Stochastic Processes

Olivier Scaillet

University of Geneva and Swiss Finance Institute

Outline

Stochastic process

Stochastic processes are used to model the state of a time dependent random system.

We work with a probability space $(\Omega, \mathcal{A}, P)$ and a function $(t, \omega) \to X(t, \omega)$ of $\mathbb{R}^+ \times \Omega$ in a space E giving the system state.

For ω fixed, that is, a particular evolution of the system, the successive states are given by the function $t \to X(t, \omega)$ that we call trajectory or path by analogy with a system modeling the position of a particle.

Definition: stochastic process

A stochastic process $X = \{X_t, t \in T\}$ is a familly of rv defined on the same probability space and indexed by the time t taking its values in a set T of indices.

- If $T = \{0, 1, 2, 3, \dots\}$ or $T = \{\dots, -2, -1, 0, 1, 2, 3, \dots\}$, the process is said to be a discrete-time process.
- If the set $T = \mathbb{R}^+$ or $T = \mathbb{R}$, it is said to be a continuous-time process.

Example: simple random walk

Let $X_1, X_2, \ldots$ be a family of independent rv, each taking value 1 with probability p and value –1 with probability $1 - p$.

Let the sum process $S_t = S_0 + \sum_{i=1}^{t} X_i$. The series $S = \{S_t, t \in \mathbb{N}\}$ is a random walk starting at S_0.

Say that S_t stands for the wealth of a gambler after t rounds, earning \$1 if she wins and paying \$1 if she loses, and starting with initial wealth $\$S_0$.

We can consider the case $P[S_t \leq 0]$ (ruin probability), or $\min\{t \in T : S_t \leq 0\}$ (stopping time), for instance.

Example: simple random walk (cont'd)

We can rewrite the random walk by $S_t = S_{t-1} + X_t$. The particle being at the position S_{t-1} at time $t-1$ can either go left 1 unit with probability p or go right 1 unit with probability $1-p$.

Properties:

1. Space homogeneity:
   $P[S_t = j | S_0 = a] = P[S_t = j + b | S_0 = a + b]$
2. Time homogeneity: $P[S_t = j | S_0 = a] = P[S_{t+h} = j | S_h = a]$
3. Markov property:
   $P[S_{t+1} = j | S_0 = s_0, \ldots, S_t = s_t] = P[S_{t+1} = j | S_t = s_t], \forall t$
   (conditionally on present, future does not depend on past)

The process is said to be memory-less. The random walk is a particular case of Markov chains.

Definition: Markov chain

We consider a discrete-time process X taking values in a countable space called state space.

The process X is a Markov chain if it satisfies the Markov condition
$P[X_t = x | X_0 = x_0, X_1 = x_1, \dots, X_{t-1} = x_{t-1}] = P[X_t = x | X_{t-1} = x_{t-1}]$
for any t.

This Markov property is equivalent to
$P[X_{t+1} = x | X_{t_1} = x_{t_1}, X_{t_2} = x_{t_2}, \dots, X_{t_k} = x_{t_k}] = P[X_{t+1} = x | X_{t_k} = x_{t_k}]$
with $0 \leq t_1 \leq \cdots \leq t_k \leq t$, or also
$P[X_{t+h} = x | X_0 = x_0, X_1 = x_1, \dots, X_t = x_t] = P[X_{t+h} = x | X_t = x_t]$
for any t, h.

Definition: Markov chain (cont'd)

We can generalize this concept to order r Markov chains:

$$P[X_t = x | X_0 = x_0, X_1 = x_1, \ldots, X_{t-2} = x_{t-2}, X_{t-1} = x_{t-1}]$$
$$= P[X_t = x | X_{t-1} = x_{t-1}, X_{t-2} = x_{t-2}, \ldots, X_{t-r} = x_{t-r}]$$

As the state space is countable, we can map it to the set of integers. We say that when $X_t = i$, the chain is in state i at time t, or also visits i or takes value i.

The chain evolution depends on the probabilities $P[X_t = j | X_{t-1} = i]$ called transition probabilities. In general, they depend on t, i, and j. However, we often consider a simple case where they depend on i and j, but are independent on t.

