Mathematical Foundations for Finance

Chapter III: Valuation and Hedging in Complete Markets

Prof. Dr. Beatrice Acciaio
beatrice.acciaio@math.ethz.ch

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Table of contents

- If in a market with (NA) we introduce a *new financial instrument* (e.g. an option) and stipulate that this should not create arbitrage opportunities, what can we say about the *price* of that instrument?

- If in a market with (NA) we introduce a *new financial instrument* (e.g. an option) and stipulate that this should not create arbitrage opportunities, what can we say about the *price* of that instrument?
- Depending on the structure of the stock's price process and of the option, the restrictions on the possible price of the option can be severe.
- In the extreme, the price of the option can be uniquely determined. However, this is the exception rather than the rule.

- If in a market with (NA) we introduce a *new financial instrument* (e.g. an option) and stipulate that this should not create arbitrage opportunities, what can we say about the *price* of that instrument?
- Depending on the structure of the stock's price process and of the option, the restrictions on the possible price of the option can be severe.
- In the extreme, the price of the option can be uniquely determined. However, this is the exception rather than the rule.

In this chapter, we consider a financial market in finite discrete time on $(\Omega, \mathcal{F}, P)$ with $\mathbb{F} = (\mathcal{F}_k)_{k=0,1,\dots,T}$, where discounted asset prices are given by $S^0 \equiv 1$ and $S = (S_k)_{k=0,1,\dots,T}$ with values in $\mathbb{R}^d$.

Attainable Payoffs

Contingent claim - definition

Definition. A *general European option* or *payoff* or *contingent claim* is a random variable $H \in L^0_+(\mathcal{F}_T)$.

Interpretation:

- H describes the *net payoff* (in units of asset 0) that the owner of this instrument obtains at time T.
- So $H \geq 0$ is natural and also avoids integrability issues.
- As H is $\mathcal{F}_T$-measurable, the payoff can depend on the entire information up to time T.
- "European" means that the time for the payoff is fixed at the end T.

Example. A *European **call** option* on *asset i* with *maturity T* and *strike K* gives its owner the right, but not the obligation, to buy at time T one unit of asset i for the price K, irrespective of what the actual future asset price S_T^i is.

Example. A *European **call** option* on *asset* i with *maturity* T and *strike* K gives its owner the right, but not the obligation, to buy at time T one unit of asset i for the price K, irrespective of what the actual future asset price S_T^i is.

- If we think in monetary terms, a rational person will *exercise* this right (or the option) if and only if $S_T^i(\omega) > K$, because it is in that, and only in that, situation that the right is more valuable than the asset.
- In that case, the net payoff is then $S_T^i(\omega) - K$. It is obtained by buying asset i at the low price K and immediately selling it on the market at the high price $S_T^i(\omega)$.

- In the other case $S^i_T(\omega) \leq K$, the option is worthless – it makes no monetary sense to pay K for one unit of asset i if one can get this on the market for less, namely for $S^i_T(\omega)$.

Contingent claim - example I

- In the other case $S^i_T(\omega) \leq K$, the option is worthless – it makes no monetary sense to pay K for one unit of asset i if one can get this on the market for less, namely for $S^i_T(\omega)$.
- So the option has a net payoff

$$H(\omega) = \max\big(0, S^i_T(\omega) - K\big) = \big(S^i_T(\omega) - K\big)^+.$$

  As a random variable, this is clearly non-negative and $\mathcal{F}_T$-measurable because S^i is adapted.
- In fact, H is even simpler here, because it only depends on the terminal asset price S^i_T. So we can write $H = h(S^i_T)$ with the function $h(x) = (x - K)^+$.

Example. If we want to *bet on* a reasonably *stable asset price evolution*, we might be interested in a payoff of the form $H = I_B$ with

$$B = \left\{ a \leq \min_{i=1,\dots,d} \min_{k=0,1,\dots,T} S_k^i < \max_{i=1,\dots,d} \max_{k=0,1,\dots,T} S_k^i \leq b \right\}.$$

- In words, this option pays at time T one unit of money if and only if all stocks remain between the levels a and b up to time T.
- This H is also $\mathcal{F}_T$-measurable, but now depends on the asset price evolution over the whole time range $k = 0, 1, \dots, T$. It cannot be written as a function of the final stock price S_T alone.

Example. A payoff of the form

$$H = I_A \, g\left(\frac{1}{T}\sum_{k=1}^{T} S_k^i\right) \qquad \text{with } A \in \mathcal{F}_T \text{ and a function } g \geq 0$$

depends on the average price (over time) of asset i, but it is only due in case that a certain event A occurs.

