

DERIVATIVES

**Professor Craig Pirrong
Energy Risk Management
Fall, 2004**

What is a Derivative?

A “derivative” is a financial contract.

Derivatives contracts get their name from the fact that they are “derived from” some other, “underlying” claim, contract, or asset.

For instance, a gold forward contract is “derived from” the underlying physical asset—gold.

Derivatives are also called “contingent claims.” This term reflects the fact that their payoff—the cash flow—is contingent upon the price of something else. Going back to the gold forward contract example, the payoff to a gold forward contract is contingent upon the price of gold at the expiration of the contract.

Although derivatives are frequently considered to be something new and exotic, they’ve been around for millennia. There are examples of derivative contracts in Aristotle and the Bible.

It is true, however, that there has been an explosion in the variety, complexity, and use of derivatives, especially in the last 30 years.

In this course, you will learn what derivatives contracts are, how you can use them, and how to price them.

The most basic derivative is a forward contract.

Basics of Forwards and Futures

A forward contract is an agreement between a buyer and a seller to transfer ownership of some asset or commodity (“the underlying”) at an agreed upon price at an agreed upon date in the future.

A forward contract is a promise to engage in a transaction at some later date.

The forward contract specifies the characteristics of the underlying. For example, for a commodity, it specifies the type of commodity (e.g., silver), the quality of the commodity (e.g., 99.9 percent pure silver), the location of delivery, the time of delivery, and the quantity to be delivered.

The primary use of a forward contract is to lock in the price at which one buys or sells a particular good in the future. This implies that the contract can be utilized to manage price risk.

Most forward contracts are traded in the “over the counter” (OTC) market.

Some forward contracts are traded on organized exchanges such as the Chicago Board of Trade or the New York Mercantile Exchange. These exchange traded contracts are called “futures contracts.”

Forward contracts traded OTC can be customized to suit the needs of the transacting parties. Exchange traded contracts are standardized. This enhances liquidity.

Performance on futures contracts are guaranteed by third parties (brokers and the clearinghouse.) Performance on OTC forwards is not guaranteed. The quality of the contractual promise is only as reliable as the firm making it.

Forwards and futures are traded on a wide variety of commodities and financial assets.

Commodity futures and forwards are traded on agricultural products (corn, soybeans, wheat, cattle, hogs, pork bellies); precious metals (silver, gold, platinum, palladium); industrial metals (copper, lead, zinc, aluminum, tin, nickel); forest products (lumber and pulp); and energy products (crude oil, gasoline, heating oil, natural gas, electricity).

Financial futures and forwards are traded on stock indices (S&P 500, Dow Jones Industrials, foreign indices); government bonds (US Treasury bonds, US Treasury notes, foreign government bonds); and interest rates (Eurodollars, EuroEuros).

More recently, forward/futures trading has begun on weather and credit risk. These are (no pun intended) the hottest areas in derivatives development.

The Uses of Derivative Markets

- Derivatives markets serve to shift risk.
- Hedgers use derivatives to reduce risk exposure. For instance, a refiner can lock in costs and revenues (i.e., lock in its margin) by buying crude oil futures and selling oil and gasoline futures.
- Speculators use derivatives to increase risk exposure in the anticipation of making a profit.
- Thus, derivatives markets facilitate the shifting of risk from those who bear it at a high cost (the risk averse) to those who bear it at a low cost (the risk tolerant).
- Speculators perform a valuable service by absorbing risk from hedgers. In return, they receive a reward—a risk premium. The risk premium is the expected profit on a derivatives transaction. Speculators may win or lose in any given trade, but on average speculators expect to profit.
- The risk premium is also the cost of hedging.

Hedging Basics

Some terminology:

Going “short” means to sell a forward/futures contract.

Going “long” means to buy a forward/futures contract.

A “short hedger” sells forwards as a hedge. For instance, the holder of an inventory of crude oil is subject to price risk. The value of the inventory declines if the price declines, but its value goes up with the price. The inventory holder can sell a forward contract as a hedge.

Hedging works because the forward price tends to move together with the value of what is being hedged.

For instance, consider a firm hedging a cargo of Dubai. The value of the cargo rises and falls with the price of Dubai crude at the location to which the cargo is being shipped. The cargo owner can sell Brent forward contracts (or Brent futures contracts). The price of Brent tends to move in tandem with the price of Dubai. Hence, when the price of Brent rises, the price of Dubai tends to rise too. Similarly, Brent prices tend to decline when the Dubai price falls.

For instance, assume that Dubai is currently selling at \$25/bbl and Brent futures are selling for \$28/bbl. The firm buys the cargo of Dubai and sells the Brent today.

By the time the cargo is delivered and sold, oil prices have moved. The firm sells the Dubai cargo at \$23/bbl. Brent is selling at \$26.08. That is, the price of Brent has fallen by \$1.92 and the price of Dubai by \$2.00. The firm loses \$2.00/bbl on its cargo, but gains \$1.92 on the Brent position (it sold at \$28, bought back at \$26.08, for a gain of \$1.92). Thus, the firm's net loss is \$.08—only 4 percent of what it would have lost if it hadn't hedged.

Alternatively, assume that when the cargo is delivered and sold, Dubai is selling for \$27.25 and Brent for \$30.19. Here the firm's cargo is now worth \$2.25 more per bbl, but it has lost \$2.19/bbl on its Brent futures, for a net gain of \$.06/bbl.

Some things to note:

- The firm still faces some risk when it hedges, but the variability of its gains and losses is smaller when it hedges than when it doesn't.
- No hedge is “perfect.” There is always some risk.
- The difference between the price of what is being hedged (e.g., the Dubai) and the price of the hedging instrument (the Brent) is called the “basis.”

- Hedgers trade “flat price risk” for “basis risk.” That is, without hedging, the firm is at risk to changes in the flat price of Dubai. When it hedges, it is at risk to changes in the difference between the prices of Dubai and Brent—i.e., it is at risk to basis changes.
- The variability of the basis—and hence the risk that the hedger faces—depends on the correlation between the prices of what’s being hedged and the hedging instrument. High correlation results in low basis risk and high hedge effectiveness.
- Therefore, hedgers care about basis risk and correlation. Correlations depend on the closeness of the match between what’s being hedged and the hedging instrument. There can be differences in quality (e.g., sweet crude vs. sour crude), differences in location (e.g., Henry Hub natural gas vs. Chicago City Gate NG), or differences in timing (spot gas vs. gas for delivery next month).

- Basis behavior varies by commodity. It can also vary over time.
- Some examples include heating oil vs. jet fuel in the first Gulf War, and California Border (Topok) gas vs. HH gas futures in 2000.

Long hedging is symmetric to short hedging. A long hedger buys forward contracts to hedge risk. As an example, an electric utility with a gas fired plant faces the risk that the price of gas may rise. It can lock in the cost of gas by purchasing a gas futures contract. If the price of gas delivered to its plant does rise, the additional cost will be offset in whole or in part by a rise in the futures contract. If the price of gas falls, the firm loses on its futures position, but this is offset in whole or in part by the lower cost of purchasing gas on the cash market to fuel its plants.

Delivery and Cash Settlement

Most commodity forwards and futures allow the short to make delivery and the long to take delivery. Typically the short has the option to choose when to deliver (within a pre-specified delivery period).

Although most futures contracts call for delivery, very few contracts actually result in delivery. Roughly 2 percent (give or take) of contracts traded actually result in delivery.

This is true because most hedgers are not dealing in the commodity deliverable against the futures contract. For instance, an oil refiner in California is not going to use WTI crude oil in Cushing, OK for its plant, but may use the WTI futures contract as a hedge. That is, most hedgers are “cross hedgers.” Similarly, speculators are just “betting” on price movements, and have no interest in owning the phys.

Therefore, most hedgers and speculators reverse their positions prior to delivery. Shorts buy back their futures prior to delivery, and longs sell back their futures.

Even though delivery doesn't occur on most contracts, delivery is important nonetheless. Delivery ties the price of the expiring future to the price of the physical commodity at delivery. That is, delivery (and the potential for delivery) forges the connection between the futures price and the price of the actual commodity. This is essential for hedging.

Relatedly, it is essential that delivery terms in a futures contract reflect physical market realities. It is important that the futures deliverable exhibit a high correlation with the prices of other varieties and locations to ensure hedging effectiveness. Also, the delivery market should have sufficient capacity to reduce its vulnerability to manipulation (more about which later).

Cash settlement is another way to tie the futures and cash markets together.

In a cash settled contract, at expiration the buyer pays the seller the difference between the fixed price established in the contract and the reference price prevailing on the payment (expiration date).

For instance, consider a cash settled contract based on the price of gas at the Chicago City Gate as published in *Inside Ferc*. The contract expires in December. When the parties enter the contract, they agree on a price (say, \$6.00/MMBTU). At expiration, the parties look at *Inside Ferc* to determine the CCG price. Let's say it's \$6.25. In this case, the seller owes the buyer \$.25 per MMBTU. If the price were \$5.80 instead, the buyer would owe the seller \$.20/MMBTU.

Cash settlement is widely used on OTC derivatives. It's important that they be based on reliable cash indices. Unfortunately, recent revelations demonstrate that in the energy industry cash indices have been anything but reliable.

EFPs and ADPs

Although few futures contracts are ultimately settled by delivery under the standard delivery terms specified in an exchange's futures contract, futures are frequently employed as part of a transfer of ownership of a commodity.

The most common technique is called an EFP—an “exchange for physicals” (also called an “ex pit” trade, or an “exchange for cash.”)

Typically an EFP involves two hedgers, one of which (“L”) has bought a futures contract in anticipation of buying the actual commodity at a future date, and the other (“S”) has sold a futures contract in anticipation of selling the actual commodity at a future date. (It is sometimes said that a futures trade is a temporary substitute for a future physical transaction).

Sometime prior to expiry of the contract, S and L decide that they want to trade the physical with one another. As soon as they agree on a price, they want to terminate their futures positions because they no longer need them as a hedge.

One thing they could do is just agree on a price for the phys trade (e.g., \$30/bbl), and then each go to the futures market to offset their positions at the prevailing market price. The problem with this is that the parties are at risk from the instant that they agree on the phys trade price to the time that they liquidate their futures. Even if they move quickly, the futures market can move a good deal before they close their positions. This is “slippage” or “execution risk.”

An EFP addresses this problem. In an EFP, the parties agree to a price at which the parties exchange the physical and the underlying. S gives up his short position and the physical commodity to L. L gives cash to S. L now has a long and a short futures position that offset—he is no longer exposed to the futures market. Similarly, S has given up his futures exposure.

The parties can transfer the futures at any price that they choose. This price reflects the value of the physical trade agreed to by the parties.

The physical good traded needn't be the same one deliverable under the contract. For instance, a natural gas buyer and seller may use an EFP involving NYMEX Henry Hub gas futures as part of a trade for Alberta NG. On NYMEX, the quantity of physical commodity traded must be “approximately equal to” the quantity of futures exchanged.

It is also possible to use an EFP to establish a futures position.

