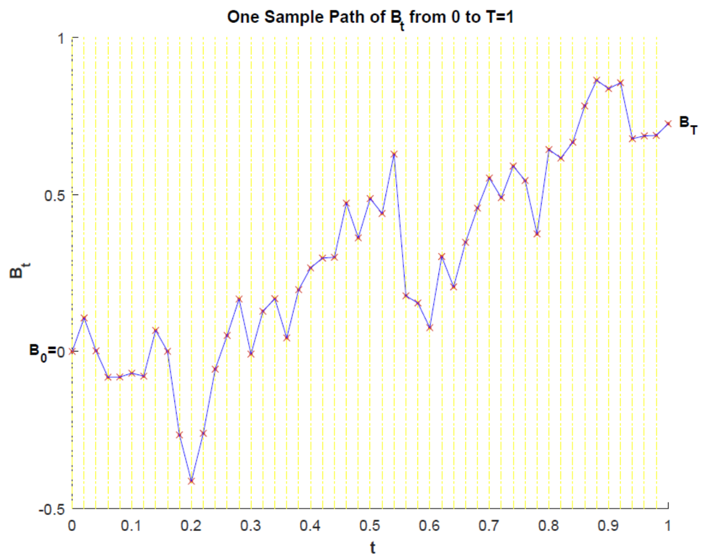

Figure 10: One sample path of a Brownian motion.

To understand the Brownian motion, let's create one. Let's start from time 0 and end in time T. Let's further chop this time interval into small increments. For example, in Figure 10, T=1 and the interval between 0 and 1 is chopped evenly into 50 smaller increments with size $\Delta = 1/50$. We can now start to create a sample path of the Brownian motion:

$$\begin{aligned}
&B_0 = 0 \\
&B_\Delta - B_0 = \sqrt{\Delta}\, \epsilon_\Delta \\
&B_{2\Delta} - B_\Delta = \sqrt{\Delta}\, \epsilon_{2\Delta} \\
&\dots \\
&B_T - B_{T-\Delta} = \sqrt{\Delta}\, \epsilon_T \,.
\end{aligned}$$

where the ϵ's are independent standard normals. In creating this sample path, we use the first two properties of the Brownian motions: independence increments and stationary normal increments. I've attached the Matlab code I used to create this plot in this note. You can run it and each time you will get a different sample path.

For our purpose of pricing a European-style option, what matters is the distribution of

B_T. But if we are interested in pricing an American-style option, then the entire path of B_t matters and at each node, we will make a decision of whether or not to exercise early. So the grid should be as fine as possible (larger N and smaller Δ).

Now back to our original process for S_t:

$$dS_t = \mu\, S_t\, dt + \sigma\, S_t\, dB_t\,,$$

where to avoid distraction, I have set the dividend yield $q = 0$. As usual, we work with $X_t = \ln S_t$ and, using the Ito's Lemma, we have

$$dX_t = \left(\mu - \frac{1}{2}\sigma^2\right)\, dt + \sigma\, dB_t\,.$$

The nice thing about working with $\ln S_t$ is that you can integrate out the process:

$$\begin{aligned}
X_T &= X_0 + \int_0^T \left(\mu - \frac{1}{2}\sigma^2\right)\, dt + \int_0^T \sigma\, dB_t \\
&= X_0 + \left(\mu - \frac{1}{2}\sigma^2\right) T + \sigma\, \left(B_T - B_0\right) \\
&= X_0 + \left(\mu - \frac{1}{2}\sigma^2\right) T + \sigma\, \sqrt{T}\, \epsilon_T\,,
\end{aligned}$$

where in the last step I use $\sqrt{T}\, \epsilon_T$ to express $B_T - B_0$. Recall that the log-return R_T is defined by $R_T = \ln X_T - \ln X_0$. We have

$$R_T = \left(\mu - \frac{1}{2}\sigma^2\right) T + \sigma\, \sqrt{T}\, \epsilon_T$$

D Change of Measure, Risk-Neutral Pricing

Under the original measure (P-measure), the process runs as

$$dX_t = \left(\mu - \frac{1}{2}\sigma^2\right)\, dt + \sigma\, dB_t\,.$$

If we were to do option pricing under this measure, we know that we cannot do

$$C_0 \neq e^{-rT} E\left(S_T - K\right)^+\,.$$

This is a big no-no in Finance because it approaches the pricing as if we were risk neutral. Interestingly, this is why the method of "risk-neutral" pricing arises. It is mostly a mathematical result. If you Google Girsonov theorem or the Radon-Nikodym derivative, you will see the related math result. But the math result has its relevance in Finance. Let me approach it this way.