Example. A payoff of the form

$$H = I_A \, g\left(\frac{1}{T}\sum_{k=1}^{T} S_k^i\right) \qquad \text{with } A \in \mathcal{F}_T \text{ and a function } g \geq 0$$

depends on the average price (over time) of asset i, but it is only due in case that a certain event A occurs.

- In *insurance*, the set A could for instance be the event of the death up to time T of an insured person.
- In this case, H would describe the payoff from an *index-linked insurance policy*.
- This is an example in which H depends on more than only the basic asset prices.

The questions studied in this chapter are the following:

- **Pricing.** Given a contingent claim $H \in L^0_+(\mathcal{F}_T)$, how can we assign to H a *value at time* $k < T$ in such a way that this creates no arbitrage opportunities (if the claim is made available for trading at that value)?

The questions studied in this chapter are the following:

- **Pricing.** Given a contingent claim $H \in L^0_+(\mathcal{F}_T)$, how can we assign to H a *value at time $k < T$* in such a way that this creates no arbitrage opportunities (if the claim is made available for trading at that value)?
- **Hedging.** Having sold H, what can we do to *insure ourselves against the risk* involved in having to pay the random, uncertain amount H at time T?

The questions studied in this chapter are the following:

- **Pricing.** Given a contingent claim $H \in L^0_+(\mathcal{F}_T)$, how can we assign to H a *value at time $k < T$* in such a way that this creates no arbitrage opportunities (if the claim is made available for trading at that value)?
- **Hedging.** Having sold H, what can we do to *insure ourselves against the risk* involved in having to pay the random, uncertain amount H at time T?

The *key idea* for answering both questions is very simple: Try to *construct* an *"artificial product"* (with the help of S^0 and S) that looks as *similar to* H as possible.

The value of this "artificial product", by absence of arbitrage, will be a good approximation for the value of H.

Contingent claim - the attainable case

The value of this "artificial product", by absence of arbitrage, will be a good approximation for the value of H.

- Consider the *ideal case*: suppose we can find a self-financing strategy $\varphi \mathrel{\widehat{=}} (V_0, \vartheta)$ such that $V_T(\varphi) = H$ P-a.s.
- Then, (the strategy of just holding) H and the strategy φ both have costs 0 at all intermediate times $k = 1, \dots, T-1$ because φ is self-financing, and at time T both have a value of H.
- To avoid arbitrage, the values of both structures must therefore coincide at time 0 as well.

The value of this "artificial product", by absence of arbitrage, will be a good approximation for the value of H.

- Consider the *ideal case*: suppose we can find a self-financing strategy $\varphi \mathrel{\widehat{=}} (V_0, \vartheta)$ such that $V_T(\varphi) = H$ P-a.s.
- Then, (the strategy of just holding) H and the strategy φ both have costs 0 at all intermediate times $k = 1, \ldots, T-1$ because φ is self-financing, and at time T both have a value of H.
- To avoid arbitrage, the values of both structures must therefore coincide at time 0 as well.

Consequences:

1) The value or price of H at time 0 must be V_0.
2) Analogously, the value or price of H at any time k must be $V_k(\varphi)$.

We have a special terminology for this "ideal case".

Definition. A payoff $H \in L^0_+(\mathcal{F}_T)$ is called *attainable* if there exists an admissible self-financing strategy $\varphi \mathrel{\widehat{=}} (V_0, \vartheta)$ with $V_T(\varphi) = H$.

The strategy φ is then said to *replicate* H and is called a *replicating strategy for H*.

We have a special terminology for this "ideal case".

Definition. A payoff $H \in L^0_+(\mathcal{F}_T)$ is called *attainable* if there exists an admissible self-financing strategy $\varphi \mathrel{\widehat{=}} (V_0, \vartheta)$ with $V_T(\varphi) = H$.

The strategy φ is then said to *replicate* H and is called a *replicating strategy for* H.

The next result formalises the previous "key idea". In addition, it provides an efficient way of computing the resulting option price.

Theorem. (Arbitrage-free valuation of attainable payoffs)
Consider a financial market in finite discrete time and suppose that S is arbitrage-free and $\mathcal{F}_0$ is trivial. Then every attainable payoff H has a unique price process $V^H = (V_k^H)_{k=0,1,\dots,T}$ which admits no arbitrage (in the extended market (S^0, S, V^H)). This is given by

$$V_k^H = E_Q[H \,|\, \mathcal{F}_k] = V_k(V_0, \vartheta) \qquad \text{for } k = 0, 1, \dots, T,$$

for any equivalent martingale measure Q for S and for any replicating strategy $\varphi \mathrel{\widehat{=}} (V_0, \vartheta)$ for H.