NYMEX also allows alternative delivery procedures, or ADPs.

In an ADP, a futures long L and a futures short S who have been matched by the clearinghouse so that S is supposed to deliver against futures to L can work out delivery terms that differ from those set out in the exchange contract.

Swaps

Swaps are the most common OTC derivative.

Most swaps are “vanilla” fixed price swaps. These are like bundles of forward contracts. (Though some swaps have only one pricing date.)

Vanilla swaps are typically cash settled.

The parties to a swap set: (a) the notional quantity; (b) the “tenor” or maturity of the swap (how long it lasts); (c) the payment dates; (d) the floating price index; and (e) the fixed price.

The mechanics of a swap are as follows: There is a fixed payer price (the “long”) and a floating price payer (the “short”). The fixed price payer makes the same payment to the counterparty (the floating price payer) every payment date. The floating payer pays the counterparty (the fixed price payer) an amount that depends on the floating price index.

The most common floating price index is the average of the last 3 days of the NYMEX NG futures settlement prices corresponding to the payment month. However, parties could agree to any price index, e.g., Chicago City Gate quoted by Platts.

The party whose payment obligation is larger pays the net amount to the counterparty. The payment equals the absolute value of the difference between fixed and floating prices times the notional quantity.

An example works best. Consider a swap for 100,000 MMBTU per month (equivalent to 10 NYMEX contracts) for each month in September, 2004-August, 2007. The fixed price in the swap is \$5.95 per MMBTU. Payments are made 5 days after the last 3 trading days of the NYMEX NG futures (each month from U04-Q07), and the floating price is based on the LTD average of NYMEX settle prices.

The September NYMEX contract ends trading on 8/27/2004. The payment date is September 1, 2004.

Assume that on 8/27/2004 the LTD September futures averaged \$5.40. The floating price is below the fixed price. Therefore, the long owes the short. The amount due is $$.45 \times 100,000 = \$45,000$.

Assume that on 12/28/2004, the LTD January 05 futures averaged \$6.84. Here the short owes the long $$.89 \times 100,000 = \$89,000$.

Determining The Fixed Price

Fixed prices are determined at the time that the swap is negotiated based on the relevant forward prices at that time. For NYMEX LTD, the NYMEX futures prices prevailing when the swap is negotiated are used to determine the fixed price.

The fixed price is set so that a swap has zero value on the date that it is created.

A hedging argument indicates how the swap price is determined. Consider a firm that buys a 3 year NYMEX LTD swap. It can hedge this swap by selling a 3 year strip of NYMEX futures. On each payment date T , this generates a cash flow of:

$$F_{t,T} - P$$

where $F_{t,T}$ is the NYMEX futures price corresponding to payment date T as of the date that the swap is negotiated (t).

This payment is known as of t . It can be discounted back to the present to determine a present value.

The value of the swap given a fixed price P is therefore:

$$\sum_{T=T1}^{TN} (F_{t,T} - P)DF_T$$

where DF_T is the discount factor (present value) corresponding to date T , $T1$ is the first payment date, and TN is the last payment date.

At initiation of the swap, P is set so that the value of the swap equals zero. This requires:

$$P = \frac{\sum_{T=T1}^{TN} F_{t,T}DF_T}{\sum_{T=T1}^{TN} DF_T}$$

Note that the swap price depends on the forward curve (the futures prices) and interest rates (which affect the discount factors).

After the swap is initiated, the forward/futures prices change, and hence the value of the swap can become either positive or negative. (Remember this is a zero sum deal, so if it is positive to one party it is negative to the other.)

The Use of Swaps

Swaps are hedging instruments (and can be used as speculative tools as well).

Consider a gas producer that wants to hedge its revenues for the next year. It can sell a 1 year swap. If prices fall, the floating payments in the swap become smaller, the fixed payments stay the same, so the net payment that the producer makes gets smaller. This smaller net payment offsets the effect of the lower gas price on the firm's revenues. Of course, if gas prices rise, the floating payments on the swap get bigger, but this is offset by higher revenues from the sales of phys gas.

Alternatively, consider an oil refiner that wants to hedge its crude costs for the next 2 years. It can buy a 2 year CL swap. If oil prices rise, the refiner receives a bigger payment from the swap counterparty, which offsets the impact of higher crude acquisition costs; on the flip side, if oil prices fall, the refiner receives a smaller payment from the counterparty, but this is offset by the fact that lower crude prices mean lower costs.

Other Kinds of Swaps

Vanilla fixed price swaps are the most common form of energy swap, but there are others.

Basis swaps have two floating price payers. The net payment is based on the difference between two different floating indices.

For instance, a basis swap may have a payoff based on the difference between the Chicago City Gate price and the Henry Hub price. Alternatively, it might be a difference between the CCG price and the NYMEX LTD price.

Combining a basis swap and a fixed price swap can create a swap that reduces basis risk. For instance, consider a large industrial gas consumer in Chicago. It can buy a NYMEX LTD swap, but since this is tied to Henry Hub prices, the firm bears basis risk (due to variations in the CCG-HH basis).

However, the firm can buy the NYMEX LTD fixed price swap, and at the same time enter a basis swap whereby the firm pays NYMEX LTD plus a differential and receives the CCG price. The NYMEX LTD payments on the fixed price swap and the basis swap cancel out. What's left is the fixed price in the NYMEX LTD swap (which the firm pays) and the CCG price plus the differential (which the firm receives). This effectively locks the firm's purchase price at the NYMEX swap fixed price plus the differential.

The swap counterparties negotiate the differential in the basis swap. For a swap in which Henry Hub or NYMEX prices determine one of the floating prices, for many locations (e.g., consuming locations such as the Midwest or Northeast) this differential is positive. This reflects the fact that it is costly to ship gas from producing regions (around HH) to consuming locations.

Market Mechanics: Trading

There are two basic ways to buy and sell forwards—on exchange, or in the Over-the-Counter (OTC) market.

The most important energy exchanges are the New York Mercantile Exchange and ICE. NYMEX is an “open outcry” exchange (with an electronic after hours trading system.) ICE is a fully electronic marketplace.

On Nymex, a firm (or individual) submits a buy or sell order to a broker on the floor of the exchange. The broker then enters the trading “pit” and endeavors to fill the order. The broker shouts out that he wants to buy (or sell, as the case may be), and other traders in the pit compete to offer the best price to fill the order.

The potential counterparties include other brokers representing other customers, or “locals” who trade on their own account.

The floor procedure is an “open outcry double auction.” It looks chaotic, but it is a very efficient way to buy and sell futures contracts.

Customers can submit different kinds of orders—market orders, limit orders, stop orders, etc.

A market order is an order to buy (or sell) at the best price prevailing in the market.

A limit order is an order to buy (or sell) at a specified price (or better).

A stop order is a market order that is only activated when the market price hits a certain level.

Spread orders are also quite common on the futures market. A spread order is an order to buy one contract and sell another related but different contract at the same time.

For instance, a time (or calendar) spread involves the sale of one month's contract and the purchase of another month in the same commodity (e.g., CL).

An intercommodity spread involves the sale of one commodity and the purchase of another. In energy, the "crack spread" is well known. This involves the purchase of crude oil futures and the simultaneous sale of gasoline and heating oil. The reverse crack does the opposite. Another example is the "spark spread"—the sale of electricity and the purchase of natural gas.

The Futures Trading Process

Customers submit orders to brokers.

Brokerage firms (“futures commission merchants,” aka FCMs) submit orders to trading system.

In open outcry market, floor brokers shout out intention to buy or sell.

Other brokers compete to take opposite side of order

In computerized market, orders that “cross” (buy price > sell price) are automatically matched and executed.

“Locals” and Liquidity

Market makers trading on their own account, often called “local” traders, also compete to execute orders.

Locals absorb order imbalances, and supply liquidity and depth to the floor.

Open outcry auction facilitates competition. Assists customers in assuring that they get the best possible price.

Only best bid and offer can be shouted out.

ICE is an electronic exchange. Buyers and sellers enter orders by computer. These orders are sent to a central computer that matches orders on the basis of price and time priorities. ICE also ensures that trades are done only with approved counterparties within established credit limits.

OTC markets are typically telephone markets. OTC dealers make markets (i.e., quote prices). Potential traders contact the dealer, negotiate a deal, and enter into a transaction.

Many OTC transactions go through brokers. This helps maintain the anonymity of the parties.

Computerized Futures Trading

Computerized trading of futures becoming increasingly common.

No open outcry exchange has opened since 1986 (MATIF); over a dozen computerized exchanges have opened in this time.

Established exchanges (e.g., CBOT, LIFFE) have computerized off-hours trading systems.

IPE has recently converted to an all electronic marketplace for trading energy derivatives (e.g., Brent crude oil futures, UK NatGas).

ICE is an all electronic exchange. It is more analogous to an electronic OTC market.

NYMEX has an after-hours electronic trading facility.

Computerized exchanges gaining market share in head-to-head competition with open outcry.

More on Computerized Trading

Computerized trading is typically cheaper.

Computerization reduces costs of accessing the market.

Considerable debate on whether computerized trading is, or ever can be, more liquid than open outcry.

In recent months representatives of open outcry exchanges have conceded the viability of computerized trading: this would have been unheard of even a year ago.

Over 80 percent of financial futures volume on the CBOT is now computerized.

Commodity markets (e.g., energy on NYMEX, grains on CBOT, meat on CME) are typically less electronic, but the IPE in London is moving to computerized trading despite the resistance of many of its local members.

Exchanges and Standardization

Standardization essential for either open outcry or computerized trading.

Many commodities not inherently standardized: e.g., wheat.

Exchanges economize on measurement costs by establishing standard grades and measures, and implementing measurement and arbitration systems. Examples include CBOT, Liverpool Cotton Exchange.

NYMEX and IPE create standardized terms for energy futures. These include delivery location terms (e.g., delivery of 1000 bbl of crude at Cushing, OK or 10,000 MMBTU of NG at Henry Hub) and quality terms (e.g., crude type, sulfur, gravity, viscosity, vapor pressure, sediment and pour point for CL).

OTC markets sometimes trade contracts that mimic standardized NYMEX contracts (“NYMEX look alike.”)

Whereas exchange trading requires standardization, OTC trading does not. Nonetheless, many OTC deals are standardized to some degree.

Credit Risk in Derivatives Markets

Standardizing credit risk is the most important function of futures exchanges.

Credit (default) risk always a concern in derivatives markets.

Prices very volatile. Losers may go bankrupt, owing large amounts to their counterparties.

Even solvent losers have an incentive to “walk away” from losing trades.

Standardizing Credit Risk

Anonymous pit or computer trading requires standardization of credit risk.

Futures markets accomplish this through “novation” or “substitution of principals.”

Clearinghouse becomes buyer to every seller, seller to every buyer.

Traders trade. Clerks “match” buy and sell. Once trades matched, clearinghouse stands as counterparty to all trades.

Clearinghouse guarantees trade performance.

Since in theory the clearinghouse stands behind every futures contract, credit risk is standardized: default risk of clearinghouse is same for all contracts.

