

Financial Derivatives

Introduction

In this class we study the pricing, hedging, and uses of *financial derivatives* or derivatives for short. A *derivative* is a financial instrument whose value depends on the value of another asset, such as a stock, a foreign currency, a futures contract, or another quantity, such as volatility. These instruments have become increasingly important in modern financial markets, serving crucial functions in risk management, price discovery, and market efficiency. Understanding their mechanics and valuation principles is essential for anyone working in finance, whether in investment banking, asset management, or corporate treasury departments.

Futures and forward contracts are derivatives that allow traders to fix the price at which an asset will trade at a given date. In this sense, futures or forward contracts give the holder the *obligation* to buy or sell at a specific price, unlike options, which provide the holder with the *right* but not the obligation to buy or sell at a particular price. This fundamental distinction affects how these instruments are priced, traded, and incorporated into investment strategies. Futures are standardized contracts traded on organized exchanges with daily settlement, while forwards are customized agreements typically traded over-the-counter between two counterparties, with settlement occurring only at maturity.

Derivatives Contracts

As for any other financial instrument, the value of a derivative is the present value of its expected payoff. A positive payoff means receiving money, whereas a negative payoff represents an outflow of cash. The valuation process requires careful consideration of multiple factors including the underlying asset's price dynamics, interest rates, time to maturity, and in some cases, assumptions about market volatility. Different pricing

models have been developed to address various types of derivatives, with the Black-Scholes-Merton model being perhaps the most famous for options pricing.

The *payoff* represents how much money you get if you buy the instrument. The *profit*, on the other hand, depends on how much you pay for the derivative. The payoff is realized at maturity for many derivatives, although this is only sometimes true.¹ Understanding the relationship between payoff and profit is critical for effective derivatives trading and risk management. Many derivatives strategies involve complex combinations of instruments with different payoff structures, designed to capitalize on specific market views or to hedge particular risk exposures. Traders must carefully analyze both the potential payoff scenarios and the initial cost to determine the risk-reward profile of any derivatives position.




Figure 1: The timeline of the cash flows.

If S_T denotes the value of a stock at maturity, the payoff of a derivative written on the stock will be a function of S_T that we denote by $f(S_T)$. An important question that we answer in this class is how to *price* this type of derivative. This fundamental challenge lies at the heart of financial mathematics and requires sophisticated mathematical techniques drawing from stochastic calculus, martingale theory, and partial differential equations. The pricing methodology we develop will provide a framework that can be applied across various derivative instruments with different payoff structures.

If we denote by ξ the relevant continuously compounded discount rate for the derivative, the value of the derivative is

$$V = e^{-\xi T} \mathbb{E} [f(S_T)]. \quad (1)$$

Even though this expression is correct, we do not generally know the right value for the discount rate ξ . For many derivatives, the only way to know ξ would be to know the value

¹American-type options can realize the payoff anytime before or at maturity.

of the derivative first.² This circular problem presents a significant obstacle in derivatives pricing under the physical measure. Furthermore, the discount rate may vary with market conditions, investor risk preferences, and the specific characteristics of the derivative instrument, making it challenging to determine ex-ante without additional assumptions about investor behavior and market equilibrium.

The seminal work of Black and Scholes (1973) and Merton (1973) showed that it is possible to price derivatives by *replication*. Continuously trading in the underlying asset and a risk-free bond can generate the same payoffs as a derivative written on the asset. If that is the case, the value of the replicating portfolio today must be the value of the derivative asset to prevent arbitrage opportunities. This groundbreaking insight shifted the focus from subjective valuation based on investor preferences to objective valuation based on arbitrage considerations. The concept of replication became the cornerstone of modern derivatives pricing theory and opened up new avenues for financial innovation and risk management techniques.

Their analysis also showed that replication works for *any value* of the discount rate used to price the underlying asset, that is, for any *probability measure* used to assign events of the underlying asset. A crucial implication of this observation is that replication should also work in a world populated by *risk-neutral* individuals. This powerful insight allows us to sidestep the complexities associated with modeling investor risk preferences and instead focus on a simplified artificial economy where all investors are indifferent to risk. In such a world, all assets would be expected to earn the risk-free rate of return, greatly simplifying the valuation process.

Under the risk-neutral probability measure, the value of the derivative is

$$V = e^{-rT} E^* [f(S_T)] \quad (2)$$

where r denotes the risk-free rate, and the expectation is taken under the risk-neutral probability measure. This elegant formulation transforms the pricing problem into a mathematical expectation under an artificial probability measure, often referred to as the equivalent martingale measure. The existence and uniqueness of this measure is closely

²For options, for example, the relevant discount rate under the historical or physical measure depends on the moneyness and the option's maturity.

tied to fundamental concepts in financial economics such as market completeness and the absence of arbitrage opportunities.