Homogeneity

We consider a simple case where transition probabilities depend on i and j, but are independent on t. A Markov chain X is said homogenous if $P[X_{t+1} = j | X_t = i] = P[X_1 = j | X_0 = i]$, $\forall t, i, j$. It is a stability condition on the probabilities in time.

The matrix $P = (p_{ij})$ with $p_{ij} = P[X_{t+1} = j | X_t = i]$ is the transition matrix of the chain. The transition matrix P is a stochastic matrix as it satisfies:

1. all entries of P are non negative: $p_{ij} \geq 0$, $\forall i, j$.
2. for each row of P, the sum of the entries gives 1: $\sum_j p_{ij} = 1$, $\forall i$.

Example: simple random walk

We move right 1 unit (go from i to $i+1$) with probability p and left 1 unit (go from i to $i-1$) with probability $1-p$.

The state space is $\{0, \pm 1, \pm 2, \dots\}$ and the transition probabilities are
$$p_{ij} = P[S_t = j | S_{t-1} = i] = \begin{cases} p & \textit{if } j = i+1, \\ 1-p & \textit{if } j = i-1, . \\ 0 & \textit{else}. \end{cases}$$

Transitions

We consider the evolution of the chain in h steps (horizon h). The h steps transition matrix $P(t, t+h)$ is the matrix whose entries give the transition probabilities in h steps:
$p_{ij}(t, t+h) = P[X_{t+h} = j | X_t = i]$.

The homogeneity hypothesis is equivalent to the condition $P(t, t+1) = P$ (the matrix is not time dependent).

Using Chapman-Kolmogorov equations, we get:
$p_{ij}(t, t+h+l) = \sum_k p_{ik}(t, t+h) p_{kj}(t+h, t+h+l)$, or using matrix notation, $P(t, t+h+l) = P(t, t+h)P(t+h, t+h+l)$.

Transitions (cont'd)

In the homogenous case, we have $P(t, t+2) = P(t, t+1+1) = P(t, t+1)P(t+1, t+1+1) = PP = P^2$. By recurrence $P(t, t+h) = P^h$ and $P(t, t+h+l) = P^h P^l = P^{h+l}$.

In particular, $P(t, t+h) = P(0, 0+h) = P^h$.

The evolution of the chain depends on the initial state at time 0 and the transition matrix P. Let $\mu_i^{(t)} = P[X_t = i]$ and the corresponding row vector $\mu^{(t)} = (\mu_i^{(t)})$.

We can show $\mu^{(t)} = \mu^{(0)} P^t$, that is, the random evolution of the chain is characterized by the transition matrix P and the initial probability mass function of $\mu^{(0)}$.

State classification

Let us consider the chain evolution as a particule movement between multiple states. We consider the time (possibly infinite) that the particle needs to come back to its starting point.

A state i is said recurrent (or persistent) if $P[X_t = i | X_0 = i] = p_{ii}(t) = 1$ for some $t \geq 1$.

A state i is said transitory (or transient) if this probability is strictly smaller than 1.

First passage time

We can also consider the probabilities of first passage $f_{ij}(t) = P[X_1 \neq j, X_2 \neq j, \ldots, X_{t-1} \neq j, X_t = j | X_0 = i]$, that is, the probability that the first visit of state j needs exactly t steps starting at i.

Let us define the probability that the chain visits state j at some point in time, starting at state i, by $f_{ij} = \sum_{t=1}^{\infty} f_{ij}(t)$.

A state j is recurrent if and only if $f_{jj} = 1$.

Let $T_j = \min\{t \geq 1 : X_t = j\}$ be the time of first passage of state j, with $T_j = \infty$ if the state is never visited.

The state j is transitory if and only if $P[T_i = \infty | X_0 = i] > 0$.

Mean time of recurrence

We also compute the mean time of recurrence m_i of state i defined by $m_i = E[T_i|X_0 = i] = \begin{cases} \sum_t tf_{ii}(t) & \text{if } i \text{ recurrent}, \\ \infty & \text{if } i \text{ transitory}. \end{cases}$

m_i can nonetheless be infinite for a recurrent state i and we say that:

- i is null recurrent if $m_i = \infty$
- i is non-null recurrent (or positive recurrent) if $m_i < \infty$

Periodicity

The period $d(i)$ of a state i is defined as the largest common divider of all values t such that $p_{ii}(t) > 0$.