Clearinghouse capital is typically large enough to make risk of default low, although there have been some close calls: October, 1987, “Silver Tuesday” in March, 1981.

Clearinghouse Operation

Things are actually a little bit more complicated. A trader first looks to his broker for performance. If the broker fails, the clearinghouse steps in.

Clearinghouse typically guarantees only *net* position of each “clearing member” broker.

In theory, therefore, if a brokerage fails, its customers may not receive all that is due them. This has happened in the past. Value Partners (COMEX clearing member) in the early 80s, Griffin (CBT clearing member) in the late 90s.

Margins and Other Safeguards

Futures markets take a “belts and suspenders” approach to managing credit risk.

Margins are an important link in the system

Margins are performance bonds. To trade futures, you must deposit money with your broker before you trade.

Losses are immediately withdrawn from margin account: gains are immediately paid in. This is called “marking to market.”

Marking-to-market based on settlement price (i.e., current market price) from the relevant futures markets.

Example. I buy 1 CL futures contract. Each contract is for 1000 bbl. The day after I buy, the price falls \$.25/bbl. I lose \$2500. This amount is withdrawn from my margin account.

If my margin balance falls below some predefined level (the “maintenance margin”), I must post additional funds with my broker.

Unless a price move wipes out my entire margin balance in one day, the broker has enough money to cover my losses.

Margins are set to ensure that margin balances seldom wiped out in a single day.

More on Margins

There is actually a chain of margins. Traders post margins with brokers. Non-clearing brokers post margins with clearing brokers. Clearing brokers post margins with the clearinghouse.

Other Safeguards

Markets adopt other safeguards to reduce the likelihood of default.

Brokers must have adequate capital.

In US and some other countries, customer funds segregated from “house” funds.

Clearing firms required to inject capital into the clearinghouse if clearinghouse’s capital is exhausted by a default.

Market Catastrophes

Failure of a clearinghouse could be catastrophic.

Extreme price moves could cause clearinghouse failure.

October, 1987 nearly resulted in severe financial problems for CME and CBOT clearinghouses.

Silver Thursday in March, 1981 almost closed COMEX clearinghouse.

Clearing system failed on Black Friday (gold corner) in 1869.

Central Banks (the Federal Reserve in the US) may be needed to supply liquidity to prevent market collapse and cascade of defaults.

Developments in Clearing

Clearing is a very dynamic area.

Many exchanges are exploring “common clearing.” This would involve multiple exchanges forming one clearinghouse. Indeed, CBT and CME have just agreed (after years of failed attempts) to combine clearing operations.

This would economize on margin costs and systems costs, and should improve information flow and coordination.

Cross-margining could also reduce demand for liquidity. This could be especially useful during market “crises” such as Black Monday, 1987.

Potential problem: failure of the common clearinghouse could bring down the entire financial system.

Initiatives to clear OTC transactions, including energy OTC transactions, have taken place in recent years. In energy markets BOTCC and EnergyClear have introduced clearing services. NYMEX also offers clearing of NYMEX look-alike OTC products through its ClearPort system.

Differences Between OTC and Exchange-Traded Derivatives

OTC products not *necessarily* standardized (although many features of basic swaps and forwards are in fact highly standardized). Traders of OTC products can customize features of these contracts; this customization not possible for exchange-traded products.

OTC products not cleared: bilateral credit risk.

OTC products not traded on centralized markets. Bilateral “search” market with intermediaries (“swap dealers”).

OTC derivatives have experienced explosive growth in recent years, far outstripping exchange traded volume growth.

Advantages of customization and scale appear to outweigh credit risk costs for many users.

Exchange markets are increasingly serving as a way for OTC dealers to lay off their unmatched risk.

OTC Derivatives and Credit Risk

Highly rated credits (e.g., AAA, AA) dominate swap intermediation in interest rate and FX markets.

A customer's credit risk in the swap market is determined by the credit risk of the dealer he contracts with.

Some discussions of central clearing in swaps/OTC derivatives market, but major dealers may have a vested interest in retaining current system.

In the late-90s, Enron (primarily through Enron Online, its electronic trading platform), became the primary dealer in OTC energy transactions.

It was somewhat surprising that Enron was able to achieve such a dominant position, given its marginal investment grade rating (BBB).

Energy markets suffering in large part because eroded creditworthiness of market participants sharply limits their ability to transact.

Clearing of OTC deals is one possibility, but it hasn't made major inroads. There is a coordination problem (clearing only adds value to the extent that most trading firms clear, so it is necessary to achieve a critical mass of participants—which is hard to do). Also, clearing isn't a panacea. In fact, clearing is primarily a means of collateralizing transactions, and many energy trading firms are short good collateral.

Other “Linear” Derivatives

Forward/contracts and futures are the simplest derivative. There are other derivative contracts that are effectively bundles of forwards. These are called swap contracts.

A “vanilla” swap contract is an agreement to exchange cash flows according to a fixed formula at certain dates in the future.

Parties to the swap agree upon (a) its maturity (“tenor”), (b) its frequency, (c) a fixed price, and (d) a reference index.

One party pays the fixed price. This is called the “fixed price payer.” The other party pays the floating price. This is called the “floating price payer.”

For instance, consider a monthly NG swap with a one year “tenor.”

A common price reference is the “prompt month” NYMEX natural gas futures price at the end of the month.

When the swap contract is initiated, the fixed price is set so that the swap has zero value. In the present case, the fixed price will be the average of the present value of the strip of 12 NYMEX futures contracts.

Each month, on the payment date, the fixed price payer owes the floating payer the fixed price, and the floating price payer owes the

fixed payer the floating price. The floating price is the NYMEX prompt month futures price on the payment date.

Swaps can be used for hedging. For instance, an independent gas producer typically sells its production at the prevailing market price, which goes up and down. To smooth out its revenues, it can enter into a swap as the floating price payer (and receive fixed).

Declines in the price of gas lower its sales revenues, but lead to higher payments received on the swap. In essence, the producer swaps a variable set of cash flows for a fixed set of cash flows, thereby hedging its risk.

The other party to the swap might be a firm (e.g., a gas utility) that buys its gas on the cash market at the prevailing market price. This party would want to receive floating payments that are high when gas prices are high and are low when gas prices are low to even out its net purchasing costs.

Although the producer and the utility could deal with one another, most swaps have a major dealer as a middleman. This is to alleviate credit/performance concerns.

Dealers often hedge their residual exposure through the futures markets.

Commodity Derivatives: Precious Metals

Cash and Carry Arbitrage

Call S_t the spot price of a commodity (silver, for instance) at time t . Moreover, $F_{t,T}$ is the futures price at t for delivery at time T , r is the riskless interest rate between time t and T , and s is the cost of storage between t and T . Assume that there is no benefit of holding inventories of the commodity (i.e., the rental/lease rate is zero). Consider the following set of transactions. Buy the spot commodity at t , store it until T . Finance this purchase and storage with borrowing. Sell the futures contract for delivery at T . Consider the cash flows from these transactions at dates t and T .

TRANSACTION	Date t	Date T
Buy spot	$-S_t$	S_T
Borrow	$S_t + s$	$-e^{r(T-t)}[S_t + s]$
Sell Futures	0	$F_{t,T} - S_T$
Pay storage	$-s$	0
Net Cash Flows	0	$F_{t,T} - e^{r(T-t)}[S_t + s]$

All of the prices that determine net cash flows at T are known and fixed as of t . Therefore, this transaction is riskless. Since this transaction involves zero investment at t , in order to avoid the existence of an arbitrage opportunity it must be the case that the cash flows at T are identically 0. Therefore, in order to prevent arbitrage, the following relation between spot and futures prices, interest rates, and storage charges must hold at t :

$$F_{t,T} = e^{r(T-t)}[S_t + s]$$

Reverse Cash and Carry Arbitrage

Cash and carry arbitrage involves buying the spot, borrowing money, and selling futures. Reverse cash and carry arbitrage involves selling the spot short, investing the proceeds, and buying futures. In order to short sell the spot asset, an investor borrows the asset and sells it on the spot market. The asset borrower must purchase the asset at T in order to return it to the lender. The cash flows from this transaction are:

TRANSACTION	Date t	Date T
Sell spot	$S_t + s$	$-S_T$
Lend	$-(S_t + s)$	$e^{r(T-t)}[S_t + s]$
Buy Futures	0	$S_T - F_{t,T}$
Net Cash Flows		
	0	$e^{r(T-t)}[S_t + s] - F_{t,T}$

Note that this implies the same arbitrage relation as cash and carry arbitrage.

Note: This analysis assumes that the arbitrageur holds inventories of the commodity. If he doesn't, he does not save on storage charges by selling the commodity. This drives a wedge (equal to the cost of storage) between the cash-and-carry and reverse-cash-and-carry arbitrage restrictions.

Implied Repo Rates

Arbitrage restrictions define relations between spot and futures prices and interest rates that must hold in order to prevent traders from earning riskless profits with no investment. Since there are many futures markets, in order to identify arbitrage opportunities in several markets simultaneously, it is convenient to convert spot-futures price relations into a common variable. Since for a given trader the same interest rate should apply to any arbitrage transaction, regardless of whether the transaction is in gold, silver, Treasury bonds or corn, the obvious common variable is an interest rate. Therefore, traders use spot and futures price to calculate an implied interest rate. This is commonly called the implied repurchase ("repo") rate, because the repo rate represents the rate at which most large traders can borrow or lend.

To calculate an implied repo rate, take natural logarithms of the basic arbitrage expression. This implies:

$$\ln F_{t,T} = r(T - t) + \ln[S_t + s]$$

Simplifying this expression, and recognizing that the difference between two logs equals the log of their ratio, produces:

$$\ln\left[\frac{F_{t,T}}{S_t + s}\right] / (T - t) = r^i$$

This is the interest rate implied by spot and futures prices. If this implied repo rate is lower than the actual repo rate at which a trader can borrow/lend, the trader can borrow cheaply through the futures market, and lend through the repo market. Conversely, if the implied repo rate exceeds the rate at which the trader can borrow/lend through the money market, he should lend through the futures market and borrow through the repo market. Moreover, by comparing implied repo rates from futures on different commodities (e.g., gold vs. S&P 500) a trader can identify cheap borrowing and rich lending opportunities.

The Pricing of Energy and Industrial Metal Futures: Spread Relations

Precious metal futures are pretty straightforward. Usually they sell at near full carry (i.e., in contango). Lease rates are typically small, and there is an active market for leasing precious metal inventories.

Pricing energy and industrial metal futures is a little more complicated. In particular, the relations between spot and futures prices, and between nearby and deferred futures prices are much more complex than is the case with precious metals.

For example, sometimes oil, heating oil, or copper trade at nearly full carry. Other times, the futures prices are far below spot prices—that is, the market is in “backwardation.” Moreover, the market can move from contango to backwardation to contango very quickly.

This raises the question: what determines the spreads between futures and spot prices for energy products and industrial metals?

The so-called “theory of storage” provides the answer to this question. In essence, this theory states that futures prices serve to ensure that consumption of these commodities is distributed efficiently over time.

