FAQ's in Option Pricing Theory

Presentation by
Peter Carr
Banc of America Securities
pcarr@bofasecurities.com

Initial version: June 7, 1999

File reference: faqpres.tex

Introduction

- Newcomers to derivatives often pose certain frequently asked questions (FAQ's), which the author has struggled with for many years.
- These same FAQ's are often posed to newcomers during interviews.
- Here is a list of FAQ's asked and hopefully answered in a paper downloadable from www.math.nyu.edu/research/carrp/papers:
  1. Why isn't an option's value just its discounted average payout?
  2. Why don't the statistical probabilities matter in the binomial model?
  3. Why doesn't the expected rate of return appear in the Black Scholes formula?
  4. Is the hedging argument given in the Black Scholes paper correct? Why doesn't one differentiate the number of shares/options held in the hedge portfolio?
  5. Can one hedge options in a trinomial model?
  6. Can jumps be hedged in a continuous time model?
  7. Why does the "market price of risk" appear in many stochastic interest rate or volatility models? Why doesn't it appear in the HJM model?
  8. How much do I lose if I hedge at the wrong volatility?

I Why Isn't an Option's Value Just Its Discounted Average Payout?

- Suppose that a non-dividend paying stock is just as likely to double in price from \$1 to \$2 as it is to halve from \$1 to 50¢:

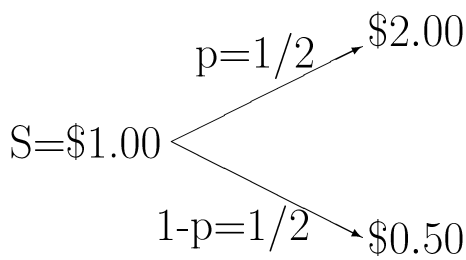

- Under zero interest rates, suppose we value an at-the-money call at its expected payoff of 50¢:

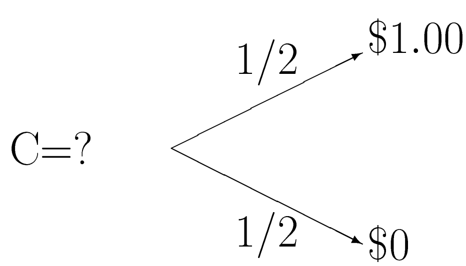

$$\Longrightarrow C = \$0.50$$

A is for Arbitrage

- Suppose we make a 2 way market in the stock at its market price of a buck and in the at-the-money call at its estimated value of 50¢:

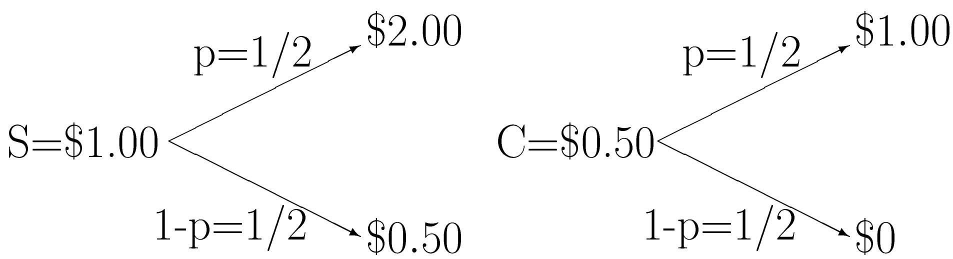

- Immediately afterwards, a hedge fund sells 2 calls to us and buys a share from us:

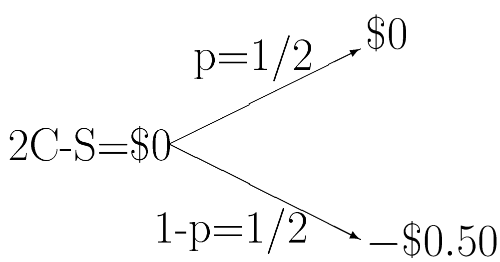

- We sold the hedge fund a free at-the-money put.

What Went Wrong?

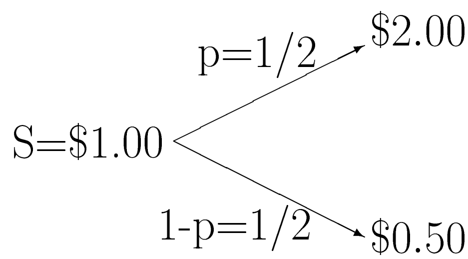

- Under zero interest rates, risk aversion (fear) has caused the stock to be priced below its expected value of $1.25.
- Since a call is even riskier than a stock, shouldn't it be priced even further below its expected payoff of 50¢:

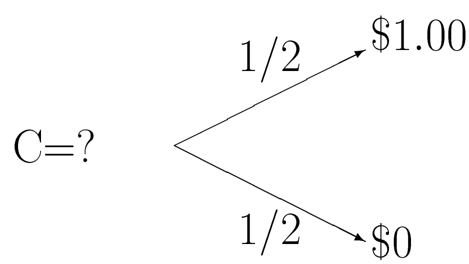

Call Pricing 101

- Recall that when we were long two calls and short one share, we were short a synthetic at-the-money put:

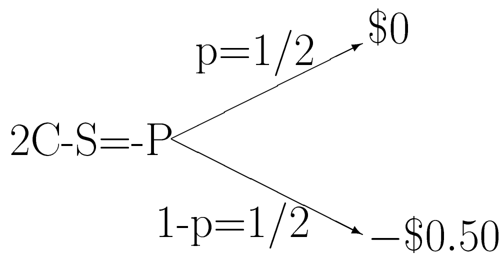

- Put-call parity implies that the synthetic at-the-money put should have the same price as the at-the-money call:

$$2C - S = -C.$$

- Since $S = 1$, the arb-free call value is $33\frac{1}{3}$¢:

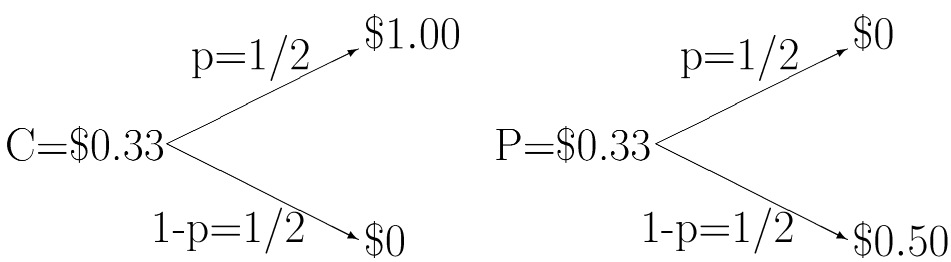

Time and State Value of Money

- Time Value of Money says that under positive interest rates, \$1 paid now is worth more than \$1 paid in a year because by depositing \$1 in the bank today, one gets more than a dollar (say \$1.10) in a year.
- Similarly, "State Value of Money" says a bad state dollar is worth more than a good state dollar because by trading today, one gets more than one good state dollar in a year. Recall the at-the-money call and put values:

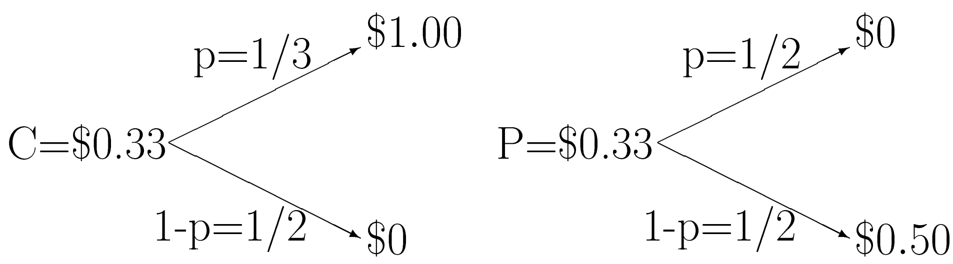

- Suppose we presently own an at-the-money put giving us .50 bad state dollars. By selling the put and buying the call, we can instead get 1 good state dollar. Whether this is a good trade or not depends on our risk aversion.