The state is said to be

- periodic if $d(i) > 1$
- aperiodic if $d(i) = 1$

We thus have that $p_{ii}(t) = 0$ unless t is a multiple of $d(i)$.

A state is said to be ergodic if it is recurrent, non-null, and aperiodic.

Chain classification

We define the class of a chain by looking at the links of its states:

- State j is accessible from state i if there exists some t such that $p_{ij}(t) > 0$
- State i communicates with state j, or $i \rightarrow j$, if j is accessible from i
- States i and j intercommunicate, or $i \leftrightarrow j$, if i communicates with j and j communicates with i.

If $i \leftrightarrow j$, then

1. i and j have the same period
2. i is transitory if and only if j is transitory
3. i is null recurrent if and only if j is null recurrent.

Chain classification (cont'd)

A set C of states is said

- closed if $p_{ij} = 0$ for all $i \in C$, $j \notin C$
- irreducible if $i \leftrightarrow j$ for all $i, j \in C$

If the chain reaches a state that belongs to a closed set, it never leaves this set.

A closed set containing exactly one element (singleton) is said to be absorbing.

Stationary distribution

What is the behavior of a Markov chain when we wait a large number of periods? We are interested in the distribution of X_t when t is large. The existence of a limit distribution for X_t as $t \to \infty$ is strongly linked to the existence of a stationary distribution.

The vector π is called stationary distribution of the chain if π has entries (π_j) such that

1. $\pi_j \geq 0$ for all j, and $\sum_j \pi_j = 1$
2. $\pi = \pi P$, that is, $\pi_j = \sum_i \pi_i p_{ij}$ for all j

Stationary distribution (cont'd)

By iterating, we get that $\pi = \pi P^t$, for any t. Additionally, we have seen that homogenous chains have the property $\mu^{(t)} = \mu^{(0)} P^t$ with $\mu^{(t)} = (\mu_i^{(t)})$ defined by $\mu_i^{(t)} = P[X_t = i]$.

Therefore, if property $\mu^{(0)} = \pi$, that is, X_0 follows the distribution given by π, then X_t also follows distribution π.

The distribution of X_t is stationary as time goes, which implies that π is also the limit distribution of X_t as $t \to \infty$.

Stationary distribution (cont'd)

An irreducible chain has a stationary distribution π if and only if all of its states are non-null recurrents.

In this case, π is the unique stationary distribution and is given by $\pi_i = m_i^{-1}$ with m_i the mean time of recurrence of i.
For irreducible and aperiodic chains, we have that $p_{ij}(t) \to m_j^{-1}$, $\forall i, j$ as $t \to \infty$.

Application in insurance, reinsurance chain

Markov chains have applications in insurance:

- bonus-malus system
- reassuring chain

Bonus-malus system

Use of Markov chains to model a problem of car insurance: bonus-malus system.

We set a finite number L of classes that define the premium paid by the insured.

Each year, the class of the insured is updated depending on the class of the previous year and the number of sinisters (or claims) during the year.

If no sinister is recorded, the insured is moved to a class whose premium is lower; conversely, if sinisters have been reported, she is moved to a class whose premium is higher.

Model: bonus-malus system

The parameters of the model are the following:

- L classes indexed by 1 to L. Classe 1 is called superbonus and class L is called supermalus. The annual premium depends on the number of insured people in each class and is computed using a fee scale.
- a fee (or premium) scale $b = (b_1, \ldots, b_L)$ with $b_1 \leq b_2 \leq \cdots \leq b_L$
- a transition rule that defines the transfer from one class to the other as a function of the number of sinisters. If k sinisters are reported $t_{ij}(k) = 1$ if the class changes from i to j, and $t_{ij}(k) = 0$ otherwise
- an initial class i_0 for any new customer.

Let $T(k) = (t_{ij}(k))$, each $T(k)$ is a $0 - 1$ matrix for which each row has exactly one 1.

Model: bonus-malus system (cont'd)

Suppose that the number of sinisters per year forms a sequence $Y_1, Y_2, \ldots$ of i.i.d. rv and with mass function $P[Y = k] = q_k$.