Consider two situations:

Case 1. Supplies of the commodity (e.g., copper) are abundant. Delivery warehouses are full. Producers are operating with excess capacity.

Under these conditions, it makes sense to store some of the commodity. Abundant supplies are on hand, and to consume them all would glut the market today, and perhaps leave us exposed to a shortage in the future.

Traders will store the commodity (i.e., hold inventories) if and only if this is profitable. By storing (rather than selling) the commodity, the trader gives up the spot price (which he could capture by selling the commodity), and incurs interest and warehousing expenses. Storage is profitable only if futures/forward prices are above the spot price by the costs of interest and storage. That is, storage is profitable only if the market is at a carry/contango. If stores are huge, the market must be at full carry.

Case 2. Current supplies are very scarce relative to expected future production and supplies. That is, there is a temporary shortage. Producers are operating at or near capacity.

In these circumstances, it is foolish to store the commodity--it is scarce today relative to expected future supplies. Therefore, we should consume available supplies and hold close to zero inventories. (This exhaustion of inventories is sometimes called a "stock-out.")

Since the commodity is scarce today relative to what we expect in the future, the current spot price may be above the forward price. That is, the spot price must rise as high as necessary to ration the limited supplies. Moreover, there is no need to reward storage. Therefore, the market need not exhibit a carry. Thus, backwardation is possible.

There are some implications of this analysis:

Implication 1. The spread between spot and futures prices, adjusted for carrying costs, should depend upon the stocks held in inventory. When stocks are low, the market should be in backwardation; when stocks are high, the market should be at nearly full carry.

The industrial metals markets illustrate this clearly.

The theory also has implications for the volatility of prices, and the correlation between spot and forward/future prices. This is of great importance in risk management and option pricing applications.

Implication 2. When stocks are short prices should be more volatile. For those who are familiar with supply and demand analysis, a shortage implies that the supply curve is inelastic, i.e., steep. The steeper the supply curve, the more volatile spot prices. Since shortages are associated with backwardation, we expect very volatile prices in an inverted market (i.e., a market with backwardation), and less volatile prices when the market shows a carry.

Implication 3. Again calling upon supply and demand analysis, we know that short run supply curves are less elastic than long run supply curves. This implies that spot prices should be more volatile than forward/futures prices. Moreover, the difference in short run and long run supply elasticities should be greatest, the shorter are supplies today. Thus, the difference between spot volatility and forward volatility should rise as the market moves towards backwardation. Furthermore, this difference should be greater, the longer the time to delivery on the deferred contract.

Implication 4. When a stockout occurs, the cash-and-carry arbitrage link between spot and forward prices is broken. Under these conditions, spot and forward prices can move independently. When stocks are abundant, arbitrage ensures that spot and futures prices move together--the futures-spot spread equals the cost of carry. Thus, spot and forward/futures prices should be highly correlated when the market is at a carry, but may exhibit very low correlation when the market is in backwardation. Moreover, since a stockout is more likely, the longer the time to expiration of the deferred forward, the correlation between spot and forward/future prices should be lower, the greater the maturity of the forward/future.

Squeezes, Hugs, and Corners

Another factor which affects spot-futures spreads is a manipulation. A manipulation--sometimes referred to as a squeeze or corner, or a “hug” for a mild manipulation--occurs when a single trader accumulates a long futures/forward position that is larger than the physical supplies that can be delivered economically against these contracts. Rather than incur large costs to acquire deliverable supplies, as contract expiration approaches shorts are willing to buy back their contracts from the long at a premium.

This causes a unique pattern in prices. The price of the expiring future/forward that is being cornered rises--sometimes precipitously--relative to the deferred futures/forward price. Since the expiring future/forward and the spot price must converge during the delivery period, this means that the spot price rises relative to the deferred forward/future too. As soon as shorts close out their positions, the spot price collapses. The effect of final liquidation of a corner on spot prices is sometimes called the “burying the corpse effect.” A sharp increase in shipments to delivery warehouses also occurs.

Squeezes/corners are not unknown in commodity derivative markets. Exchanges and governments attempt to prevent or deter them, but they still occur from time to time. The experience of the zinc market in 1989 provides a good illustration of the effects of a squeeze on spreads.

Moral of the story: When trading commodities, if you are short you must always be aware of the possibility of a manipulation. Monitor futures market and cash market activity continuously to make sure that you are not unexpectedly caught in a squeeze.

Non-Storable Commodities

- There are also futures and forward contracts traded on non-storable commodities.
- These include: electricity, weather, bandwidth, live animals (e.g., hogs and cattle).
- Non-storability has a big impact on price dynamics and forward pricing.
- Storage mitigates price volatility—inventories are accumulated when demand is low and supply is high, thereby reducing the magnitude of price declines under these conditions, and are drawn from when demand is high and supply is low, thereby mitigating the magnitude of price increases.
- Without storage, inventories cannot “smooth” the effects of supply and demand shocks.
- This implies that prices of non-storables—notably electricity—can be **extremely** volatile.

Forward Prices for Non-Storables

- Due to non-storability, cash-and-carry arbitrage is impossible—you can't hold an inventory of electricity from the nighttime to the afternoon.
- This contributes to considerable intra-day variation in prices for non-storables with systematic intra-day variation in demand or supply.
- It also makes cash-and-carry arbitrage pricing methods—“preference free” pricing techniques—impossible for non-storables.
- Therefore, estimating forward prices (and forward curves) for non-storables must take a different tack.
- Forward price=expected spot price + risk premium

Electricity Forward Pricing

- For electricity, for markets in which good spot price and demand (load) data are available, such as PJM, California, or Australia, can estimate expected spot prices using traditional statistical techniques.
- Statistical distribution of demand can be estimated accurately.
- Use observed spot-price/load data to estimate a function that relates prices to load.
- Combine demand distribution and spot price/load relation to estimate an expected spot price.
- Estimating risk premium is trickier—need to use market forward price data.
- Essential to take risk premium into account when pricing power forwards because this risk premium is huge, especially on-peak.

Power Derivatives Markets

- Power derivatives markets have grown, albeit somewhat slower than had been anticipated in the mid-1990s.
- Virtually all power forward and options trading done over-the-counter.
- Exchange markets have languished.
- Heterogeneity of electricity (especially locational differences) have impeded development of liquid power forward markets—heterogeneity fragments liquidity.
- Credit issues (note effects of 1998 Midwest price spike, California crisis) have also impeded market development.
- Finally, lack of integration of financial and physical markets has impeded development.
- Until these issues are resolved, power derivatives trading may prove treacherous.

Weather Derivatives

- Weather derivatives are a new frontier in derivatives trading.
- Weather derivatives trading almost exclusively OTC, although there are exchange-listed contracts on the CME.
- Typical weather derivative product is based on heating- or cooling degree days.
- One heating degree day occurs when the low temperature is one degree below 65° for one day. A cooling degree day occurs when the high temperature is one degree above 65° for one day.
- Most weather derivatives are heating or cooling degree day options. For instance, you could have a contract that pays \$1 per cooling degree day times the difference between total cooling degree days in Chicago in July, 2002 and 250, if that difference is positive and zero if it is not.

- Weather derivatives can be used to manage quantity risk. For instance, the quantity of power or natural gas sold by an energy company depends on weather conditions. So does the price. Can use weather derivatives to manage this risk.
- Retailers can use them to manage risks due to weather. Retail sales for certain products are very sensitive to weather.
- Although degree-day options are the most common, weather derivatives can be based on rainfall, snowfall, or any other measurable weather variable.
- Like power derivatives, due to lack of storability arbitrage-based pricing of weather derivatives not possible. Need to utilize historical data on weather (corrected for global warming???) to estimate expected payoffs, then adjust by a risk premium.

Determining the Risk Premium

- For storable goods, our arbitrage analysis shows that the risk premium is irrelevant to determining the relation between spot and forward prices. That is, for these goods we can use “risk preference free” pricing to determine these relative prices. However, even for storable goods the risk premium affects the level of futures and spot prices, and their average movements through time.
- The earliest theory of the risk premium is due to Keynes. Keynes posited that hedgers are typically short futures. That is, they are typically holders of inventories of a commodity (e.g., corn) and they sell futures as a hedge.
- Since hedgers are net short, speculators must be net buyers in equilibrium (since total buys = total sells). Speculators will not absorb risk unless they are rewarded by profiting on average. Buying futures is profitable on average if futures prices rise on average. That is, speculators will enter the market only if the futures price is below the spot price expected at contract expiration. In this theory, the futures price tends to “drift up” over time—the speed of the drift measures the risk premium.

- In Keynes' theory, futures prices are "downward biased."
- In Keynes' theory, this downward bias makes short hedging costly (since short hedgers lose on average) so they will tend to hedge less than 100 percent of their risk.
- If long hedgers outnumber short hedgers (at a futures price equal to the expected spot price) futures prices must be upward biased to attract speculative interest.
- Thus, net hedging interest determines whether futures prices are upward or downward biased.
- For some markets (exchange traded futures) there is data on net hedging interest in the form of the CFTC's "Commitment of Traders Reports" available on the [www](http://www.cftc.gov).

The Magnitude of the Risk Premium

- The foregoing implies that the size of the risk premium depends on (a) net hedging demand at a price equal to the expected spot price, and (b) the risk aversion of speculators.
- The greater the hedging imbalance, the greater the risk premium (in absolute value).
- The more risk averse the speculators, the greater the risk premium (in absolute value).
- The risk aversion of speculators depends on (a) how well the futures market is integrated with the broader financial markets, and (b) the correlation between the futures price and movements in the market portfolio.
- Usually it is the case that risk premia will be smaller if the futures market is well integrated with broader financial markets because integration makes it possible for diversified speculators to participate in the market; diversification reduces speculator risk exposure.
- CAPM-type models imply that the higher the beta between the futures price and the overall market, the greater the upward drift in the futures price.

Estimating the Risk Premium

- The risk premium is the difference between the futures price and the expected spot price. Also, the risk premium affects the expected change in the futures price.
- Thus, estimation of the risk premium requires either estimation of the expected spot price, or estimation of the drift in futures prices.
- Since the risk premium affects both the costs of hedging and the benefits of speculation, both require estimation of expected spot prices. Thus, spot price forecasting is an important part of hedging and speculation.
- For most goods and commodities it is hard to estimate expected spot prices with accuracy. Electricity and weather may be exceptions.

Credit Derivatives

- Credit derivatives are another new type of financial contract. This is a rapidly growing market.
- Credit derivatives allow firms to hedge the risks of default on loans. For example, a bank that loans money to a customer loses if that customer defaults. The bank may be able to sell-off some portion of the credit risk without selling the loan, by entering into a credit derivatives transaction.
- Almost all credit derivatives are traded over the counter.
- There are a variety of credit derivative products in wide use.