II Why Don't Statistical Probabilities Matter in the Binomial Model?

- Again consider valuing an at-the-money call on a stock priced at a dollar which can double or halve over a single period.
- Since neither the stock nor the call were priced at expected value, perhaps probabilities don't matter.
- Let's now suppose we don't know the probabilities of the two states.

No Hockey Sticks!

- The graph below shows the 2 possible payouts of the at-the-money call:

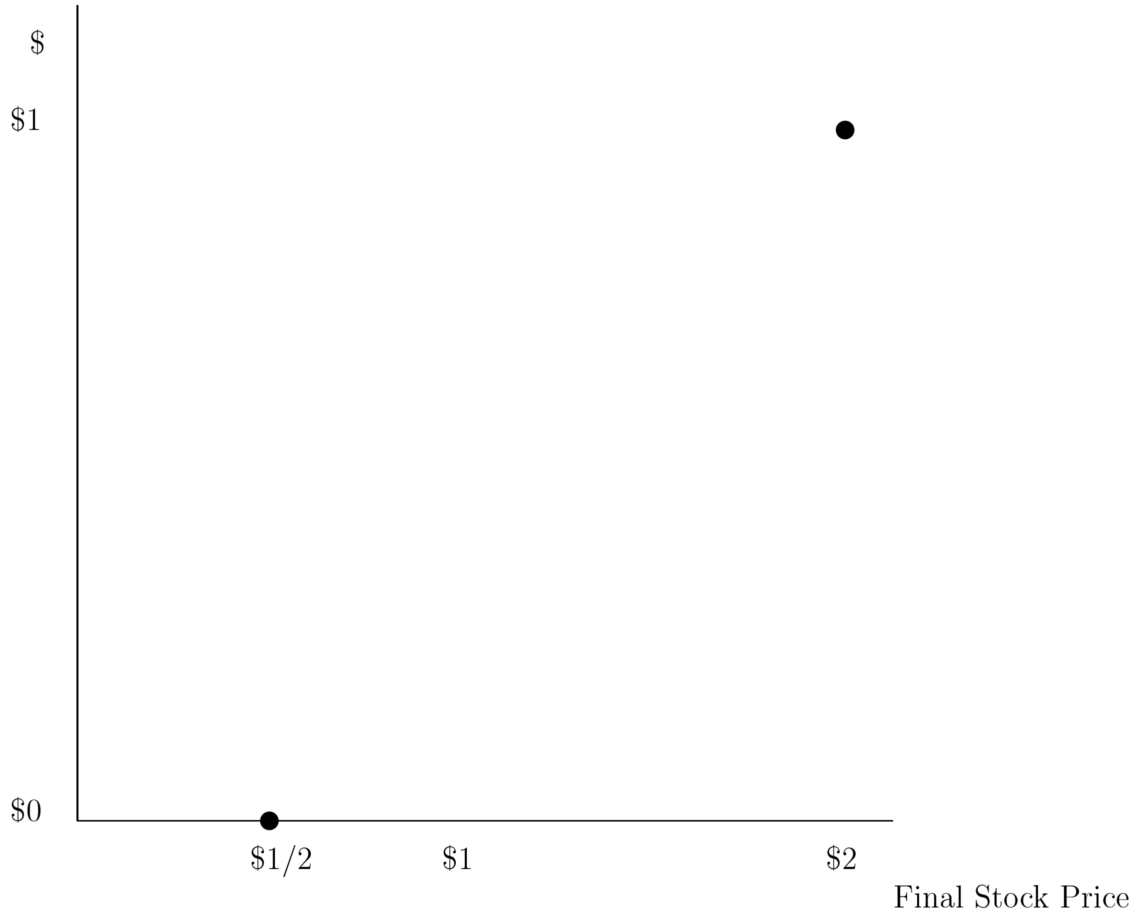

Figure 1: Payoffs in Binomial Model.

- If we draw a line through the two points, then the slope is 2/3 so the vertical intercept is -1/3.
- Since only two points on the call's hockey stick payoff are considered possible, one is indifferent between the linear payoff and the nonlinear payoff. However, the linear payoff is easily created by buying 2/3 of a share and borrowing 1/3 of a unit bond.
- It follows that the value of the option is the cost of this replicating portfolio, which is independent of the probabilities of up and down.

Do Probabilities Matter?

- To be clear, if market probability assessments change, then the stock and bond prices may change, and if they do, then option prices will almost surely change.
- However, the functional relationship between the option price and the stock and bond prices will remain independent of the probabilities.
- In essence, these prices are sufficient statistics for the statistical probabilities.
- In this model, probabilities do matter for prices; they don't matter for pricing.

Local Linearity

- Now consider a two period recombining binomial tree as shown below:

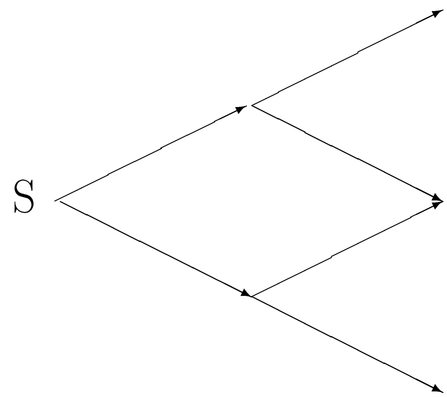

- Although an at-the-money call maturing at the end of the second period has a nonlinear payoff, at each intermediate node, the payoff is *locally* linear.
- Thus terminal values can be propagated back one period to generate intermediate values.
- If these intermediate values are regarded as the payoffs of a single period claim, this payoff is again locally linear when viewed from the root node and hence susceptible to valuation by replication.
- **ALL** arbitrage-free pricing models exploit this local linearity principle.

III Why Doesn't the Expected Rate of Return Appear in the Black Scholes Formula?

- The most important feature of the Black Scholes/Merton (BSM) model is that the expected return doesn't matter for pricing (but it does matter for prices).
- This is usually attributed to the fact that probabilities don't matter for pricing in the binomial model, and that the binomial model converges to the BSM model.
- Why then does volatility matter for pricing in the Black Scholes model? A re-assignment of probabilities in the binomial model *will* cause volatility to change. However, if the stock and bond prices don't change, then neither will the option price.
- This paradox suggests that the usual explanation isn't the whole story.

The Linearity Principle in Continuous Time

- We now attempt to explain *why* the stock's expected return over the option's life $[0, T]$ is irrelevant for valuing an option at some fixed future time $t \in [0, T]$, *given* the stock price at t.
- Our explanation relies heavily on the following continuous time version of the:
- **Linearity Principle:** *Consider a derivative security whose payoff at T is linear in the price of a stock:*

$$D_T = N_0 + N_1 S_T,$$

where N_0 and N_1 are known at some prior time $t < T$. For simplicity, suppose a constant interest rate r and a constant dividend yield of δ. Then the absence of arbitrage requires that the time t price of the derivative must be:

$$D_t = N_0 e^{-r(T-t)} + N_1 e^{-\delta(T-t)} S_t, \qquad t < T.$$

- The proof actually is obvious.
- Given the values of $N_0, e^{-r(T-t)}, N_1$, and $e^{-\delta(T-t)} S_t$, one does not need to know *anything* about the parameters governing changes in the price.
- If the price follows a diffusion, this means that one does not need to know the expected rate of return *or the volatility*, given the values of N_0, $e^{-r(T-t)}$, N_1, and $e^{-\delta(T-t)} S_t$.
- However, the requirement that the payoff be linear appears to be quite restrictive...