We denote by $X_1, X_2, \ldots$ the successive classes through which the insured has gone, and suppose that the class of a particular year only depends on the class of the previous year and the number of sinisters during the year. We can describe the evolution of $\{X_t\}$ using the recurrence equation $X_t = \phi(X_{t-1}, Y_t)$, with $\phi(i, k) = j$ if and only if $t_{ij}(k) = 1$.

Model: bonus-malus system (cont'd)

This evolution corresponds to a Markov chain whose probabilities of transition are given by $p_{ij} = \sum_{k=0}^{\infty} q_k t_{ij}(k)$.

In practice, we suppose that the number of sinisters follows a Poisson law with parameter λ, which leads to $p_{ij} = \sum_{k=0}^{\infty} \frac{\lambda^k}{k!} e^{-\lambda} t_{ij}(k)$.

Model: bonus-malus system (cont'd)

In a bonus-malus system, the fee b_i is applied every time an insured visits class i. He may then consider

- the cumulative sum of the premia paid until t: $R_t = \sum_{k=0}^{t-1} b_{X_k}$
- the discounted cumulative sum of the premia paid until t: $R_t^a = \sum_{k=0}^{t-1} \frac{b_{X_k}}{(1+r)^k}$, with interest rate r
- the discounted cumulative sum of the premia paid with an infinite horizon: $R_\infty^a = \sum_{k=0}^{\infty} \frac{b_{x_k}}{(1+r)^k}$.

Model: bonus-malus system (cont'd)

We may show that when the chain has an initial class $X_0 = i_0$, we have

- $E[R_t] = e_{i_0} \sum_{k=0}^{t-1} P^k b'$
- $E[R_t^a] = e_{i_0}(Id - (1+r)^{-1}P)^{-1}(Id - (1+r)^{-n}P^n)b'$
- $E[R_\infty^a] = e_{i_0}(Id - (1+r)^{-1}P)^{-1}b'$,

where e_{i_0} denotes the row vector with zeros everywhere except on the i_0th entry which which is equal to 1.

We may also show that if the chain has a stationary distribution π, $E[R_t^a]$ is a linear combination of the stationary fee, i.e. $\bar{b} = \pi b'$.

Chain of reinsurance

Suppose that an asset such as a plane needs insurance. Its value is too high for a unique company to bear the risk.

We start at time 0 with an insurer of first rank. Then, the insurer negotiates with second rank insurers in order to sell them part of the original insurance. Each on the second rank companies does the same and negotiates part of its own insurance policy share with a third rank insurer. This sequence corresponds to a (simplified) chain of reinsurance.

Model: chain of reinsurance

Suppose that each company needs a certain time interval (e.g. administrative tasks) with distribution function F before signing simultaneously all its contrats of reinsurance.

The number of reinsurance companies is assumed to be randomly distributed with some probability mass function.

Each company initiates its chain of searches, independently on the history of other companies. All administrative durations are i.i.d. following distribution F.

Model: chain of reinsurance (cont'd)

The sequence $\{X_t\}$ counting the number of companies in each node of the chain at time t (tth rank) is called branching process of Galton-Watson-Bienaymé.

Let $R_{t,i}$ be the number of companies linked to company i in node t and suppose that all $R_{t,i}$ are i.i.d. The process $\{X_t\}$ defined by $X_1 = 1$, $X_{t+1} = \sum_{i=1}^{X_t} R_{t,i}$ is a Markov chain.

We may show that the average number of companies at link $t = n$, that is, $E[X_n]$, is equal to $E[X_n] = m^n$ where m is the average number of companies contacted at any link and by any company.

Application in finance

Markov chains have multiple applications in finance, e.g.:

- credit ratings
- binomial model for derivatives pricing

Change of credit rating

A change in the credit rating of an issuer relates to the increase or decrease of the probability of default. It has an impact on the assets (debts) it might issue.

We may go bankrupt, or have the credit quality increase or decrease. It is important to correctly estimate these changes in a context of quantitative risk management with risky assets.

If returns follow approximately a normal distribution, mean and variance are sufficient to identify risk-return arbitrages. Returns are no longer normally distributed when there is a risk of default, as a default event might wipe out the entire value of the assets.

Change of credit rating (cont'd)

The distribution is skewed because of the increased risk of large loss and has more extreme events than the normal distribution (kurtosis).