Credit Event Derivatives

- A credit event derivative between party A and party B involves (a) a periodic payment from A to B, (b) the definition of a “credit event”, and (c) a payment from B to A if a “credit event” occurs.
- A credit event may be a bankruptcy or ratings downgrade, for instance. Although defining a credit event seems straightforward, it is not. The events in Russia in 1998 provide an illustration of the difficulties of defining a credit event.
- A may be a bank that wants to reduce its exposure to default by a particular borrower. B may be an insurance company or other financial firm that is willing to bear default risk.
- The size of the payment made in the event of a credit event is related to the impact of the event on the value of the loan.
- The size of periodic payment depends on the price the market charges to bear credit risk. Yield spreads between bonds of differing credit risk measure the market price of credit risk. Thus, the periodic payment should be related to yield spreads.

Total Return Swaps

- Total return swaps are another common type of credit derivative.
- In a total return swap, A and B exchange payments, where the magnitude of the payments swapped depend on the total returns on instruments of different credit-worthiness.
- For instance, the swap may involve A paying B the total return (interest plus capital gain/loss) on a BBB bond, and B paying A the LIBOR rate.
- This would make sense for A if he owned BBB bonds and didn't want to bear the credit risk. In essence, this deal passes the credit risk to B without selling the actual bond.

Why Use Credit Derivatives?

- It seems somewhat weird that firms would alter credit risk exposures through credit derivatives—why don't they just trade the underlying loans? That is, why would a bank that wants to reduce credit exposure enter a total return swap instead of just selling off credit-risky loans?
- Taxes, accounting, and regulatory arbitrage. Sales of loans may have adverse impact on taxes or reported earnings or the balance sheet. For instance, a gain or loss must be recognized on sale for tax or accounting purposes, but may not be recognized if credit risk is transferred through a credit derivatives transaction. Also, some intermediaries may operate under regulations that limit their ability to purchase below-investment grade bonds, but that do not limit their ability to use credit derivatives.
- Liquidity. Credit derivatives may be more liquid than the underlying securities. This is typically the case since most credit derivatives have shorter maturities (e.g., one year) than the underlying securities (e.g., five years). Information asymmetries are plausibly smaller for shorter term instruments, making them more liquid

OPTIONS BASICS

1. A call option gives the owner the right, but not the obligation, to buy the underlying asset (e.g., a stock, a bond, a currency, a futures contracts) at a fixed price. This fixed price is called the "strike price." Call options have a fixed expiration date. Time to expiration can range between days and years.
2. A put option gives the owner the right, but not the obligation to sell the underlying asset at a fixed price.
3. There are two basic types of options: European and American. The holder of a European option can exercise it only on the expiration date. That is, there is no early exercise for European options. In contrast, the holder of an American option can exercise it any time prior to the expiration date as well as on the expiration date itself. Most exchange traded options are American. Many over the counter options are European.
4. Energy options traded on NYMEX and IPE.
5. OTC energy options markets are also important.
6. Many contracts have options embedded in them. Power supply contracts in particular have many embedded options. Many NG contracts also embed options (e.g., swing options).

HEDGING WITH OPTIONS

Recall that futures and forward (or swap) hedges allow you to lock in a price (abstracting from basis risk).

Options allow you to place a ceiling or a floor on price exposure.

The seller of a commodity can buy a put to place a floor on its revenues. The strike price of the put (minus the cost of buying it) is a floor on the price that the seller will receive. For instance, a seller of natural gas buys a put struck at \$5.00 per MMBTU for \$.25/MMBTU. The put expires in April. If the price of NG in April is \$4.25, the firm exercises its put, gets \$5.00, and delivers the gas. Net of the option premium, it receives \$4.75. The same is true if the price of gas is \$4.00 in April. However, if the price of gas is above \$5.00, the firm gets to pocket the entire amount.

The buyer of a commodity can buy a call to place a ceiling on its costs. A utility buying gas can buy a call struck at \$6.00 (say) that expires in January. If the price of gas is above \$6.00, the firm exercises its call, receives the gas, and pays \$6.00 plus the premium regardless of how high the price rises. Conversely, if the price of gas is below \$6.00 in January, the firm lets the call expire unexercised and benefits from the lower gas price. Thus, the firm caps its price exposure at \$6.00 plus the option premium.

Options can be used in various combinations to create customized payoff patterns. For instance, a gas seller can reduce the cost of obtaining downside protection by giving up some upside potential. It does so by buying a put (which costs money) but selling a call (which generates a cash inflow). This is called a “collar” or a “range forward.”

ARBITRAGE RESTRICTIONS

Futures/Forward options prices must obey certain arbitrage restrictions.

1. Put-call parity. Consider *European* options forward contract with current forward price F_t . A position consisting of a call struck at K and a short put with strike K provides the same payoffs as a forward contract with a forward price equal to K . Both options expire at time T . Thus, the present value of a forward contract with a forward price of K equals $e^{-r(T-t)}(F_t - K)$. Since the long call-short put position gives the same payoff as a forward contract, it must be the case that the value of this position equals the value of the forward contract. That is:

$$c(F_t, K, t, T) - p(F_t, K, t, T) = e^{-r(T-t)}(F_t - K)$$

This is called the put-call parity relation.

2. Given put and call prices, the forward price, the strike price, and the time to expiration, it is possible to solve the put-call parity expression for an implied interest rate in order to determine whether these prices present an arbitrage opportunity.

EARLY EXERCISE OF OPTIONS

Early exercise has one clear disadvantage: by exercising an option, a trader gives up the value of the option over its remaining life. This value must be positive. Therefore, early exercise is desirable if and only if there is some off-setting benefit. There is only one possible source of such a benefit. Namely, there is a potential value to early exercise if this allows the owner of the option to receive cash flows earlier.

Case 1. Call on a non-dividend paying stock.

Assume the owner of a call with strike price K and time to expiration T exercises the option at $t < T$, and borrows money in order to pay the strike price. Then, at T , the wealth of the trader is $S_T - e^{r(T-t)}K$. If the trader had not exercised the option, his wealth at T would equal $\max[0, S_T - K]$. It is clear that the trader's wealth is greater if he does not exercise early. This is true for two reasons, first, if the price of the stock falls below the strike price between t and T , the trader won't exercise the option at T , and thus isn't "stuck" with a less valuable stock. Second, by deferring exercise, the call owner saves the interest on the strike price over the period t to T . Conclusion: Never exercise a call on a non-dividend paying stock early.

Case 2. Call on a dividend paying stock.

Assume the firm pays a dividend equal to D at $t < T$. If the call owner exercises at t , his wealth at T equals:

$$S_T - e^{r(T-t)}(K - D)$$

This may (but may not) exceed $\max[0, S_T - K]$ because by exercising the call early, the owner receives the dividend. Therefore, unlike the case with a call on a non-dividend paying stock, we can't say for certain that the payoffs to early exercise are always lower than the payoffs from exercise at expiration. If this dividend is large enough, it may exceed the interest on the strike price and the option value foregone from early exercise. Thus, you may exercise a call on a dividend paying stock early. If you do exercise early, you will only do so immediately before the payment of a dividend (i.e., on the cum dividend date).

If the future value of the dividend is smaller than the value of the interest paid on the strike price from t to T , early exercise will not be optimal. The value of this interest equals $e^{r(T-t)}K - K$. If $e^{r(T-t)}K - K > e^{r(T-t)}D$ then

$$S_T - e^{r(T-t)}(K - D) < S_T - K < \max[0, S_T - K]$$

and therefore early exercise is not profitable.

Intuitively, this means that if the dividend received by exercising early is not large enough to compensate the option

holder for the interest incurred on the strike price due to early exercise, early exercise cannot be optimal.

Case 3. Put on a non-dividend paying stock.

Consider a trader who owns a put and a share of the underlying stock. If he exercises the put at $t < T$ and invests the strike proceeds, his wealth at t equals $e^{r(T-t)}K$. There are some values of the stock price such that this exceeds the trader's wealth at T if he does not exercise the put at t , $\max[0, K - S_T] + S_T$. By exercising early, the put owner receives cash flows (i.e., the proceeds from exercise) earlier; the interest earned on the strike price over this interval of time. This may compensate the trader for the gains he would earn by holding onto the stock if the price of the stock were to rise above the strike price between t and T . Early exercise of a put can be optimal at any time prior to expiration.

Case 4. Put on a dividend paying stock.

All else equal, dividend payments cause the price of the stock to decline. Therefore, if the put is not protected against dividend payments, dividends tend to induce the holder of a put to defer exercise in order to take advantage of this predictable price decline. It still may be the case, however, that the advantages of receiving the proceeds from exercise early (i.e., the interest earned on the strike price) offsets this effect.

Case 5. Puts and calls on futures contracts.

When the holder of an option on a futures contract exercises, she receives a cash payment equal to the difference between the strike price and the futures price at the time of exercise and a futures position. For example, if the futures price at exercise equals F_t , the holder of a put receives a cash payment equal to $K - F_t$ and a short futures position at the market price. The value of the short position equals 0. Early exercise of either a put or a call therefore leads to an acceleration of cash flows. Therefore, one cannot rule out early exercise of futures options.

Case 6. Calls on foreign currency.

Foreign currencies can be invested at interest; recall that this is equivalent to a continuous dividend yield. Therefore, there may be advantages to exercising a call on a foreign currency early in order to receive this cash flow over a longer period of time.

Early exercise requires a modification in the put-call parity expression. In particular, we can no longer derive an equality restriction, but only an inequality restriction.

First consider the case of American puts and calls on non-dividend paying stocks. We know that the value of an American call on a non-dividend paying stock equals the value of a European call on the same stock (with the same strike and time to expiration). Moreover, we know that an American put may be exercised prior to expiration. This option to exercise early has value, so an American put must be more valuable than a European put.

Therefore,

$$C(F_t, K, t, T) - P(F_t, K, t, T) < e^{-r(T-t)} (F_t - K)$$

A BINOMIAL MODEL OF FUTURES PRICE MOVEMENTS

In order to derive a formula for futures option prices, we must first specify a model that describes how futures prices move.

Our first model is a so-called binomial model because at any point in time, it is assumed that the percentage change in the futures price--the stock "return"--can take only two values, $u > 1$ (an "up" move) or $d < 1$ (a "down" move). The probability of an up move equals q , and the probability of a down move equals $1 - q$. We will see that these actual probabilities are irrelevant to the pricing of options.

The binomial model divides time between the present and the expiration date of an option into discrete intervals of equal length. Each interval is Δt in length. By allowing more intervals between now and expiration, this interval becomes shorter. In the limit, with an infinite number of intervals, the length of an interval becomes vanishingly small, and equal to dt .

That is, if the futures price equals F today, it may equal either $uF > F$ or $dF < F$ at the end of the next interval of time.

It is possible to show that as the length of a time interval becomes vanishingly small, the distribution of futures prices is lognormal. We will utilize this fact in deriving the Black formula for pricing options.

USING THE BINOMIAL MODEL TO PRICE AN OPTION

Consider a portfolio of a short position in a single call on a futures, and Δ shares of the futures underlying the call. Assume the value of the call in Δt units of time equals c_u if the futures price rises over this interval, and equals c_d if the futures price falls. Thus, at $t+\Delta t$ the value of the portfolio equals $\Delta uF - c_u$ if the futures price rises, and equals $\Delta dF - c_d$ if the futures price falls.