In Finance, we develop this concept of pricing kernel or the stochastic discount factor. Armed with this pricing kernel ξ_T, we can do our pricing:

$$C_0 = e^{-rT} E\left(\frac{\xi_T}{\xi_0}\left(S_T - K\right)^+\right) .$$

Under the Black-Scholes setting, the markets are complete and the pricing kernel is unique. In fact, as an application of the Girsonov theorem, this pricing kernel is of the form

$$\xi_T = \frac{dQ}{dP} = e^{-\gamma\, B_T - \frac{1}{2}\gamma^2\, T} .$$

This ξ_T is what the mathematician would call the Radon-Nikodym derivative. Notice that by construction $E(\xi_T) = 1$.

As mentioned earlier, the pricing kernel is unique under the Black-Scholes setting. So the constant γ is uniquely defined. In Finance, we call this parameter the market price of risk and for the Black-Scholes setting, it is $\gamma = (\mu - r)/\sigma$, which in fact is the Sharpe ratio. In a more general setting, γ can itself be a stochastic process. Also notice that with a positive market price of risk, $\gamma > 0$, ξ_T is negatively correlated with B_T (hence negatively correlated with X_T and S_T). This is what you were taught in Finance 15.415. When S_T experiences a positive stock, the stochastic discount factor is smaller; when S_T experiences a negative stock, the stochastic discount factor is bigger. This asymmetry has its origin in the fact that investors are risk averse and the risk in S_T is systematic (undiversifiable).

It turns out that we can create a new measure Q, called the equivalent martingale measure, for the original P and the pricing becomes,

$$\begin{aligned} C_0 &= e^{-rT} E\left(\frac{\xi_T}{\xi_0}\left(S_T - K\right)^+\right) \\ &= e^{-rT} E^Q\left(\left(S_T - K\right)^+\right) , \end{aligned}$$

and the link between these two measures is $\xi_T = dQ/dP$.

Now let's construct this new Q-Brownian:

$$\begin{aligned} dX_t &= \left(\mu - \frac{1}{2}\sigma^2\right) dt + \sigma\, dB_t^p \\ &= \left(r - \frac{1}{2}\sigma^2\right) dt + \sigma \left(\frac{\mu - r}{\sigma} + dB_t^P\right) \\ &= \left(r - \frac{1}{2}\sigma^2\right) dt + \sigma\, dB_t^Q\,, \end{aligned}$$

where the Q-Brownian is defined as

$$dB_t^Q = \frac{\mu - r}{\sigma} + dB_t^P\,.$$

And this change of measure, from P to Q, is the essence of the risk-neutral pricing.

The name of "risk-neutral" pricing is ironical: the whole thing arises from the observation that we cannot do

$$C_0 \neq e^{-rT} E^P \left(S_T - K\right)^+\,.$$

But if we are willing to change our probability measure from P to Q, under which

$$dX_t = \left(r - \frac{1}{2}\sigma^2\right) dt + \sigma dB_t^Q\,,$$

then we can indeed do

$$C_0 = e^{-rT} E^Q \left(S_T - K\right)^+\,.$$

E Change of Measure, One More Application

This mathematical tool can be further exploited. Recall that we need to do this calculation in our Black-Scholes option pricing,

$$e^{-rT}\, E^Q \left(S_T \mathbf{1}_{S_T > K}\right)$$

What if we can drop S_T and change it to

$$S_0\, E^{?} \left(\mathbf{1}_{S_T > K}\right)$$

That would make our math very simple.

In fact, we can drop S_T like the way we dropped ξ_T. As long as the process is positive, there is an equivalent martingle measure waiting for us to help us simplify the math. This is where the new measure QQ comes from. You can start with the observation that

$$S_T = e^{X_T} = e^{\sigma\, B_T + \text{other deterministic terms}}$$

You can then check

$$\begin{aligned} dX_t &= \left(r - \frac{1}{2}\,\sigma^2\right)\, dt + \sigma dB_t^Q \\ &= \left(r + \frac{1}{2}\,\sigma^2\right)\, dt - \sigma^2\, dt + \sigma dB_t^Q \\ &= \left(r + \frac{1}{2}\,\sigma^2\right)\, dt + \sigma\, \left(-\sigma\, dt + dB_t^Q\right) \end{aligned}$$

So if we define

$$dB_t^{QQ} = -\sigma\, dt + dB_t^Q\,,$$

under which

$$dX_t = \left(r + \frac{1}{2}\,\sigma^2\right)\, dt + \sigma dB_t^{QQ}\,.$$

Then we can indeed get

$$e^{-rT}\, E^Q\left(S_T \mathbf{1}_{S_T > K}\right) = S_0 E^{QQ}\left(\mathbf{1}_{S_T > K}\right)\,.$$

I am being a bit sloppy in my notation, but I trust a careful and thorough student would fill in the details (including the *other deterministic terms*).