Proof.

- By the DMW theorem, $\mathbb{P}_e(S)$ is non-empty because S is arbitrage-free. So there is at least one EMM Q.

Proof.

- By the DMW theorem, $\mathbb{P}_e(S)$ is non-empty because S is arbitrage-free. So there is at least one EMM Q.
- By assumption, H is attainable. So there is at least one replicating strategy φ.

Proof.

- By the DMW theorem, $\mathbb{P}_e(S)$ is non-empty because S is arbitrage-free. So there is at least one EMM Q.
- By assumption, H is attainable. So there is at least one replicating strategy φ.
- Because φ and H provide the same payoff structures and by absence of arbitrage (in the extended market), they must have the same value processes, $V^H = V(\varphi)$. This holds for any replicating φ.

- Note that V_0 is constant, because $\mathcal{F}_0$ is trivial.
- As any replicating φ is admissible by definition, the process $V(\varphi) = V_0 + \vartheta \cdot S = V(V_0, \vartheta)$ is uniformly bounded from below, and thus it is a Q-martingale (from a theorem previously proven), for any $Q \in \mathbb{P}_e(S)$.
- Since the final value of such a martingale is $V_T(\varphi) = H$, we get

$$V_k^H = V_k(\varphi) = E_Q[H \mid \mathcal{F}_k] \qquad \text{for all } k.$$

In terms of *efficiency*, the last theorem is a substantial achievement:

- In a first step, we need to check whether or not the basic model for S is arbitrage-free, and this is e.g. done by exhibiting or constructing an EMM Q for S.
- Then, for any attainable payoff we compute its price process by taking conditional expectations under Q, without needing to take care of finding a replication strategy.

In terms of *efficiency*, the last theorem is a substantial achievement:

- ▷ In a first step, we need to check whether or not the basic model for S is arbitrage-free, and this is e.g. done by exhibiting or constructing an EMM Q for S.
- ▷ Then, for any attainable payoff we compute its price process by taking conditional expectations under Q, without needing to take care of finding a replication strategy.

However:

- For hedging purposes, we are often interested in actually knowing and then also using a replicating strategy.
- Also, how can we decide for a given payoff whether or not it is attainable, without exhibiting or constructing a replicating strategy?

Theorem. (Characterisation of attainable payoffs)
Consider a financial market in finite discrete time and suppose that S is arbitrage-free and $\mathcal{F}_0$ is trivial. For any payoff $H \in L^0_+(\mathcal{F}_T)$, the following assertions are equivalent:

1) *H is attainable.*

2) *$\sup_{Q \in \mathbb{P}_{e,loc}(S)} E_Q[H] < \infty$ is attained for some $Q^* \in \mathbb{P}_{e,loc}(S)$, that is, the supremum is finite and a maximum (equal to $E_{Q^*}[H]$).*

3) *The mapping $\mathbb{P}_e(S) \to \mathbb{R}$, $Q \mapsto E_Q[H]$ is constant, that is, H has the same and finite expectation under all EMMs Q for S.*

Remarks.

- We omit the proof, which requires results not covered so far, in particular the optional decomposition theorem. For the case where prices S are nonnegative, see Föllmer/Schied (2011). The general case is more delicate; the simplification for $S \geq 0$ is due to the fact that the sets $\mathbb{P}_e(S)$ and $\mathbb{P}_{e,loc}(S)$ then coincide.
- In continuous or infinite discrete time, the implications 3) $\Rightarrow$ 1) $\Leftrightarrow$ 2) still hold (though with a slightly different definition of attainability), while 2) $\Rightarrow$ 3) may fail.

Contingent claim - summary

In summary, the approach to valuing and hedging a given payoff H in a financial market in finite discrete time (with $\mathcal{F}_0$ trivial) is:

1) Check if S is arbitrage-free by finding at least one ELMM Q for S.
2) Find all ELMMs Q for S.
3) Compute $E_Q[H]$ for all ELMMs Q for S and determine the supremum of $E_Q[H]$ over Q.

4a) If the supremum is finite and a maximum, i.e. attained in some $Q^* \in \mathbb{P}_{e,loc}(S)$, then H is attainable and its price process can be computed as $V_k^H = E_Q[H \,|\, \mathcal{F}_k]$, for any $Q \in \mathbb{P}_e(S)$.

4b) If the supremum is not attained (or, equivalently for finite discrete time, there is a pair of EMMs Q_1, Q_2 with $E_{Q_1}[H] \neq E_{Q_2}[H]$), then H is not attainable.