Note that we can choose Δ to make the value of the portfolio the same regardless of whether the futures price rises or falls.

Formally, choose Δ such that

$$\Delta uF - c_u = \Delta dF - c_d$$

or

$$\Delta = \frac{c_u - c_d}{F(u - d)}$$

Since the value of the portfolio doesn't depend upon the futures price change, the portfolio is riskless. Therefore, the return on the portfolio must equal the risk free rate. That is, if the value of the call today equals c , then:

$$e^{r\Delta t}(-c) = \Delta F(u-1) - c_u = \Delta F(d-1) - c_d$$

This expression is somewhat different from that relevant for a futures option. Buying shares to hedge the option costs money. Buying the forward/futures contract requires no initial outlay of funds.

Solving for c implies:

$$c = e^{-r\Delta t} [pc_u + (1-p)c_d]$$

where

$$p = \frac{1-d}{u-d}$$

Note that the value of the call at t is equal to the expected present value of the call at $t+\Delta t$, using p to measure the probability of an up move and $(1-p)$ to measure the probability of the down move.

It is essential to recognize that p is not equal to q , the *true* probability of a futures price increase. Instead, p is the probability of an up move such that today's futures price equals what the futures price is expected to be at $t+\Delta t$. This would occur in a market where traders are risk neutral. That is, if there is no risk premium.

This analysis implies that we can price options as if we are in a risk neutral world! Put differently, we don't need to know the expected change in the futures price to price options on that futures. This is true because we can form a portfolio consisting of the future and the option to eliminate all risk.

In order to determine the price of the call, we work backwards from the end of the binomial tree because at the end of the tree we know the values of the option.

That is, we use “backward induction” repeatedly to value an option on an underlying future. We go backwards because we know the payoffs to the option at the expiration date, and can therefore apply our formula repeatedly by proceeding from the end of time to the beginning.

The main choice you must make in establishing a binomial tree is for u and d . We choose these parameters such that the theoretical value of the variance of the stock given by the binomial model equals the actual value of the variance of the future we are interested in.

Remember that the variance of a future’s return equals the expected value of the squared deviation between the realized return and the expected return. The expected return in our risk neutral model equals 0. The return in an upmove equals u and the return in a down move equals d . Therefore, the variance equals:

$$p(u - 1)^2 + (1 - p)(d - 1)^2 =$$
$$p(1 - p)(u - d)^2 = \sigma^2 \Delta t$$

where σ is the actual standard deviation of the future’s return.

The first equality follows from the fact that

$$F = puF + (1 - p)dF$$

We have already found the relation between p and u and d . Thus, we have one equation in two unknowns. We eliminate one unknown by choosing $d=1/u$.

If we solve all of this for u , we get: $u = e^{\sigma\sqrt{\Delta t}}$ if Δt is small.

DERIVATION OF THE FORMULA FOR p .

First, define $a = e^{r\Delta t}$.

Then:

$$\Delta F - u\Delta F + c_u = ac$$

Substituting for Δ implies:

$$\frac{(c_u - c_d)}{u - d} - \frac{u(c_u - c_d)}{u - d} + c_u = ac$$

Gathering terms with c_u and c_d :

$$\frac{c_u}{u - d} (1 - u + u - d) + \frac{c_d}{u - d} (u - 1) = ac$$

$$\frac{1 - d}{u - d} c_u + \frac{u - 1}{u - d} c_d = ac$$

Define $p = (1 - d)/(u - d)$ then note that

$$\frac{u - 1}{u - d} = 1 - \frac{1 - d}{u - d} = 1 - p$$

Thus,

$$ac = pc_u + (1 - p)c_d$$

Dividing both sides by a produces the expression presented earlier. It is important to note that if the probability of an upmove in the futures price equals p , then the expected futures price change is zero.

To see why, note that if the probability of an upmove equals p , the expected value of the future next period equals:

$$\begin{aligned} F[pu + (1 - p)d] &= \\ F[p(u - d) + d] &= \\ F\left[\frac{1 - d}{u - d}(u - d) + d\right] &= \\ F & \end{aligned}$$

This implies that the futures price does not drift up or down. In other words, there is no risk premium.

Use of the Binomial Model

A major advantage of the binomial model is that it can be used to value American options for which early exercise is possible, and hence valuable. An example illustrates this.

First consider the example of puts on a futures contract.

In this example, $\sigma=.40$, $T-t=.4167$, $r=.1$, $F=30$, $K=30$. If we divide the time until expiration into 5 segments, we get $\Delta t=.4167/5=.0833$.

Given these values, we can determine $p=.4712$, and $1-p=.5288$.

Also $u=1.122$, $d=.8909$, and $e^{-r\Delta t} = 1/1.0084 = .9917$.

Let's first value a European put.

We start by creating our tree of futures prices. Each node on the tree corresponds to the futures price at some date in the future. The dates we look at are equally spaced.

We look at 5 time steps to expiry, so there are six possible futures prices at expiration, ranging from 53.44 to 16.84.

We now start at these expiration date futures prices and see what the payoff to our put is for each of them. When the futures price is 53.44, we aren't going to exercise, so here the payoff is 0. When the futures price is 16.84, we will exercise, and get $30 - 16.84 = 13.16$.

We now step Δt (i.e., one time step) closer to the present and apply our formula. For instance, when the futures price 4 weeks from now has gone up 4 weeks in a row, and equals 47.61, the formula gives us a value of zero, because regardless of whether the futures price goes up or down in the next step, the value of the option will be zero.

However, if the futures price has gone up two times and down two times (and the order doesn't matter) and the futures price is 30 again, the payoff to the option is zero if the futures price goes up in the last week, but is 3.27 if the futures price goes down. Therefore, at this node the value of the option is

$$.9917[.4711*0 + .5288*3.27] = 1.72.$$

We perform this exercise for each of the 5 possible futures prices 4 periods from now.

Given these 5 values, we now proceed to 3 periods from the present and again apply our formula. Working back this way gives our option value, which in this case is 4.54.

The procedure for an American option is pretty similar, except that it is necessary to perform an additional check at each node of the tree. Specifically, we check whether it is better to hold the option at that node, or exercise it.

To determine the value we get from holding the option, we apply our formula. For instance, when the futures price is 30 in 4 periods, the formula tells us that the value of the option is 1.72. If we exercise, we get nothing. Therefore, we'll hold. We then plug in 1.72 as the value of our option as we proceed back to 3 periods prior to expiry.

Conversely, if the futures price is 18.90 one period before expiration, the value we get from holding the option is:

$$.9917[.4711*8.78+.5288*13.16] = 11.00.$$

Exercising the option generates a payoff of $30 - 18.90 = 11.10$. It's better to exercise than hold, so we exercise early. We use 11.10 as the value of the option after 4 down moves when proceeding backwards through the tree.

This demonstrates how we can use the binomial option to calculate American option prices.

Analytically, the binomial model is straightforward. In order to increase the accuracy of this approach, however, it is necessary to make Δt fairly small by increasing the number of time intervals. This can increase the cost of computing options prices using the binomial method.

Thus, when pricing an option, we face a trade off: the binomial method can handle early exercise easily, but is computationally cumbersome.

This raises the question, is there a more computationally tractable model?

The answer is yes: If we are willing to consider only European options, it is possible to produce an option pricing model that is very easy to use--the Black model.

THE BLACK MODEL FOR FUTURES OPTIONS

The Black-Scholes model is essentially the same as the binomial pricing model when the number of time intervals approaches infinity, i.e., as Δt becomes arbitrarily close to zero. This is sometimes called a "continuous time" model in contrast to the "discrete time" binomial model because we no longer divide the time line into several discrete periods, but instead consider time as a continuum.

It is possible to show that as the number of time steps approaches infinity, the return on the underlying future obeys a normal distribution. The normal distribution is simply the well-known bell shaped curve.

Recall that the return on a stock equals the percentage change in price on the future. Also note that over a very small time interval dt the return on a stock equals:

$$\ln F_{t+dt} - \ln F_t$$

Thus, if the return on the futures in the continuous time world is normally distributed, then the futures price in the future is lognormally distributed.

It is also essential to remember that in valuing options we can assume that the expected change in the futures price is zero.

Again, this is because we can construct a portfolio including the future and the option that is riskless.

This can be represented formally. If the stock price at T (which may be the expiration date of an option) is lognormally distributed, then we can write:

$$F_T = F_t e^{(-.5\sigma^2)(T-t) + \sigma\sqrt{T-t}Z}$$

where Z is a normally distributed variable with expected value (i.e., mean) equal to 0.

(You can check that this expression is correct by taking the natural logs of both sides. You will find that the difference in the logs is normally distributed because Z is normally distributed.)

We can now value a European option that expires at T . Recall that the value of the option at t is the expected present value of the payoffs of the option at T , where we use the riskless interest rate to discount these payoffs. Also remember that any expected value is the sum of the possible payoffs multiplied by the probability of receiving a given payoff. In our analysis, the size of the payoff depends upon the realization of Z because Z determines the stock price at T . Moreover, because Z is normally distributed, the probability of observing any given Z is

$$n(Z) = \frac{e^{-.5Z^2}}{\sqrt{2\Pi}}$$

Note that the Π in this expression is the mathematical constant
Pi=3.14159 . . .

We can use this information to value an option. Consider a call option with strike price K . We know that the option's payoff is positive if and only if:

$$F_T = F_t e^{(-.5\sigma^2)(T-t) + \sigma\sqrt{T-t}Z} \geq K$$

There is a value of Z , call it Z^* , such that this expression holds with equality. This is the "critical value" of Z : For larger Z , the call is in the money at expiration, for smaller Z , the option is out of the money. Therefore:

$$\ln F_t + (-.5\sigma^2)(T-t) + \sigma\sqrt{T-t}Z^* = \ln K$$

Solving for Z^* implies:

$$Z^* = \frac{\ln(K / F_t) - (-.5\sigma^2)(T-t)}{\sigma\sqrt{T-t}}$$

Given this result, the expected present value of this option's payoffs (and hence its price) is:

$$c = e^{-r(T-t)} \left\{ \int_{Z^*}^{\infty} [F_T(Z) - K]n(Z)dZ + \int_{-\infty}^{Z^*} 0n(Z)dZ \right\}$$

This expression holds because the payoff to the option is 0 when $Z < Z^*$.