Local Linearity in Continuous Time

- By assuming continuity in time of both the price process and trading opportunities, BSM were able to apply the Linearity Principle in a situation where the derivative's future price D_T is apparently non-linear in the stock's future price S_T.
- BSM further showed that when given time t information such as the stock price S_t, then the intercept N_0 and the slope N_1 depend on the stock's volatility, but not on its expected rate of return.
- Under the continuity assumptions, BSM realized that when viewed very closely, nonlinear functions can be treated as "locally linear". (Pass the Nobel Please).

A Serious Lemma

- To show that the future value of any derivative is locally linear in the BSM model, suppose that the stock price process is:

$$dS_t = [\mu(S_t, t) - \delta] S_t dt + \sigma(S_t, t) S_t dW_t, \qquad t \in [0, T].$$

- Assuming that the time t value of the derivative D_t is a sufficiently smooth function V of only S_t and t, a bivariate Taylor series expansion implies:

$$\begin{aligned}
&\triangle D_t \equiv V(S_t + \triangle S, t + \triangle t) - V(S, t) \\
&= \frac{\partial V}{\partial t}(S_t, t)\triangle t + O(\triangle t)^2 \\
&+ \frac{\partial V}{\partial S}(S_t, t)\triangle S_t + \frac{1}{2}\frac{\partial^2 V}{\partial S^2}(\triangle S_t)^2 + O(\triangle S)^3 + O((\triangle S)^2 \triangle t), t \in [0, T].
\end{aligned}$$

- If one substitutes $\triangle S = \int_t^{t+\triangle t} [\mu(S_u, u) - \delta] S_u du + \sigma(S_u, u) S_u dW_u$, then μ appears in all four terms in the last row.

- However, if we focus on an infinitesimally small time step, then Itô's lemma implies:

$$dD_t = \frac{\partial V}{\partial t}(S_t, t)dt + \frac{\partial V}{\partial S}(S_t, t)dS_t + \frac{1}{2}\frac{\partial^2 V}{\partial S^2}(dS_t)^2, \qquad t \in [0, T].$$

- Since the higher order terms have disappeared, so has the appearance of μ in them. However, substitution of the top equation in the bottom one shows that μ still affects the *change in* the value.

A Shell Game

- Recall the application of Itô's lemma in explaining the change in a derivative security's price:

$$dD_t = \frac{\partial V}{\partial t}(S_t, t)dt + \frac{\partial V}{\partial S}(S_t, t)dS_t + \frac{1}{2}\frac{\partial^2 V}{\partial S^2}(dS_t)^2, \qquad t \in [0, T].$$

- The future value of the derivative D_{t+dt} is apparently *quadratic* in the future value of the stock S_{t+dt}:

$$D_{t+dt} - D_t = \frac{\partial V}{\partial t}(S_t, t)dt + \frac{\partial V}{\partial S}(S_t, t)(S_{t+dt} - S_t) + \frac{1}{2}\frac{\partial^2 V}{\partial S^2}(S_t, t)(S_{t+dt} - S_t)^2, \qquad t$$

- However, Itô's recognition that $(S_{t+dt} - S_t)^2 = \sigma^2(S_t, t)S_t^2 dt$ linearizes this equation:

$$\begin{aligned} D_{t+dt} &= \underbrace{D_t + \left[\frac{\partial V}{\partial t}(S_t, t) + \frac{\sigma^2(S_t, t)S_t^2}{2}\frac{\partial^2 V}{\partial S^2}(S_t, t)\right] dt - \frac{\partial V}{\partial S}(S_t, t)S_t}_{N_0} \\ &\quad + \underbrace{\frac{\partial V}{\partial S}(S_t, t)}_{N_1} S_{t+dt}, \qquad t \in [0, T]. \end{aligned}$$

LP Implies PDE

- Recall that the future value of the derivative D_{t+dt} is actually linear in the future value of the stock S_{t+dt}:

$$D_{t+dt} = \underbrace{D_t + \left[\frac{\partial V}{\partial t}(S_t,t) + \frac{\sigma^2(S_t,t)S_t^2}{2}\frac{\partial^2 V}{\partial S^2}(S_t,t)\right]dt - \frac{\partial V}{\partial S}(S_t,t)S_t}_{N_0} + \underbrace{\frac{\partial V}{\partial S}(S_t,t)}_{N_1} S_{t+dt}, \qquad t \in [0,T].$$

- If one replaces S_{t+dt} with $S_t + \mu(S_t,t)S_t dt + \sigma(S_t,t)dW_t$, then μ still appears.

- However, applying the Linearity Principle:

$$D_t = \left\{D_t + \left[\frac{\partial V}{\partial t}(S_t,t) + \frac{\sigma^2(S_t,t)S_t^2}{2}\frac{\partial^2 V}{\partial S^2}(S_t,t)\right]dt - \frac{\partial V}{\partial S}(S_t,t)S_t\right\}e^{-rdt} + \frac{\partial V}{\partial S}(S_t,t)S_t e^{-\delta dt}, \qquad t \in [0,T].$$

- Replacing e^{-rdt} with $1 - rdt$, $e^{-\delta dt}$ with $1 - \delta dt$, and eliminating terms of order $(dt)^2$ leaves the Black Scholes p.d.e.:

$$\frac{\partial V}{\partial t}(S_t,t) + \frac{\sigma^2(S_t,t)S_t^2}{2}\frac{\partial^2 V}{\partial S^2}(S_t,t) + (r-\delta)S_t\frac{\partial V}{\partial S}(S_t,t) - rV(S_t,t) = 0.$$

From PDE to μ Independence

- Recall that the Linearity Principle lead to the BSM PDE:

$$\frac{\partial V}{\partial t}(S_t,t)+\frac{\sigma^2(S_t,t)S_t^2}{2}\frac{\partial^2 V}{\partial S^2}(S_t,t)+(r-\delta)S_t\frac{\partial V}{\partial S}(S_t,t)-rV(S_t,t)=0.$$

- Solving this p.d.e. subject to the terminal condition and the appropriate boundary conditions allows us to determine $V(S,t)$ for all $t\in[0,T]$.

- Since μ does not appear in the PDE or the boundary conditions, D_t is completely independent of μ, given the information at t such as S_t.

- Also recall the linear decomposition of the future value of the derivative:

$$\begin{aligned} D_{t+dt} &= \underbrace{D_t+\left[\frac{\partial V}{\partial t}(S_t,t)+\frac{\sigma^2(S_t,t)S_t^2}{2}\frac{\partial^2 V}{\partial S^2}(S_t,t)\right]dt-\frac{\partial V}{\partial S}(S_t,t)S_t}_{N_0} \\ &\quad +\underbrace{\frac{\partial V}{\partial S}(S_t,t)}_{N_1}S_{t+dt}, \qquad t\in[0,T]. \end{aligned}$$

- Thus, when replicating the payoff, the bank balance at $t+dt$ is:

$$\begin{aligned} N_0(S,t+dt) &= V(S_t,t)-\frac{\partial V}{\partial S}(S_t,t)S_t+\left[\frac{\partial V}{\partial t}(S_t,t)+\frac{\sigma^2(S_t,t)S_t^2}{2}\frac{\partial^2 V}{\partial S^2}(S_t,t)\right. \\ &= V(S_t,t)-\frac{\partial V}{\partial S}(S_t,t)S_t+r\left[V(S_t,t)-S_t\frac{\partial V}{\partial S}(S_t,t)\right]dt+\delta S_t \end{aligned}$$

while the number of shares held from t to $t+dt$ is:

$$N_1(S_t,t)=\frac{\partial V}{\partial S}(S_t,t), \qquad t\in[0,T].$$

- It follows that N_0 and N_1 are also independent of μ, given the information at t such as S_t.