A crucial difference between (1) and (2) is the discount rate used to price the derivative. In (1), we need to know both the discount rate of the underlying asset to compute the expectation and the discount rate of the derivative to compute its price. In (2), on the other hand, we can discount the expected payoffs of any asset at the risk-free rate. This remarkable simplification is one of the most important results in financial economics and has profound implications for how we approach derivative pricing. It essentially decouples the valuation problem from the need to estimate risk premiums, allowing us to focus instead on modeling the dynamics of the underlying asset under the risk-neutral measure.

From equation (2), we can see that to price a derivative, it is essential to understand what is a good proxy for the *risk-free rate* and also how to compute the expected payoff under the risk-neutral measure. In practice, government securities such as Treasury bills or bonds are often used as proxies for the risk-free rate, although the appropriate choice may depend on the time horizon of the derivative and the currency in which it is denominated. Computing the expected payoff requires specifying a stochastic process for the underlying asset under the risk-neutral measure, which typically involves assumptions about volatility and other parameters.

The payoff function is linear for some derivatives, such as forward contracts, simplifying the expected payoff computation. For other derivatives, such as options, the payoff function is nonlinear and generally more complex to price than linear payoffs. These nonlinear payoffs create convexity effects that are central to understanding options behavior and pricing. The mathematical tools required to handle these nonlinearities include numerical methods such as Monte Carlo simulation, binomial trees, and finite difference schemes for solving the associated partial differential equations. Each approach has its strengths and limitations, making the choice of methodology an important consideration in practical applications.

Example 1 (A derivative with a linear payoff). A *forward contract* is a commitment to purchase or sell an asset at maturity for K . The payoff of a *long forward* is the difference

between the price of the asset at maturity and the price agreed in the contract; that is, the payoff is a linear function of the stock price:

$$f(S) = S - K$$

Because the payoff can be positive or negative depending on the sign of S , the value of the contract can be positive or negative. Usually, the contract has zero value at inception. Later on, the value of the contract will change and might become positive or negative. $\square$

Example 2 (A derivative with a nonlinear payoff). An *option* gives the holder the right but not the obligation to purchase or sell an asset at maturity for a given price of K . The payoff of an option is a nonlinear function of the asset price at maturity.

For example, the buyer of a call option receives:

$$f(S) = \begin{cases} 0 & \text{if } S < K \\ S - K & \text{if } S \geq K \end{cases}$$

Since the payoff is non-negative, the holder of an option must pay a premium to the seller. $\square$

More Complex Derivatives

The financial engineering revolution in the 1990s revolved around the idea that we could package simpler derivatives and build new products. Pricing many more complex derivatives involves knowing how to price the basic building blocks used to make the financial product.

Derivatives with Periodic Payments

A classical way to create a more complex and helpful derivative is to put derivatives with different expirations together, creating a product that involves the payment of cash flows periodically over time.

For example, *interest rate swaps* involve the exchange of a fixed interest rate for a floating interest rate or vice-versa. They allow corporations to convert a loan with floating payments into a bond with fixed cash flows. Another example is a *credit default swap* (CDS), which involves the exchange of periodic payments in exchange for protection in case of a bond default.

Pricing derivatives with periodic payments is similar to pricing a derivative with a single payment. If we denote by $p(\cdot)$ the pricing functional, we have that:

$$p(f(S_{t_1}) + f(S_{t_2}) + \dots + f(S_{t_n})) = p(f(S_{t_1})) + p(f(S_{t_2})) + \dots + p(f(S_{t_n})).$$

The same logic applies to the pricing of bonds. The price of a bond must be the sum of the present value of its coupons and its face value; otherwise, there would be a straight-forward arbitrage opportunity. In asset pricing, we call this property the *law of one price*, which says pricing must be linear to prevent arbitrage opportunities.

Assets with Embedded Derivatives

Many assets, such as corporate bonds, have *embedded* options. For example, many bonds in financial markets are *callable*; the issuer has the right to pay the bondholder the principal before maturity. To value a callable bond involves modeling the evolution of the term structure of interest rates and analyzing which states it is profitable to call the bond.

For many callable bonds, the issuer has the right to call a bond at any time, starting on the first date the bond is callable until its maturity. Therefore, the call option embedded in these bonds is an American-type option. The pricing of these bonds requires using

numerical techniques such as binomial trees since there is no close-form solution for the price of an American-type option.

Other bonds are convertible into shares of the issuing company at a fixed price. Thus, *convertible* bonds contain a call option on the company stock, which might be valuable.

Do We Need Other Payoffs?