Doing a little substitution, we get

$$c = e^{-r(T-t)} \int_{Z^*}^{\infty} [F_t e^{(-.5\sigma^2)(T-t) + \sigma Z \sqrt{T-t}} - K] \frac{e^{-.5Z^2}}{\sqrt{2\pi}} dZ$$

Consider the exponentials. We can add the exponential terms to get the following exponent:

$$\begin{aligned} & (-.5\sigma^2)(T-t) + \sigma Z \sqrt{T-t} - .5Z^2 = \\ & -.5[\sigma^2(T-t) + Z^2 - 2\sigma Z \sqrt{T-t}] \end{aligned}$$

Define a new variable

$$y = Z - \sigma \sqrt{T-t}$$

Note that y is normally distributed (because Z is). Moreover,

$$y^2 = \sigma^2(T-t) + Z^2 - 2\sigma Z \sqrt{T-t}$$

In addition, the call option is in the money if

$$y = Z - \sigma\sqrt{T-t} \geq Z^* - \sigma\sqrt{T-t}$$

$$y \geq Z^* - \sigma\sqrt{T-t} =$$

$$\frac{\ln(K / F_t) - (-.5\sigma^2)(T-t)}{\sigma\sqrt{T-t}} - \sigma\sqrt{T-t} =$$

$$\frac{\ln(K / F_t) - .5\sigma^2(T-t)}{\sigma\sqrt{T-t}} \equiv y^*$$

Then we can rewrite the first term in our integral as

$$e^{-r(T-t)} \int_{y^*}^{\infty} F_t \frac{e^{-.5y^2}}{\sqrt{2\Pi}} dy =$$
$$e^{-r(T-t)} \int_{-\infty}^{-y^*} F_t \frac{e^{-.5y^2}}{\sqrt{2\Pi}} dy$$

This is just equal to $N(-y^*)$, where $N(x)$ gives the area under a normal distribution curve to the left of x . Thus, the first term in our integral is:

$$e^{-r(T-t)} F_t N\left[\frac{\ln(F_t / K) + .5\sigma^2(T-t)}{\sigma\sqrt{T-t}}\right]$$

Now consider the second term in our integral:

$$e^{-r(T-t)} K \int_{Z^*}^{\infty} n(Z) dZ =$$

$$e^{-r(T-t)} K \int_{-\infty}^{-Z^*} n(Z) dZ =$$

$$e^{-r(T-t)} KN(-Z^*) =$$

$$e^{-r(T-t)} KN \left[\frac{\ln(F_t / K) - .5\sigma^2 (T-t)}{\sigma \sqrt{T-t}} \right]$$

Putting the two terms together produces the Black formula for the call:

$$c = e^{-r(T-t)} [F_t N(d_1) - KN(d_2)]$$

where

$$d_1 = \frac{\ln(F_t / K) + .5\sigma^2(T-t)}{\sigma\sqrt{T-t}}$$

and

$$d_2 = d_1 - \sigma\sqrt{T-t}$$

We can then use put-call parity to get the value of a put. Recall:

$$c - e^{-r(T-t)} (F_t - K) = p$$

Then

$$p = e^{-r(T-t)} [F_t (N(d_1) - 1) + K(1 - N(d_2))]$$

Note that $N(x) = 1 - N(-x)$. Therefore,

$$p = e^{-r(T-t)} [KN(-d_2) - F_t N(-d_1)]$$

OPTIONS ON OTHER ASSETS

One can rewrite the Black-Scholes formula for a call as follows:

$$c = e^{-r(T-t)} [e^{r(T-t)} S_t N(d_1) - KN(d_2)]$$

Note that the term multiplying $N(d_1)$ is the forward price of a non-dividend paying stock.

One can also rewrite the term inside the normal function in the model as:

$$d_1 = \frac{\ln(Se^{r(T-t)} / K) + .5\sigma^2(T-t)}{\sigma\sqrt{T-t}}$$

Note that the numerator in the logarithm term is also the forward price of a non-dividend paying stock. In general, it is possible to show that to find the value of a call on some other type of asset, one simply uses the forward price of the asset wherever one finds the forward price of the stock in the Black-Scholes equation. For example, for a European call on a dividend paying stock, we use

$$d_1^* = \frac{\ln[(Se^{r(T-t)} - FVD) / K] + .5\sigma^2(T-t)}{\sigma\sqrt{T-t}}$$

where FVD is the future value of the dividend on the stock. As a result, we can write the price of a European call on a dividend paying stock as:

$$c = e^{-r(T-t)} \{ [e^{r(T-t)} S_t - FVD] N(d_1^*) - KN(d_2^*) \}$$

We can use similar reasoning to price options on other assets. For example, consider an option on a futures or forward contract. The forward price of the futures (or forward) price is just the futures (forward) price itself. Therefore, if the futures price equals $F_{t,T}$ then the price of a call option on this future is:

$$c = e^{-r(T-t)} [F_{t,T} N(d_1') - KN(d_2')]$$

where

$$d_1' = \frac{\ln(F_{t,T} / K) + .5\sigma^2(T-t)}{\sigma\sqrt{T-t}}$$

This formula is called the “Black Model.”

Finally, consider the value of an option on a foreign currency. If the foreign interest rate is r_F and the spot price of the currency is S_t then

$$c = e^{-r(T-t)} [e^{(r-r_F)(T-t)} S_t N(d_1'') - KN(d_2'')]$$

where

$$d_1'' = \frac{\ln(S_t / K) + (r - r_F + .5\sigma^2)(T - t)}{\sigma\sqrt{T - t}}$$

DYNAMIC HEDGING: DELTA, GAMMA, AND VEGA

The Black-Scholes model also describes how the price of the option changes when certain underlying variables change. For example:

$$\frac{\partial c}{\partial F_t} = e^{-r(T-t)} N(d_1) \equiv \Delta$$

The interpretation for this "Delta" is exactly the same as the interpretation of the Δ in the binomial model: It gives the number of shares of stock to hold with a short call position in order to form a riskless position. To see why, note that if one holds Δ shares of stock and one short call, the change in the value of the portfolio is

$$\frac{\partial(-c + e^{-r(T-t)} \Delta F_t)}{\partial F_t} = e^{-r(T-t)} [-N(d_1) + N(d_1)] = 0$$

Note that the value of this position doesn't change even if the futures price changes. Thus, it is riskless.

Put differently, if one wants to hedge a single long call, one *sells* Δ futures.

It is possible to show that if one changes Δ continuously in response to changes in the futures price, one's position will be riskless. This is called "dynamic hedging." It is necessary to change Δ every instant for this strategy to work perfectly.

Of course, in the real world, it is impossible to make such adjustments. Moreover, transactions costs make this strategy prohibitive. Therefore, it is necessary to make adjustments less frequently.

It is desirable to know, however, when it is necessary to make changes most frequently. The "Gamma" of an option provides this information.

Formally,

$$\Gamma = \frac{\partial^2 c}{\partial F_t^2} = \frac{\partial \Delta}{\partial F_t}$$

That is, Γ tells you how rapidly Δ changes when the futures price changes. It quantifies the curvature of the option price function.

This is important information for a trader managing the risk of an option position. Left untended, a position with substantial $\Gamma > 0$ may become overhedged (i.e., the trader's actual stock position becomes substantially larger than the optimal Δ) or may become

unerhedged (i.e., the traders actual futures position becomes substantially smaller than the optimal Δ) in response to even small changes in the stock price. Put differently, you are likely to need to change a hedge position quickly if there is a lot of Gamma. Gamma neutral hedgers (i.e., hedgers with positions with $\Gamma=0$) can sleep better than hedgers with substantial Gamma (either positive or negative).

The other important variable a hedger needs to track is Vega. This gives the sensitivity of the option price to changes in σ . That is,

$$Y = \frac{\partial c}{\partial \sigma}$$

Remember that Black assumes that volatility is constant over time. In reality, however, volatility can change rapidly. If so, a trader with substantial Vega is subject to large gains or losses. Just as a hedger wants to know his vulnerability to futures price moves, he should also track his risk to volatility changes.

These "Greeks" are important for another reason. Consider a trader who wants to purchase or sell an option that is not traded on an organized exchange, or on the OTC market. (For example, an individual might want to buy a 5 year call on WTI). Even though the trader cannot trade the option, he can replicate the payoffs of the option synthetically by dynamically buying or selling the futures underlying the option.

That is, at every instant of time, the trader holds Δ units of the future underlying the option, where Δ gives the delta of the option he wants to replicate. If the trader can adjust the position continuously and costlessly, such a strategy will produce cash flows at option expiration that are exactly identical to those of the option.

You may also want to replicate an option that is traded. For example, if you want to buy a put, but you think puts are overpriced in the market (e.g., the implied volatility of traded puts is greater than your forecast of volatility), you may wish to replicate the put through a dynamic trading strategy. Similarly, you can arbitrage the market by buying (selling) underpriced (overpriced) options, and dynamically replicating off-setting positions.

In reality, it is impossible (and very costly!) to adjust the portfolio continuously in this fashion. This is particularly true if the option is near the money, so has a lot of Gamma. Moreover, the dynamic hedger is vulnerable to changes in volatility. The problems faced by "portfolio insurers" (traders who attempted to replicate calls on the stock market through dynamic hedging techniques) on 10/19/87 provides a graphic illustration of the potential problems.

In order to address these problems, the trader can add traded options to his portfolio. The trader's objective is to match the Delta, Gamma, and Vega of his portfolio to the Delta, Gamma, and

Vega of the option the trader desires to replicate. This reduces the amount of adjustment needed, and provides some insurance against volatility shocks.

For example, consider replicating an option with $\Delta=\Delta^*$ and $\Gamma=\Gamma^*$. You can use the underlying stock and an option with $\Delta=\Delta^{**}$ and $\Gamma=\Gamma^{**}$ to construct your replicating portfolio. To get a delta and gamma match, you have two equations and two unknowns. The unknowns are N_s , the number of shares of stock to trade, and N_o , the number of options to buy or sell as part of your hedge portfolio. Formally, solve:

$$\Delta^* = N_s + N_o \Delta^{**}$$

and

$$\Gamma^* = N_o \Gamma^{**}$$

You could also construct a delta, gamma, and vega matched position. To do so require two options and the stock. You have to solve three equations in three unknowns here.

Note that even a gamma matched or gamma and vega matched position must be adjusted over time. Recently, researchers have developed strategies that allow you to create “fire-and-forget” hedges. That is, using these techniques, you can replicate an option by constructing a portfolio at the initiation of the trade, and never adjusting the portfolio until the expiration date of the option you want to replicate. As you might guess, this requires you to trade in a large number of options.

FORMULAE FOR GAMMA, VEGA, AND THETA
(The π in the formulae is the mathematical constant 3.14159)

The Gamma of a Euro futures call or put equals:

$$\Gamma = \frac{e^{-r(T-t)} e^{-.5d_1^2}}{F\sigma\sqrt{2\pi(T-t)}}$$

The Vega of a Euro futures call or put equals:

$$\Lambda = e^{-r(T-t)} F\sqrt{T-t} \frac{e^{-.5d_1^2}}{\sqrt{2\pi}}$$

The Theta of a Euro call equals:

$$\Theta = -\frac{e^{-r(T-t)} F e^{-.5d_1^2} \sigma}{2\sqrt{2\pi(T-t)}} - rKe^{-r(T-t)} N(d_2)$$

The Theta of a Euro put equals:

$$\Theta = -\frac{e^{-r(T-t)} F e^{-.5d_1^2} \sigma}{2\sqrt{2\pi(T-t)}} + rKe^{-r(T-t)} N(-d_2)$$

Note that Theta is unambiguously negative for a Euro call, but may be positive for a Euro put.