Why μ Disappears

- Recall the Taylor series expansion:

$$
\begin{aligned}
& \triangle D_t \equiv V(S_t + \triangle S, t + \triangle t) - V(S, t) \\
= \; & \frac{\partial V}{\partial t}(S_t, t)\triangle t + O(\triangle t)^2 \\
+ \; & \frac{\partial V}{\partial S}(S_t, t)\triangle S_t + \frac{1}{2}\frac{\partial^2 V}{\partial S^2}(\triangle S_t)^2 + O(\triangle S)^3 + O((\triangle S)^2 \triangle t), t \in [0, T].
\end{aligned}
$$

- The μ appearing in the $\triangle S_t$ term disappears for a reason completely different from the reason that the μ appearing in $(\triangle S_t)^2$ and higher powers disappears.

- The appearance of μ in the $\triangle S_t$ term is handled by the Linearity Principle, which has nothing to do with continuous trading or a continuous price process. In other words, given that the future value of the derivative at $t + \triangle t$ is linear in the future stock price $S_{t+\triangle t}$ for any $\triangle t > 0$, the μ appearing in $\triangle S_t \equiv S_{t+\triangle t} - S_t$ has no effect on the value (and neither does volatility).

- In contrast, the appearance of μ in $(\triangle S_t)^2$ is handled by Itô's recognition that the distinction between variation about zero and variation about the mean vanishes as $\triangle t \downarrow 0$. It is interesting to note that volatility only matters for value because stochastic calculus is being used rather than ordinary calculus.

- Similarly, the appearance of μ in terms like $(\triangle S_t)^3$ is handled by the result of stochastic (and ordinary) calculus that such terms are negligible in comparison to the leading terms.

IV Is the Hedging Argument Given in the Black-Scholes Paper Correct?

- The textbook derivation of the Black Scholes hedging argument assumes frictionless markets, no arbitrage, a constant interest rate r, no dividends, and that the stock price process $\{S_t, t \in [0,T]\}$ is GBM:

$$\frac{dS_t}{S_t} = \mu_t dt + \sigma dW_t, \qquad t \in [0,T].$$

- Letting $\{C_t, t \in [0,T]\}$ denote the European call price process, further assume $C_t = C(S_t, t), t \in [0,T]$ for some $C^{2,1}$ function $C(S,t)$. Then by Itô's lemma:

$$dC_t = \left[\frac{\partial C}{\partial t}(S_t,t) + \frac{\sigma^2 S_t^2}{2}\frac{\partial^2 C}{\partial S^2}(S_t,t)\right] dt + \frac{\partial C}{\partial S}(S_t,t) dS_t, \quad t \in [0,T].$$

- Consider a hedge portfolio consisting of long one call and short $\frac{\partial C}{\partial S}(S,t)$ shares:

$$H_t \equiv C_t - \frac{\partial C}{\partial S} S_t.$$

- The textbook argument is that:

$$\begin{aligned} dH_t &= dC_t - \frac{\partial C}{\partial S} dS_t \qquad (***) \\ &= \left[\frac{\partial C}{\partial t}(S_t,t) + \frac{\sigma^2 S_t^2}{2}\frac{\partial^2 C}{\partial S^2}(S_t,t)\right] dt \\ &= r\left[C(S_t,t) - \frac{\partial C}{\partial S}(S_t,t) S_t\right] dt, \end{aligned}$$

by the absence of arbitrage. Equating coefficients on dt yields the Black Scholes PDE:

$$\frac{\partial C}{\partial t}(S,t) + \frac{\sigma^2 S^2}{2}\frac{\partial^2 C}{\partial S^2}(S,t) + r\frac{\partial C}{\partial S}(S,t)S - rC(S_t,t) = 0, t \in (0,T), S > 0.$$

- Is (***) correct?

Gains Process

- Recall the textbook argument that if:

$$H_t \equiv C_t - \frac{\partial C}{\partial S} S_t, \qquad t \in [0, T]$$

then:

$$dH_t = dC_t - \frac{\partial C}{\partial S} dS_t, \qquad t \in [0, T].$$

- However, integration by parts requires that the second equation should instead be:

$$dH_t = dC_t - \frac{\partial C}{\partial S} dS_t - d(\frac{\partial C}{\partial S}) S_t - d\langle \frac{\partial C}{\partial S}, S \rangle_t, \quad t \in [0, T].$$

- The additional two terms represent the additional investment needed to maintain the strategy. These terms are differentials of processes of unbounded variation, so one cannot argue that dH_t is riskless.

- By doing the math right, we can't get the PDE. Could BSM be wrong?

- Many authors argue that the result is right, but the derivation is wrong. They then derive the PDE by more complicated means, achieve tenure, and make the world safe from derivatives.

- Others argue that the derivation is right since the number of shares held is "instantaneously constant". To a mathematician, this argument must be perplexing since the total variation of the number of shares held in any finite time interval is in fact infinite. In fact, the number of shares is changing so fast that the ordinary rules of calculus do not apply.

- Is the derivation right?

Yes!(Almost)

- I believe the BSM derivation is right up to a typo. If the hedge portfolio value at time t is given by:

$$H_t \equiv C_t - \frac{\partial C}{\partial S} S_t, \qquad t \in [0, T],$$

then the *gain* gH_t on the hedge portfolio defined by:

$$\begin{aligned} gH_t &\equiv dC_t - \frac{\partial C}{\partial S} dS_t \\ &= \left[\frac{\partial C}{\partial t}(S_t, t) + \frac{\sigma^2 S_t^2}{2}\frac{\partial^2 C}{\partial S^2}(S_t, t)\right] dt \\ &= r[C(S_t, t) - \frac{\partial C}{\partial S}(S_t, t) S_t] dt, \end{aligned}$$

is riskless and thus should grow at the riskfree rate to annihilate arbitrage.

- Equating coefficients on dt again yields the Black Scholes PDE.

- To make the world safe for derivatives, the above argument replaces the mathematical operation of computing a total derivative with the financial operation of defining a gain.

- It is worth noting that the portfolio consisting of the option and stock is *not* self-financing. Similarly, positions in the riskless asset are not self-financing. Nonetheless, by showing that the gains between two non-self-financing strategies are always equal under no arbitrage, the value of the derivative security can be determined.

V Can One Hedge Options in a Trinomial Model?

- The binomial model of Rendleman and Bartter(1979) and Cox, Ross, and Rubinstein(1979) is a pedagogical marvel.
- An even more natural model is the trinomial model which captures the simple idea that the price of the asset can either move up, move down, or stay the same:

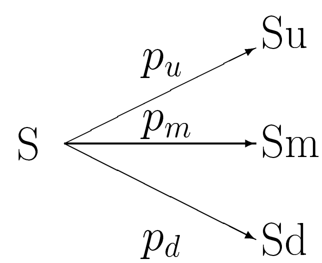

- Trinomial models are widely used in practice since they converge faster than binomial models (i.e. take less time to achieve a given accuracy).
- The extra degree of freedom is also handy in allowing one to place nodes on barriers.
- The only problem is that one can't hedge options using dynamic trading in just a stock and a bond. So how do we price?