The standard deviation (or variance) is a symmetrical risk measure and is not adapted to this context. We prefer using quantiles of the portfolio return distribution (or the value distribution).

We estimate quantiles using simulations of the asset values across time, taking into account market conditions (interest rates, exchange rates, ...) and also changes in the credit risk (default, credit rating change)

Model: change of credit rating

We use a Markov chain whose states correspond to credit ratings.

We model the probabilities of transition between ratings, as well as the default probability, e.g. probability of going from rating AAA to rating AAB, or probability to go from rating BBB to bankruptcy. These parameters are estimated using historical data on the change of ratings and defaults of firms.

Model: change of credit rating(cont'd)

We then simulate:

1. the evolution of the Markov Chain to obtain the change of ratings of the portfolio assets
2. the evolution of the financial variables in order to keep track of the changes in financial markets

We build a histogram of the portfolio PnL with a maturity of 1 year, 2 years, ... We compute the associated portfolio quantiles.

We define the economic capital as the difference between an extreme loss quantile (99.9%) and the average loss. This value is a buffer used to absorb the impact of large unexpected losses (defaults).

Options valuation

An option is a financial asset giving the right (but not the obligation) to its holder to buy or sell some predefined amount of another asset, at some particular date and price.

Multiple factors define an option:

- option kind: *call* (option to buy), *put* (option to sell)
- underlying asset (stock, debt, interest rate, exchange rate, ...) whose value at time t is denoted by S_t)
- quantity
- maturity: expiration date (denoted by T)
- exercise price: price at which the exercise takes place (denoted by K)
- kind of exercise: European (at maturity), American (until maturity)
- premium: price of the option

Payoff of a European option

Consider a call option:

- if $S_T > K$, we exercise the option and the profit is $S_T - K$ as we pay a price K for an asset that has value S_T
- if $S_T \leq K$, we do not exercise the option and there is no profit, as we do not buy the underlying at a price higher than its market price.

Thus, the payoff (or value of the call option) at maturity is $\max(S_T - K, 0) = (S_T - K)_+$. Symmetrically, the payoff of a put option at maturity is given by $(K - S_T)_+$.

Pricing and hedging

The option seller faces two problems:

1. a pricing problem: determine the value today of the future random payoff
2. hedging problem: be sure to be able to honor her part of the contract at maturity

The basic assumption in derivatives pricing is absence of arbitrage (no free lunch), that is, we cannot make a profit (above riskfree rate) without investing an initial capital and taking risks.

Put-call parity

We can deduce the put-call parity relationship from the absence of arbitrage assumption. It takes the form $C_t - P_t = S_t - Ke^{-r(T-t)}$, i.e., the difference between the price of a call option and a put option (with same strike price and maturity) is equal to the difference of the underlying price and the discounted strike price.

We may also deduce price boundaries:

- $S_t > C_t > \max(S_t - Ke^{-r(T-t)}, 0)$
- $Ke^{-r(T-t)} > P_t > \max(Ke^{-r(T-t)} - S_t, 0)$

These results are only due to the absence of arbitrage hypothesis, in particular there is no assumption on the stochastic process followed by the assets.

European call pricing

Unfortunately, this assumption is not sufficient to compute a unique price, and boundaries are large.

We need to model the evolution of the underlying asset. The most famous model is Black-Scholes model (1973). It is a continuous-time model which can be defined as the limit of a discrete-time binomial model as the time interval between price changes goes to zero.

One-period binomial model

- One period (two dates, t et $t+1$),
- Two states: high ($S_t \rightarrow S_{t+1} = uS_t$) and low ($S_t \rightarrow S_{t+1} = dS_t$), where $d < u$ are multiplicative factors,
- Two financial assets: a stock (S_t), and a risk free asset ($1 \rightarrow 1 + r$ in any state).

We show that the probability of the two states does not have any impact on the option price.

One-period binomial model (cont'd)

Hypothesis: $dS_t < K < uS_t$. Indeed, if $K < dS_t < uS_t$, the buyer always exercises his option, thus nobody would like to sell such a contract (and inversely for $K > dS_t > uS_t$).