For a call, Gamma, Delta, and Theta are related by this equation:

$$rC = \Theta + .5\sigma^2 F^2 \Gamma$$

A similar expression holds for a put.

Using the Black Model

- Most of the inputs to the Black model are observable. These include: the strike price, the underlying price, the interest rate, and time to expiration. One key input is **not** observable—the volatility.
- Good option values depend on good volatility values.
- Where do volatility values come from? There are two sources: historical data and implied values.
- Historical volatilities are calculated using historical data on returns. For instance, to calculate the volatility for NGk, one would plug historical data on NG percentage price changes into a spreadsheet and calculate the standard deviation (adjusted for the data frequency).
- If the Black model were strictly correct, this approach would be the right one. Moreover, if the Black model were strictly correct—that is, if volatility was truly constant—you would want to get as much historical data as is available to get the most precise estimates of volatility possible.
- An alternative approach is to use “implied volatility.” In this approach, you find the value of σ that sets the value of the option given by the Black formula equal to the value of the option observed in the marketplace.

- Implied volatility can be estimated with extreme accuracy using the “Solver” function of Excel. There are very good approximations that you can use to determine the implied volatility for an at-the-money option.
- Obviously, if one uses the implied volatility, one is assuming that (a) the model is correct, and (b) the option is correctly priced. Thus, implied volatility cannot be used to determine whether a particular option is mis-priced. So what is it good for?
- If the Black model is correct, all options should have the same implied vol. Therefore, comparing implied volatilities across options allows one to determine whether options on the same underlying are mis-priced relative to one another. Buy the option with low volatility, sell the option with high volatility.
- Also, you can use the implied volatility to get more accurate measures of the “Greeks.”
- In practice, most options are quoted using an implied volatility (that is then plugged into the Black formula). Therefore, most options traders think in vol terms rather than in price terms, and option trading is essentially volatility trading.

How Well Does the Black Model Work?

- The Black model is the most influential theory ever introduced in finance, and perhaps in all of the social sciences. Despite this success, it has its weaknesses. Practitioners have learned how to patch the holes in Black.
- We know that the Blackmodel is not strictly correct. First, contrary to the model's predictions, there are systematic differences in implied volatility by strike price and time to expiration.
- Most important, deep in-the-money and deep out-of-the-money options have consistently different volatilities than at-the-money options. This is referred to as the “volatility smile.”
- This implies that σ cannot be constant.
- This is confirmed by empirical evidence. This evidence shows that volatility varies systematically over time. Moreover, volatility “clumps”—big price moves tend to follow big price moves, and small price moves tend to follow small price moves.

Dealing With the Smile

- There are two basic approaches to the smile.
- The first approach posits that volatility is a function of the underlying price and time. That is, $\sigma(F, t)$. Using very advanced techniques, one can find such a function that fits a variety of options prices. (Many practitioners use a crude approach that gives a very unreliable and unstable estimate of this function).
- The second approach is to post that volatility is random and to write down a specific stochastic process for volatility.
- The second approach is probably more realistic, but raises many difficulties. Most important, volatility risk cannot be hedged using the underlying. Consequently, any pricing formula must include a volatility risk premium.
- Some practitioners and academics assume that this risk premium is zero. This results in convenient pricing formulae, but is inconsistent with a great deal of evidence. Most notably, options hedged against moves in the underlying price earn a positive risk premium—this wouldn't happen if vol risk premia equal zero.
- Thus, the smile is a knotty issue that most practitioners address in an *ad hoc* manner.

How Exotic

- So far we have considered “vanilla” options—basic puts and calls.
- Other kinds of options—“exotics”—are traded in the OTC market.
- Although the variety of exotics is limited only by the imagination of traders, some exotics are more common than others.
- Sometimes exotics are bundled implicitly in securities.

Digital or “Bet” Options

- A digital option pays a fixed amount of money if a certain event occurs.
- For instance, a digital call on Microsoft struck at \$75 that pays \$10, pays \$10 if the price of MSFT at expiration exceeds \$75, regardless of whether the price at expiration is \$75.01 or \$175. Similarly, a digital put struck at \$75 that pays \$10, pays \$10 if the price of the underlying at expiration is \$74.99 or \$0.
- The value of a digital is easy to determine. A digital call value is $e^{-r(T-t)} PN(d_2)$ where P is the payoff and d_2 is the same as in the B-S formula. Similarly, a digital put value is $e^{-r(T-t)} PN(-d_2)$.
- Digitals are traded in the OTC market, but perhaps the most (in)famous example of digital options was embedded in bonds bought by Orange County CA in the early 1990s. To raise yields the Orange Cty treasurer bought bonds that had embedded short digital options on interest rates.
- A “one touch” option is an American digital. It is called a one touch because it is optimal to exercise as soon as the underlying price hits (“touches”) the strike price (do you know why?)

Knock-Options

- Knock options come in several varieties.
- Knock-in options. These are options with an underlying option that comes into existence only if some condition is met. There are up-and-in and down-and-in varieties. For example, an up-and-in call with a strike price of \$75 and a “knock barrier” of \$100 has a payoff at expiration if-and-only-if the underlying price hit or exceeded \$100 some time prior to expiration.
- Knock-out options. These are options that go out of existence if some condition is met. For example, a down-and-out call struck at \$75 with a knock barrier of \$50 expires worthless if the underlying price hits or is less than \$50 at any time during the option’s life. Thus, even if the stock price is \$100 at expiration, the option is worthless if the stock price hit \$50 prior to expiration.
- Knock-options are “path dependent” because their payoff depends not only on the value of the underlying at expiration, but the path that the underlying price follows prior to expiration.
- Manipulation can be an issue with these “barrier options.”

Asian Options

- An Asian option has a payoff that depends on the average price of the underlying over some time period.
- For instance, an Asian call on crude oil with a monthly averaging period has a payoff that depends on the average price of crude oil over a one-month-long period.
- There are average price options where the strike price is set prior to expiration, and the payoff is based on the difference between an average price and the pre-agreed strike price.
- There are also “average strike” options where the strike price is the average price on the underlying during some averaging period, and the payoff is based on the difference between the value of the underlying at expiration and this strike price.
- There are no closed form solutions for Asian options with arithmetic averaging. Various numerical techniques can be used to price them.

Some Other Exotics

- Compound options. These are options on an option, such as a put on a call, or a call on a call. There are formulae for valuing such options given Black-Scholes assumptions.
- Exchange options. These give you the right, but not the obligation to exchange one asset for another. There are formulae for valuing such options given Black-Scholes assumptions.
- “COD” (cash-on-delivery) options. Here you only pay the premium if they wind up in the money. These options can have a negative payoff.
- Quantos. Cross-currency options are the most common example. An example is a contract on the Nikkei Index with the payoff converted to USD. This conversion can occur at either an exchange rate set when the option is written or the exchange rate prevailing at expiration. The payoff to the option is $X \max[S-K, 0]$ where X is the exchange rate, S is the value of the Nikkei, and K is the strike price (in JY). Here there are two sources of risk—the Nikkei index (in yen), and the JY-USD exchange rate.
- Compos. These are options with payoffs $\max[XS-K, 0]$ where the strike price is in the domestic currency.

- Lookbacks. Lookbacks are path dependent options because the payoff depends on the maximum (or minimum) price reached by the underlying during the option's life. For instance, a lookback call with a strike price K has a payoff $\max[\max(S)-K, 0]$ where $\max(S)$ is the maximum price achieved by the underlying during the life of the option. Obviously this lookback is more valuable than a vanilla call.

Issues With Exotics

- Exotics frequently present serious hedging difficulties, especially when they are near the money and close to expiration.
- For example, a digital option's gamma is positive when it is out of the money and negative when it is the money. Right before expiration, the option's gamma is positive infinity at a price immediately below the strike price and negative infinity at a price immediately above the strike price. This behavior makes it hard to hedge.
- As another example, a knock-out option has an infinite gamma when the underlying price is at the knock-out boundary.
- Compound options have high gammas when they are at the money because they are options on options—the convexity in the underlying option compounds the convexity in the “top” option.
- Recall that convexity also has ramifications for the sensitivity option value to volatility. Hence, some exotics may have substantial vol (vega) risk.

Get Real

- Heretofore we have discussed options that are financial claims.
- The world is full of non-financial—“real”—options.
- In particular, virtually all business investment and operational decisions involve choice—i.e., optionality.
- Option to defer investment.
- “Time-to-build” option—build a project in stages with the choice to abandon after each stage.
- Option to increase or reduce investment scale.
- Option to abandon.
- Option to switch inputs or outputs (operational flexibility).
- Growth options (R&D, leases on property or resources).
- Combinations of the above options.

Valuing Real Options

- Traditional capital budgeting approaches ignore real options. This is unfortunate as real options may have a large affect on the profitability of investment.
- Firms may underinvest, overinvest, or invest at the wrong time if they ignore real options.
- Increasingly firms are using option pricing techniques to evaluate investment projects.
- Given the intense informational demands, it is often not practical to use these techniques to get precise estimates of real option value.
- Nonetheless, the use of options valuation techniques forces firms to analyze the characteristics of their investments more rigorously.

An Example

- Consider a firm that has the option to defer investment in drilling a gas well. It can invest today, spend \$104, and get a well worth \$100 (i.e., which generates \$100 in net present value). From an NPV perspective, this project is a turkey.
- However, the investment is risky. In one year, if good news about the value of gas arrives, the well is worth \$180. If bad news arrives, it is worth \$60. The interest rate is 8 percent per year.
- In terms of our binomial model, $u=1.8$ and $d=.6$. The “pseudo-probability” is $(1.08-.6)/(1.8-.6)=.4$.
- The option to delay the project one year is therefore $[.4(180-104)+.6(0)]/1.08=28.15$
- Exercise: Show that the value to delay the project for up to two years is worth \$31.82, and the value to delay the project for up to three years is worth \$45.01.

Intuition

- The intuition here is pretty straightforward—the option to delay allows you to wait for the arrival of information, namely whether gas prices are going to rise or fall. Since better information means better investment decisions, you the option to delay is valuable.
- The simple example here implies that you will never invest if you always have the option to delay. This is an American call option. Under the assumptions of the analysis (with no intermediate cash flows) standard analysis implies that deferral is always the best choice.
- In the real world, investments throw off benefits only if the firm invests in the asset. These benefits are like dividends and can induce the firm to exercise the call option early.
- In commodity markets, for instance, prices tend to be “mean reverting”—that is, if the price is higher (lower) than average today, it is expected to fall (rise). This mean reversion can make it optimal to invest now rather than waiting.

Implications

- Real options have important implications.
- Under some circumstances, the value of investment options is increasing in uncertainty (i.e., the volatility of the value of the asset).
- An increase in uncertainty induces firms to defer investment.
- Policy uncertainty (e.g., uncertainty in tax, regulatory or monetary policy) can affect firm investment strategies. One would expect to see greater levels of investment in “stable” jurisdictions (countries, states) than “unstable” ones