Three is Not a Crowd

- Just as two assets (eg. stock and bond) can be used to hedge a derivative when there are two states, three assets can be used to hedge a derivative when there are three states.

- This is most obvious in a single period setting. Suppose that interest rates are zero and that a stock can double, halve, or remain at its initial price of \$2. To hedge a call struck at \$3, consider using the bond, the stock, and an at-the-money call with initial price $\frac{2}{3}$:

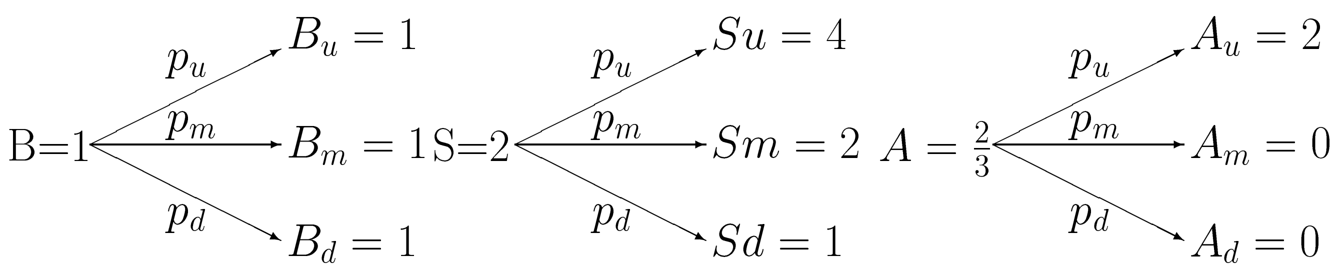

- Let N^b, N^s, and N^a denote the number of bonds, shares, and at-the-money calls held in the hedge portfolio. Equating portfolio values to desired payoffs in the 3 states gives 3 linear equations in the 3 unknowns:

$$\begin{aligned} \text{up} \quad N^b 1 + N^s 4 + N^a 2 &= \; 1 \\ \text{same} \quad N^b 1 + N^s 2 + N^a 0 &= \; 0 \\ \text{down} \quad N^b 1 + N^s 1 + N^a 0 &= \; 0. \end{aligned}$$

- The unique solution to this system is $N^b = 0$, $N^s = 0$, and $N^a = \frac{1}{2}$. Thus, given that the initial price of the at-the-money call is $\$\frac{2}{3}$, the arbitrage-free value of the call struck at \$3 is $\$\frac{1}{3}$.

But Two Dates Can Be Trouble

- Now consider a two period recombining trinomial tree as shown below:

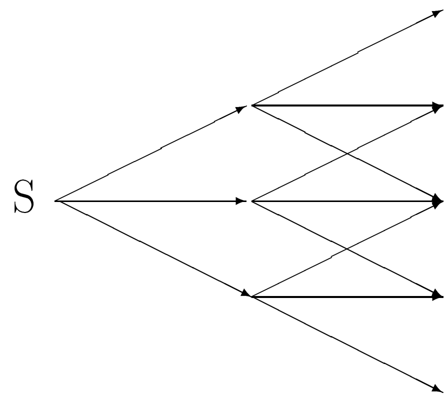

- Assume for simplicity that the gross periodic stock returns $u > 1$, $m = 1$, and $d = 1/u$ are constant over time and that interest rates and dividend yields are zero.

- If we mimic the backward recursion used in the binomial model, we quickly face a hurdle. Imagine placing ourselves at the upper node and the middle date, i.e. when the stock price has risen over the first period. Then since only 3 possible prices yawn before us, the situation appears to be identical to the single period problem.

- However, there is a subtle difference. In the single period problem, the price of the at-the-money was known to us and used to derive the value of a call with another strike. In the two period case, the time 1 price of any option is unknown because the only known prices are those at the root of the tree. At the middle date upper node, there are 3 possible future stock prices but only two known time 1 prices.

LP to the Rescue

- The way out of this morass is to pull out the Swiss army knife of derivative pricing: the Linearity Principle.
- In our two period binomial model with zero interest rates and dividends, this principle says that if the derivative's payoff at the end of the second period is linear in the stock price:

$$D_2 = N_0 + N_1 S_2,$$

  then the initial value must be linear in the initial stock price to avoid arbitrage:

$$D_0 = N_0 + N_1 S_0.$$

- How can we use this principle to price derivatives with nonlinear payoffs?

Example: A Quadratic Payoff

- For example, suppose that a derivative's payoff at the end of the second period is the following quadratic in the stock price:

$$Q_2(S_2) = (S_2 - 2)^2, \quad S_2 > 0.$$

- The graph below shows the 5 possible payouts given that the stock price starts at $2 and over each period, it can double, halve, or remain the same.

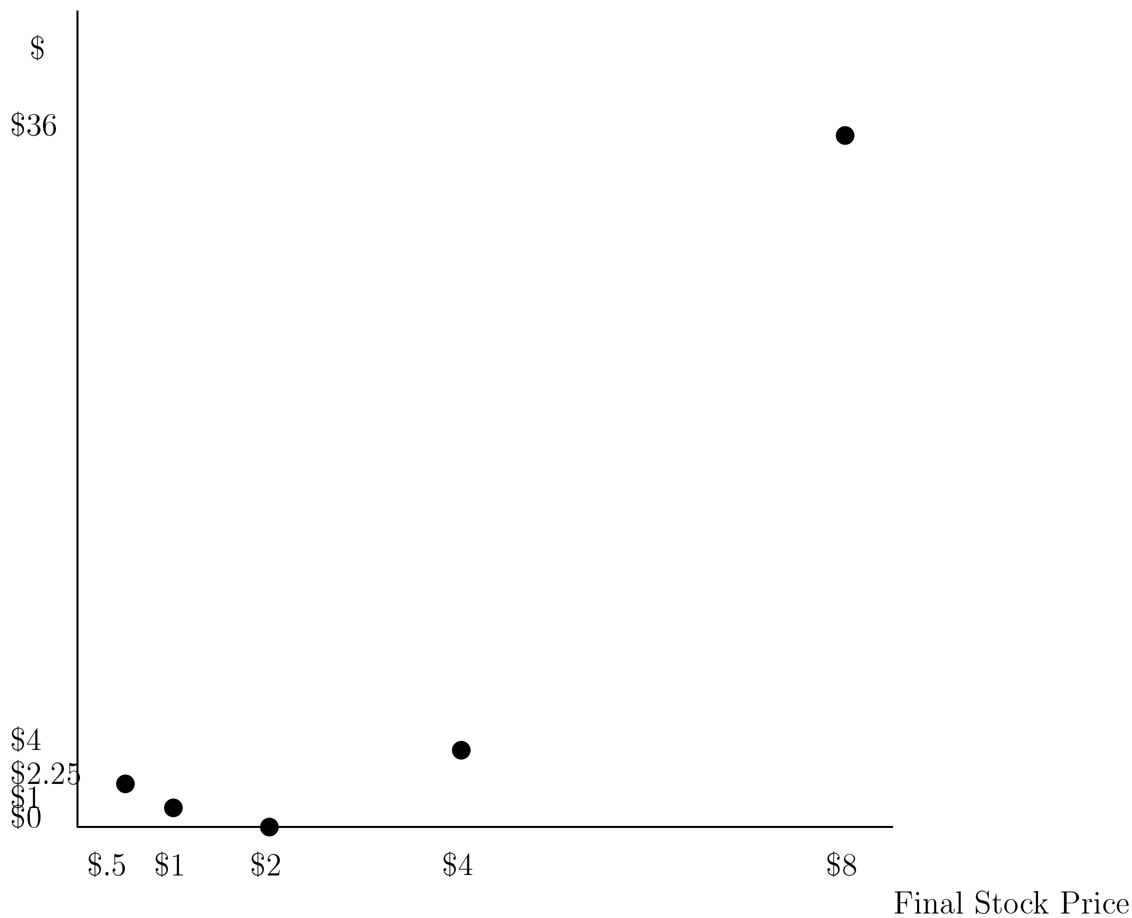

Figure 2: Payoff of Square of Stock Price Minus Two.