We want to determine the call price at time t.The payoff function at $t+1$ is $\max(S_{t+1} - K, 0)$, that is

- $\max(uS_t - K, 0) = uS_t - K$ in the up state
- $\max(dS_t - K, 0) = 0$ in the down state

Pricing by replication

In order to evaluate the option price, we build a portfolio with the underlying risky asset and the risk free asset that replicates the option payoff.

If there is no arbitrage between the option and such a portfolio, their prices should be the same as they both have the same payoff at maturity in any state of the world.

If they do not have the same price, it means that there is an arbitrage opportunity and it is profitable to sell the expensive asset and buy the cheap one. An immediate and risk free profit is made as the payoffs of this strategy offset each other in any state of the world.

Pricing by replication (cont'd)

We denote by α_1 the amount of stocks, α_2 the amount of risk free asset.

- At time t, the portfolio value is $\alpha_1 S_t + \alpha_2$
- At time $t+1$, the portfolio value is $\alpha_1 S_{t+1} + \alpha_2(1+r)$, that is, $\alpha_1 u S_t + \alpha_2(1+r)$ in the up state, and $\alpha_1 d S_t + \alpha_2(1+r)$ in the down state.

We determine the value of α_1 and α_2 such that the portfolio replicates the option payoff:

- $\alpha_1 u S_t + \alpha_2(1+r) = u S_t - K$
- $\alpha_1 d S_t + \alpha_2(1+r) = 0.$

Pricing by replication (cont'd)

We obtain a system of equations with two equations and two unknowns whose solution is $\alpha_1^* = \frac{uS_t - K}{uS_t - dS_t}$ and $\alpha_2^* = -\frac{(uS_t - K)dS_t}{(uS_t - dS_t)(1+r)}$.

The current option price must be equal to the portfolio price whose weights are given by α_1^*, α_2^* in order to avoid arbitrages. The option price is therefore given by
$C_t = \alpha_1^* S_t + \alpha_2^* = \frac{uS_t - K}{uS_t - dS_t} S_t - \frac{(uS_t - K)dS_t}{(uS_t - dS_t)(1+r)}$.

Pricing with risk-neutral expectation

Defining $\pi = \frac{(1+r)S_t - dS_t}{uS_t - dS_t} = \frac{(1+r)-d}{u-d}$, we may rewrite the option pricing formula as $C_t = \frac{1}{1+r}[\pi(uS_t - K) + (1-\pi)0]$.

This expression corresponds to the expectation of a binomial law taking value $\frac{uS_t - K}{1+r}$ with probability π and value $\frac{0}{1+r}$ with probability $1-\pi$. The option price is thus an expectation of the discounted future payoffs, where the probabilities have been shifted.

We call π risk-neutral probability of the up state, in contrast to the true or historical probability p. The probability p is not involved in the pricing equation and is therefore irrelevant in this context!

Risk-neutral probability

Let $(p, 1-p)$ be the historical probabilities of an up and down move of the stock. As there are only two possible states, S_{t+1} follows a binomial law that takes value uS_t with probability p and value dS_t with probability $1-p$.

The expectation is $E^p[S_{t+1}] = puS_t + (1-p)dS_t$ and the expected return is given by $\frac{E^p[S_{t+1}]}{S_t} = pu + (1-p)d$.

By adding $0 = (1+r) - [p(1+r) + (1-p)(1+r)]$ to this expression, we get:

$$\begin{aligned}\frac{E^p[S_{t+1}]}{S_t} &= (1+r) - [p(1+r) + (1-p)(1+r)] + pu + (1-p)d \\ &= (1+r) + [u-(1+r)]p + [d-(1+r)p](1-p)\end{aligned}$$

Risk-neutral probability (cont'd)

Therefore, the term
$\frac{E^p[S_{t+1}]}{S_t} - (1+r) = [u-(1+r)]p + [d-(1+r)p](1-p)$
corresponds to a risk premium, that is, an additional return paid as a compensation for risk.

Doing the same calculations with the risk-neutral probability π rather than p, we get $E^\pi[S_{t+1}] = \pi u S_t + (1-\pi) d S_t = (1+r)S_t$.

The expected return is then $\frac{E^\pi[S_{t+1}]}{S_t} = 1 + r$, i.e., computing the expected return with probability π gives a risk free return.

The risk premium is thus zero: $\frac{E^\pi[S_{t+1}]}{S_t} - (1+r) = 0$.