Theoretically, we could design a derivative with any payoff function $f(S)$ that one might think of. For example, we could choose $f(S) = S^2$ or $f(S) = \ln(S)$.

With forwards and options, it is possible to build any payoff a trader might want. We will see that having options and forwards with different strikes can *complete the market*. Combining options and forwards is usually called *options strategies*.

Derivatives Markets

Uses of Derivatives

Derivatives allow investors to obtain payoffs that are useful to achieve specific objectives. For example, some commodity producers use derivatives to *hedge* their future production by fixing the price they will sell today. Other traders like derivatives because they can obtain custom design payoffs, allowing them to *speculate* in specific ways.

Therefore, derivatives make both types of traders, hedgers and speculators, better off by expanding their trading opportunity set and thus increasing their utility.

The interaction between hedgers and speculators has been widely studied in commodity markets. Keynes and Hicks (Hicks 1937) see speculators primarily as risk-takers who provide insurance to risk-averse hedgers, who pay premiums to speculators. Working (1953), in contrast, sees commodity producers as strategic agents who want to maximize profits, making sophisticated trading decisions rather than just automatically hedging everything. Telser (1958) supports and extends Working's view with additional theoretical and empirical analysis. He shows that hedging costs affect where hedgers choose to

hedge (market selection) and demonstrates that hedgers will often accept “poor” hedges if they are cheaper.

The Market for Derivatives

Derivatives are always in zero-net supply. Every long position has a corresponding short position. As such, the demand for derivatives can be positive or negative depending on the traders’ intentions generating the demand.

The demand for derivatives comes from buy-side traders who want to use derivatives for hedging or speculative purposes. For example, many hedgers in commodity markets are commodity producers who want to fix the price at which they can sell their products. Hedging their exposure using futures involves selling these contracts to whoever is willing to take the opposite side of the trade.

The *net demand*, positive or negative, is balanced by sell-side investors or *market makers* that provide liquidity to the rest of the market.

Pricing and Hedging of Derivatives

Market makers often will *hedge* their exposure by dynamically trading the underlying asset and risk-free bonds. One of the main results of modern asset pricing is that a perfectly hedged portfolio should earn a risk-free interest rate. Otherwise, there would be an *arbitrage opportunity*. This is the main intuition behind the Black-Scholes model for option pricing. Therefore, we must first learn how to hedge or replicate the position to price an option or a forward contract.

For certain derivatives, such as options, the hedging recipe depends heavily on modeling the stock price evolution over time. In modeling the evolution of the underlying asset, time can be discrete or continuous. The choice of how to model time usually depends on how difficult it is to solve the model.

The distribution of random shocks will affect the evolution of stock prices over time. Some models, such as the geometric Brownian motion, assume that the instantaneous rate of

return follows a normal distribution. It is possible to introduce more complex elements, such as stochastic volatility or jumps, to make the modeling closer to actual financial markets.

When perfect replication is not feasible, the pricing of derivatives might depend on supply and demand forces. For example, Garleanu, Pedersen, and Poteshman (2008) show how in an incomplete market the net-demand for options can influence their price. The intuition for this result is simple. If risk-averse market-makers cannot hedge all the risk of the opposite trade they just engaged with a customer, they will ask for a risk compensation that will grow larger as the net-demand for the derivative increases.

Practice Problems

For solutions go to lorenzonaranjo.com/fin451.

Problem 1. Which of the statements below is (are) correct? Select all alternatives that apply.

- a. A financial derivative always pays a positive cash flow to its holder.
- b. The price of a financial derivative cannot be negative, i.e., it should always cost something.
- c. Arbitrage opportunities for derivatives should not exist.
- d. A derivative is an instrument whose value depends on the value of another asset.

Problem 2. Which of the following financial assets have embedded options? Select all alternatives that apply.

- a. A callable bond.
- b. A forward contract.
- c. A convertible bond.
- d. A company's stock.

References

- Black, Fischer, and Myron Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81 (3): 637–54.
- Garleanu, Nicolae, Lasse Heje Pedersen, and Allen M. Poteshman. 2008. "Demand-Based Option Pricing." *Review of Financial Studies* 22 (10): 4259–99.
- Hicks, John R. 1937. "Mr. Keynes and the 'Classics'; A Suggested Interpretation." *Econometrica*, 147–59.
- Merton, Robert C. 1973. "Theory of Rational Option Pricing." *The Bell Journal of Economics and Management Science*, 141–83.
- Telser, Lester G. 1958. "Futures Trading and the Storage of Cotton and Wheat." *Journal of Political Economy* 66 (3): 233–55.
- Working, Holbrook. 1953. "Futures Trading and Hedging." *American Economic Review* 43 (3): 314–43.