Butterfly Spreads

- There are a couple of ways to linearize the quadratic payoff without resorting to passing to the continuous time limit.
- Both methods are made easier to understand by introducing the fundamental notion of a butterfly spread whose payoff is graphed below:

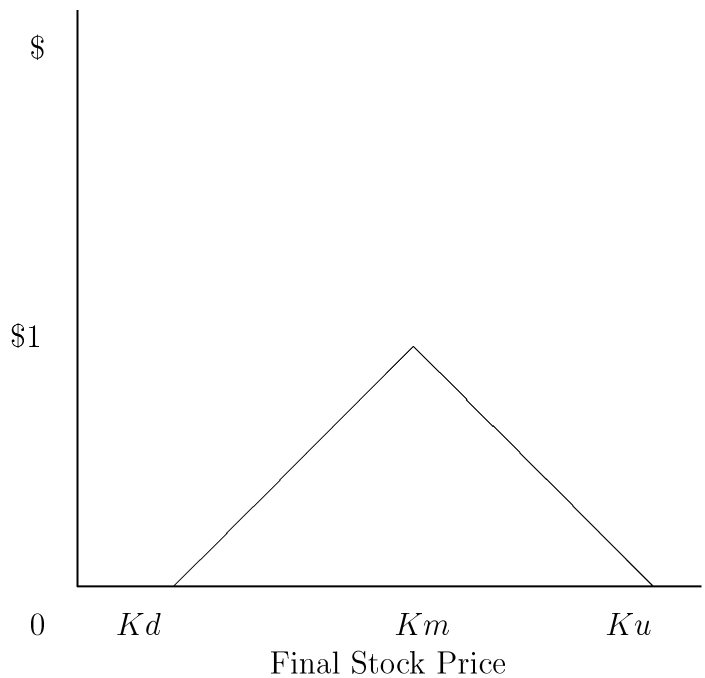

Figure 3: Payoff of a Butterfly Spread.

- Although the kinks can be placed at any option strike and the height is arbitrary, we will restrict attention to butterfly spreads with kinks only at possible stock prices and with heights of one.

Linearizing with Butterfly Spreads

- In the first method, linearization of the quadratic payoff is achieved by drawing a line through any 2 of the 5 possible payoffs and then using (static) positions in butterfly spreads centered at the other 3 to draw the payoff into line.

- As before, let N_0 and N_1 denote the respective intercept and slope of the linearized payoff:

$$L_2(S_2) = N_0 + N_1 S_2, \quad S_2 > 0.$$

- For example, if the second and third points from the left are used to draw a line, then the line goes through $(Sdm, D_{dm}) = (1, 1)$ and $(Smm, D_{mm}) = (2, 0)$ and so has equation:

$$L_2(S_2) = 2 - S_2, \quad S_2 > 0.$$

  with intercept $N_0 = 2$ and slope $N_1 = -1$.

- The table below calculates the number of butterfly spreads to hold struck at each of the remaining three possible stock prices:

Stock Price	$L_2(S_2)$	$Q_2(S_2)$	# Butterfly Spreads
$\frac{1}{2}$	$1\frac{1}{2}$	$2\frac{1}{4}$	$-\frac{3}{4}$
4	-2	4	-6
8	- 6	36	-42

- The second column computes the desired linear payoff at each of the 3 remaining possible stock prices, while the third column computes the quadratic payoff. The fourth column calculates the number of butterfly spreads to hold at each strike of interest as the difference between the second and third columns.

Static Replication

- Since the quadratic payoff has been linearized by shorting butterfly spreads, the Linearity Principle can be invoked.
- In our example, under zero interest rate and dividend yield, the linear payoff has zero initial value:

$$L_0 = 2 - 1(2) = 0,$$

since the initial stock price is $S_0 = 2$.

- To recover the value of the quadratic payoff, we need to add back the known premiums of the butterfly spreads:

$$Q_0 = L_0 + \frac{3}{4} BS_0\left(\frac{1}{2}\right) + 6BS_0\left(4\right) + 42BS_0\left(8\right).$$

- For example, if it so happens that all 3 butterfly spreads cost $\$\frac{1}{9}$, then the arbitrage-free value of the quadratic payoff is:

$$Q_0 = 0 + \frac{1}{9}\left[\frac{3}{4} + 6 + 42\right] = \frac{195}{36}.$$

- This method has statically replicated the quadratic payoff using stocks, bonds and three butterfly spreads. In fact, one can hold almost any 3 options in place of the butterfly spreads.

Dynamic and Static Replication

- The second method for linearizing involves combining dynamic trading in stock and bond with static positions in options.
- To obtain risk-neutral probabilities, we use forward induction.
- Specifically, suppose that we know that the initial price of a single period at-the-money call is \$4/7 and that the initial stock price is \$2 and finally that the single period unit bond price is \$6/7. Then, the one period forward prices are 2/3, 7/3, and 1 for the call, stock, and bond respectively.
- The three risk neutral probabilities q_u, q_m, and q_d solve the system:

$$\begin{array}{rrl} \text{bond} & q_u + q_m + q_d = & 1 \\ \text{stock} & q_u 4 + q_m 2 + q_d 1 = & 7/3 \\ \text{call} & q_u 2 + q_m 0 + q_d 0 = & 2/3. \end{array}$$

- It is clear by inspection that the three risk-neutral probabilities for moves over the first period are all 1/3.

Forward Induction

- Recall that the first period risk-neutral probabilities were all $\frac{1}{3}$:

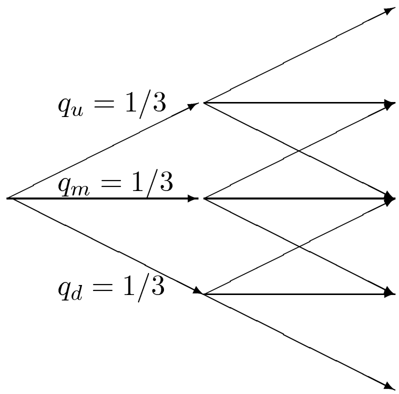

- Now place yourself at the middle date upper state.
- To get the risk-neutral probabilities of the 3 successor nodes, assume that we know the initial price of a 2 period call struck at 4 is 64/49. Since the call's only payoff of 4 occurs only if the stock goes up twice, we know:

$$\frac{1}{3} q_{uu} (\frac{6}{7})^2 4 = 64/49.$$

  and hence $q_{uu} = 1/3$.
- By stationarity, the other two probabilities are $q_{um} = q_{ud} = 1/3$.
- Similarly, if we know the initial price of a two period European put struck at 1, we can back out q_{dd} and then q_{dm} and q_{du}.
- Finally, if we know the price of a European at-the-money call or put, the remaining three probabilities can be inferred.
- Since the risk-neutral price of each path is now known, it is straightforward to value the quadratic payoff.