Risk-neutral probability (cont'd)

Using the risk-neutral probability π rather than the historical probability p neutralizes the risk premium, i.e., investors ask for the same return as truly risk-neutral agents. The general result follows:

In a one-period model with two states of the world, there exists a unique probability π called risk-neutral probability, such that for each asset $x_{t+1} \in \mathbb{R}^2$ its price and at time t is equal to $P_t = \frac{1}{1+r} E^\pi[x_{t+1}]$, that is, the discounted risk-neutral expectation of the future random payoffs.

Risk-neutral probability (cont'd)

Note that there is no notion of utility function with this approach.

This approach uses the absence of arbitrage hypothesis to value assets, which does not require notions of equilibrium and modeling of economic agents with utility functions.

We make the assumption of a perfect market:

1. we can buy or sell an asset at the same price (no transaction fees on the risk free asset)
2. we can borrow or lend at the same interest rate (no fee on the risk free asset)

In this model, we only have one period (or two dates), which is a very short Markov chain! We generalize this result to a multi-period binomial model.

Multi period binomial model

- Multiple periods: n periods, or $n + 1$ dates
- Two states of the world from one period to the other: up and down
- Two financial assets: a stock and a risk free asset

Let $n = 2$, that is, three dates $0, 1, 2$. At the end of the first period, there are two possibilities: $S_0 \to S_1 = uS_0$ or $S_0 \to S_1 = dS_0$.

At the end of the second period, there are three possibilities: $S_1 \to S_2 = uuS_0$ or $S_1 \to S_2 = udS_0 = duS_0$ or $S_1 \to S_2 = ddS_0$.

We thus have a recombining tree: the price is the same after an upward and downward move that it is after a downward and upward move. Also, the price only depends on the previous period (Markov property).

Pricing by replication

The payoff function at maturity is

- after two ups $C_{uu} = \max(S_0 uu - K, 0)$
- after one up and one down (and vice versa) $C_{ud} = \max(S_0 ud - K, 0)$
- after two downs $C_{dd} = \max(S_0 dd - K, 0)$

Pricing by replication (cont'd)

Consider the second period, and the asset having the following payoffs:

- in up state $C_{uu} = \max(S_0uu - K, 0)$
- in down state $C_{ud} = \max(S_0ud - K, 0)$.

We see that we are back to a one-period model and that we can apply the same tools as in the previous slides, using the risk-neutral valuation.

We thus have:

- $C_u = \frac{1}{1+r}[\pi C_{uu} + (1-\pi)C_{ud}]$ after an upward move in the first period
- $C_d = \frac{1}{1+r}[\pi C_{ud} + (1-\pi)C_{dd}]$ otherwise.

Pricing by replication (cont'd)

By applying the same method recursively for the first period, we get $C = \frac{1}{1+r}[\pi C_u + (1-\pi)C_d]$, that is:

$C = \frac{1}{(1+r)^2}[\pi^2(S_0uu - K)_+ + 2\pi(1-\pi)(S_0ud - K)_+ + (1-\pi)^2(S_0dd - K)_+]$.

This expectation corresponds to a binomial law $B(2, \pi)$ whose events are $(S_0uu - K)_+$, $(S_0ud - K)_+$and $(S_0dd - K)_+$ with probabilities π^2, $2\pi(1-\pi)$ et $(1-\pi)^2$, respectively.

Again, the asset price is given by the discounted expectation of the future payoffs under the risk-neutral probability measure.

Generalization

We may apply the same technique with an arbitrary number of periods n. We get:

$C = \frac{1}{(1+r)^n} \sum_{j=0}^{n} \binom{n}{j} \pi^j (1-\pi)^{n-j} (u^j d^{n-j} S_0 - K)_+$, with $\binom{n}{j} = \frac{n!}{j!(n-j)!}$.

The expectation is taken with respect to the risk-neutral distribution $B(n, \pi)$.

Generalization (cont'd)

Taking an horizon T and dividing the interval $[0, T]$ in n subintervals of length T/n, we get that the binomial formula converges to the Black-Scholes formula as n tends to infinity, that is, a continuous-time model where the underlying asset is assumed to follow a geometric Brownian motion.

This kind of numerical approximation of a continuous-time model is extremely useful when computations are too difficult to perform explicitely.