Valuing Path-dependent Options

- This second approach of combining dynamic trading in stock and bond with static replication in options can also be used to price path-dependent claims.
- To illustrate, consider valuing an at-the-money American put in our two period setup.
- If the stock goes down initially, then the put should be exercised since the final payoff if held alive is locally linear and thus the volatility value is zero. On all other paths, the put is either exercised only at maturity or not at all.
- The standard dynamic program begins by assigning final values to terminal nodes:

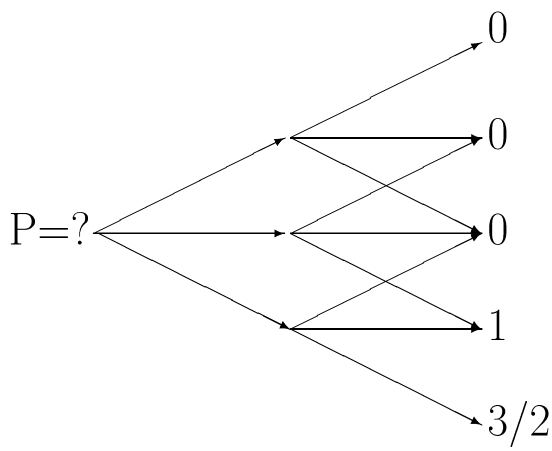

- Assuming that we know the price of 4 European options, the risk-neutral probabilities can again be determined.
- The American put is then valued in the standard manner.

How Many Options?

- In comparing the two approaches, we see that four European option prices were needed to determine the price of each path, while only three prices were needed to price the path-independent quadratic payoff.
- The market is said to be complete when the four prices are known and is said to be incomplete when only three are known. However, we can see that all European-style claims can still be priced in our incomplete market with only three option prices.
- In fact, the path-dependent American put can be priced knowing only three option prices, since the local linearity of the payoff at the bottom three nodes implies that low strike options are not needed.
- The hedge for the American put would involve static positions in the three options along with dynamic trading in stock and bond. The positions can be determined by linearizing locally or even globally.
- Under global linearization, the linear payoff across the 5 terminal nodes is propagated back one period and then the known option values are subtracted off to get the continuation value of the American put. The larger of this value and the exercise value is computed and then linearized using the single period option. We can again progagate back and subtract off the known single period option value.

Dynamic Trading in Options

- We note that in many options markets, there may not be enough options maturing at each date.
- Furthermore, as the number of periods grows beyond two, the number of options held grows exponentially. This shortage of available strikes can be addressed using dynamic replication with options as we now explore in the two period setup.
- Suppose that only the at-the-money European option price is available and we wish to price the quadratic payoff.
- To completely specify the price process for the at-the-money option, we need the three prices at the intermediate date.
- These can be obtained either directly or indirectly.
- In the direct approach, we guess three option values and hope that no arbitrage is introduced.
- Since it is actually quite easy to inadvertently introduce arbitrage, we advocate the indirect approach in which we specify the risk-neutral process directly and compute the values of the three hedging instruments at each node by discounting expected values. There are many ways to specify a risk-neutral process which in general trade off the competing requirements of tractability and realism.
- Just as one need only dynamically trade a stock and bond in a binomial model, one *can* only dynamically trade in three assets in a trinomial model.

Calibrating to All Liquid Option Prices

- In both the binomial and trinomial models, dynamic trading in only two or three assets ignores useful market information.
- A better approach is to use all of the information in the initial option prices and then to impose further criteria on the risk-neutral process such as smoothness or distance from a prior to select one.
- To minimize model risk, the hedge strategy should use all of the available instruments. Multivariate linear regression can be used recursively to choose the positions so that next period values of the target claim are as close to linear in hedging asset prices as possible.
- In practice, the use of all available instruments is costly due to transactions costs/illiquidity.
- Modelling the tradeoff between model risk and transactions costs is an interesting open problem.

VI Can Jumps be Hedged in a Continuous Time Model?

- In short, yes. The binomial and trinomial models can be viewed as continuous time jump models with known jump times.
- At the opposite end of the spectrum are Poisson processes with completely unanticipated jump times. Cox and Ross (1976) showed that options on stocks following such processes can be hedged if the jump size is known.
- Note that the jump size does not have to be the same over time and in general can depend on the path.
- By analogy, a jump process with two possible jump sizes can be hedged by dynamic trading in a bond, stock, and call.
- By extension, a continuous jump size distribution as in Merton's lognormal jump model requires continuous trading in a continuum of assets in order to make markets complete.
- However, as in the discrete time case, many interesting contracts may still be spanned in incomplete markets.
- Furthermore, not all options are used up in the hedge/calibration, so that for example it is still theoretically possible to price options of one maturity in terms of options of another maturity.

VII Why does the "market price of risk" appear in many stochastic interest rate model or volatility models? Why doesn't it appear in the HJM model?

- In single factor spot rate models (eg. Vasicek), the values of interest rate derivatives depend on the market price of interest rate risk. Yet no such concept appears in the Black model for a bond option or in the HJM model. Why?
- The reason is that the Black model relates option prices to bond *prices*, while the HJM model relates option prices to a known mathematical function of bond prices, namely forward rates.
- In a single factor spot rate model, if one writes down the PDE for the function relating the option price to the bond price (rather than to the spot rate), then the market price of interest rate risk drops out.

Goodbye λ

- If the statistical process for the spot rate r is given by:

$$dr_t = b(r_t, t)dt + a(r_t, t)dW_t,$$

then the absence of arbitrage implies that the price $C(r,t)$ of any non-coupon paying claim obeys:

$$\frac{\partial C}{\partial t} + \frac{a^2}{2}\frac{\partial^2 C}{\partial r^2} + [b + \lambda a]\frac{\partial C}{\partial r} = rC, \tag{1}$$

where $\lambda(r,t)$ is the market price of interest rate risk. It follows that the price of a pure discount bond, $P(r,t)$ obeys the same PDE:

$$\frac{\partial P}{\partial t} + \frac{a^2}{2}\frac{\partial^2 P}{\partial r^2} + [b + \lambda a]\frac{\partial P}{\partial r} = rP, \tag{2}$$

Suppose that we relate the claim price to the bond price by defining the function $V(P,t)$ by:

$$C(r,t) = V(P(r,t),t). \tag{3}$$

Differentiating C w.r.t. t and r and substituting in (1) yields the following generalization of the Black model PDE:

$$\frac{\partial V}{\partial t} + \frac{\sigma^2 P^2}{2}\frac{\partial^2 V}{\partial P^2} + r\frac{\partial V}{\partial P} = rV, \tag{4}$$

where:

$$\sigma(r,t) \equiv \frac{a(r,t)}{P(r,t)}\frac{\partial P}{\partial r}(r,t)$$

is the volatility of bond returns, which is being modelled here as a known function of the spot rate and time.

- The market price of risk λ does not appear in (4)!

Other Models

- The market price of risk also does not appear whenever option prices are related to known functions of the bond price.
- Brennan and Schwartz(1979) showed that the market price of long rate risk does not appear in their valuation PDE when the long rate is used as a state variable, since it is known to be the reciprocal of the perpetuity value.
- Similarly, since the term structure of forward rates $f_t(T) \equiv -\frac{\partial \ln P_t(T)}{\partial T}$, is a known function of the term structure of bond prices, the market price of risk does not appear in the HJM model.
- The same observations apply for stochastic volatility models. Thus, one can develop stochastic volatility models which do not depend on the market price of volatility risk by relating option prices to either other option prices or to known functions of option prices such as "forward local volatility" or Black Scholes implied volatility.
- The paper shows how the market price of volatility risk drops out when an option price is related to another option price.

VIII How much do I lose if I hedge at the wrong volatility?

- The glib answer is nothing of value.
- Assume that the futures price process is continuous and that the true vol is given by some unknown stochastic process σ_t:

$$\frac{dF_t}{F_t} = \mu_t dt + \sigma_t dW_t, \quad t \in [0, T].$$

- Assume that a claim on the futures price is sold for an initial implied vol of σ_i and that delta-hedging is conducted using the Black model delta evaluated at a constant hedge vol σ_h.
- Applying Itô's lemma to $V(F, t)e^{r(T-t)}$:

$$\begin{aligned} V(F_T, T) &= V(F_0, 0)e^{rT} + \int_0^T e^{r(T-t)} \frac{\partial V}{\partial F}(F_t, t) dF_t \\ &+ \int_0^T e^{r(T-t)} \frac{1}{2} \frac{\partial^2 V}{\partial F^2}(F_t, t)(dF_t)^2 \\ &+ \int_0^T e^{r(T-t)} \left[\frac{\partial V}{\partial t}(F_t, t) - rV(F_t, t) \right] dt. \end{aligned}$$

- Since the futures has stochastic volatility, $(dF_t)^2 = \sigma_t^2 F_t^2 dt$ and so:

$$\begin{aligned} V(F_T, T) &= V(F_0, 0)e^{rT} + \int_0^T e^{r(T-t)} \frac{\partial V}{\partial F}(F_t, t) dF_t \\ &+ \int_0^T e^{r(T-t)} \left[\frac{\sigma_t^2 F_t^2}{2} \frac{\partial^2 V}{\partial F^2}(F_t, t) + \frac{\partial V}{\partial t}(F_t, t) - rV(F_t, t) \right] dt. \end{aligned}$$

Delta-Hedging at a Constant Vol in a Stochastic Vol World (con'd)

- Recall:

$$\begin{aligned} V(F_T, T) &= V(F_0, 0)e^{rT} + \int_0^T e^{r(T-t)} \frac{\partial V}{\partial F}(F_t, t) dF_t \\ &+ \int_0^T e^{r(T-t)} \left[\frac{\sigma_t^2 F_t^2}{2} \frac{\partial^2 V}{\partial F^2}(F_t, t) + \frac{\partial V}{\partial t}(F_t, t) - rV(F_t, t) \right] dt. \end{aligned}$$

- Suppose we have chosen our function $V(F,t)$ to be the function $V(F,t;\sigma_h)$ which solves the Black PDE:

$$\frac{\sigma_h^2 F^2}{2} \frac{\partial^2 V}{\partial F^2}(F, t; \sigma_h) + \frac{\partial V}{\partial t}(F, t; \sigma_h) - rV(F, t; \sigma_h) = 0,$$

and the terminal condition:

$$V(F, T; \sigma_h) = f(F).$$

- Substitution implies:

$$\begin{aligned} f(F_T) &= V(F_0, 0; \sigma_h)e^{rT} + \int_0^T e^{r(T-t)} \frac{\partial V}{\partial F}(F_t, t; \sigma_h) dF_t, \\ &+ \int_0^T e^{r(T-t)} [\sigma_t^2 - \sigma_h^2] \frac{F_t^2}{2} \frac{\partial^2 V}{\partial F^2}(F_t, t; \sigma_h) dt. \end{aligned}$$

Delta-Hedging at a Constant Vol in a Stochastic Vol World (con'd)

- Recall:

$$\begin{aligned} f(F_T) &= V(F_0, 0; \sigma_h)e^{rT} + \int_0^T e^{r(T-t)} \frac{\partial V}{\partial F}(F_t, t; \sigma_h) dF_t, \\ &+ \int_0^T e^{r(T-t)} [\sigma_t^2 - \sigma_h^2] \frac{F_t^2}{2} \frac{\partial^2 V}{\partial F^2}(F_t, t; \sigma_h) dt. \end{aligned}$$

- Adding $V(F_0, 0; \sigma_i)e^{rT}$ to both sides and re-arranging:

$$P\&L_T \equiv V(F_0, 0; \sigma_i)e^{rT} - f(F_T) + \int_0^T e^{r(T-t)} \frac{\partial V}{\partial F}(F_t, t; \sigma_h) dF_t,$$

  where:

$$\begin{aligned} P\&L_T &= [V(F_0, 0; \sigma_i) - V(F_0, 0; \sigma_h)]e^{rT} \\ &+ \int_0^T e^{r(T-t)} \frac{F_t^2}{2} \frac{\partial^2 V}{\partial F^2}(F_t, t; \sigma_h)(\sigma_h^2 - \sigma_t^2) dt. \end{aligned}$$

- In words, when we sell the claim for an implied vol of σ_i initially, the total P&L from delta-hedging with the constant vol σ_h over $(0, T)$ is the future value of the difference in Black Scholes valuations plus half the accumulated dollar gamma weighted average of the difference between the hedge variance and the true variance.

- Note that the P&L vanishes if $\sigma_t = \sigma_h = \sigma_i$.

- If $\frac{\partial^2 V}{\partial F^2}(F_t, t; \sigma_h) \geq 0$ as is true for options, and if $\sigma_i = \sigma_h < \sigma_t$ for all $t \in [0, T]$, then you sold the claim for too low a vol and a loss results, regardless of the path. Conversely, if $\sigma_i > \sigma_t$ for all $t \in [0, T]$, then delta-hedging at $\sigma_h = \sigma_i$ guarantees a positive P&L.

The Irrelevance of Hedge Vol for Risk-Neutral Expected P&L

- Recall the definition of P&L:

$$P\&L_T = V(F_0, 0; \sigma_i)e^{rT} - f(F_T) + \int_0^T e^{r(T-t)} \frac{\partial V}{\partial F}(F_t, t; \sigma_h) dF_t.$$

- Define the risk-neutral expected value of a random variable as the initial cost of creating the random variable using a self-financing strategy. Then:

$$E_0^Q P\&L_T = V(F_0, 0; \sigma_i)e^{rT} - E_0^Q f(F_T),$$

  since $E_0^Q dF_t = 0$ for all $t \in [0, T]$.

- The RHS is independent of σ_h. It follows that the hedge vol is irrelevant for risk-neutral expected P&L. You lose (and gain) nothing of value if you delta-hedge at the wrong volatility.

- For a call, varying the hedge vol between 0 and ∞ causes the delta hedge to vary between 0 and 1 contracts. In fact, the general result here is that the risk-neutral expected P&L is independent of the hedge. It is determined at inception and equals the difference between the future value of what the claim was sold for, $V(F_0, 0; \sigma_i)e^{rT}$, and its fair forward price, $E_0^Q f(F_T)$.

- Determining this fair forward price when volatility is stochastic is a major research thrust.

IX Summary

- We did our best to answer several FAQ's.
- Here are a few more to try on your own:
  1. Does risk-neutral valuation work over infinite horizons? Bear in mind that shorting a stock appears to be an arbitrage opportunity since the probability is one in many models (eg. Black Scholes) that a short investor can cover for a dollar lower later.
  2. To what extent does expected return matter in incomplete markets? For example, consider a pure jump process and suppose an investor chooses an equivalent martingale measure and delta-hedges in the standard way. While the P&L distribution depends on expected return, Monte Carlo simulation suggests that there are many paths over which this dependence is mild.
  3. How does the inability to withstand arbitrarily large losses impact standard valuation procedures? Why would anyone lend at the riskfree rate to imperfect credits engaging in risky strategies? If the lending rate varies with the credit risk, how should the risk-neutral drift of the underlying and the option reflect this credit spread?