Contingent claim - summary

- In case 4a), the previous results tell us how to *value* H. But if we also want to find a *replicating strategy*, then more work is required.

Contingent claim - summary

- In case 4a), the previous results tell us how to *value* H. But if we also want to find a *replicating strategy*, then more work is required.
- In case 4b), we are faced with a *genuine problem*: It is impossible to replicate H, so our whole conceptual approach up to here breaks down.

- In case 4a), the previous results tell us how to *value* H. But if we also want to find a *replicating strategy*, then more work is required.
- In case 4b), we are faced with a *genuine problem*: It is impossible to replicate H, so our whole conceptual approach up to here breaks down.
- We then have the difficult problem of *valuation and hedging* for a *non-attainable payoff*, and there are in the literature several competing approaches to that, all involving in some way the specification of *preferences* or *subjective views* of the option seller.

Remark. The valuation from the last theorem does not involve preferences, but only the assumption of absence of arbitrage It is often called *risk-neutral valuation*, and an EMM Q for S is called a *risk-neutral measure*.

Warning: In large parts of the literature, the terminology "risk-neutral valuation" is used for computing conditional expectations of a given payoff H under some EMM Q. This is *problematic* for two reasons:

1) $V_k^{H,Q} := E_Q[H \mid \mathcal{F}_k]$ typically depends on Q if H is not attainable. So when following that approach, one should at the very least think carefully about which $Q \in \mathbb{P}_e(S)$ one uses, and why.

2) If H is not attainable, it is at best not clear how to hedge H in any reasonably safe way, and at worst, this may be impossible to achieve.

Both of these issues are often ignored in the literature. One area where this used to be particularly prominent is credit risk.

Complete Markets

Definition. A financial market model (in finite discrete time) is called *complete* if every payoff $H \in L^0_+(\mathcal{F}_T)$ is attainable. Otherwise it is called *incomplete*.

Theorem. (Valuation and hedging in complete markets)
Consider a financial market model in finite discrete time and suppose that $\mathcal{F}_0$ *is trivial and* S *is arbitrage-free and complete. Then for every payoff* $H \in L^0_+(\mathcal{F}_T)$*, there is a unique price process* $V^H = (V^H_k)_{k=0,1,\dots,T}$ *which admits no arbitrage. It is given by*

$$V^H_k = E_Q[H \,|\, \mathcal{F}_k] = V_k(V_0, \vartheta) \qquad \text{for } k = 0, 1, \dots, T$$

for any EMM Q *for* S *and any replicating strategy* $\varphi \mathrel{\widehat{=}} (V_0, \vartheta)$ *for* H.

The latter theorem raises the question of *how to recognise a complete market*, since completeness is a statement about *all* payoffs $H \in L^0_+(\mathcal{F}_T)$.

Fortunately, the next theorem gives a very simple criterion:.

Theorem. (Second fundamental theorem of asset pricing)
Consider a financial market model in finite discrete time and assume that S is arbitrage-free, $\mathcal{F}_0$ is trivial and $\mathcal{F}_T = \mathcal{F}$. Then S is complete if and only if there is a unique equivalent martingale measure for S. In brief:

$$(NA) + \textit{completeness} \iff \mathbb{P}_e(S) \text{ is a singleton.}$$

Proof. "$\Longleftarrow$": If $\mathbb{P}_e(S)$ contains only one element, then $Q \mapsto E_Q[H]$ is of course constant over $Q \in \mathbb{P}_e(S)$ for any $H \in L^0_+(\mathcal{F}_T)$. Hence H is attainable.

[To be accurate, one also needs to check a priori some integrability issues, namely that $E_Q[H] < \infty$ for at least one $Q \in \mathbb{P}_e(S)$; see Föllmer/Schied (2011), Theorems 5.30 and 5.26, for details.]

"$\Longrightarrow$": For any $A \in \mathcal{F}_T$, the payoff $H := I_A$ is attainable. So we have for any pair of EMMs Q_1, Q_2 for S that

$$Q_1[A] = E_{Q_1}[H] = V_0^H = E_{Q_2}[H] = Q_2[A].$$

So Q_1 and Q_2 coincide on $\mathcal{F}_T = \mathcal{F}$, which means that there is at most one EMM for S. By the DMW theorem, there is at least one EMM because S is arbitrage-free, and so the proof is complete. **q.e.d.**

Combining this result with the first FTAP, we have a very simple and beautiful description of financial market models in finite discrete time:

- ▷ **Existence** of an EMM is equivalent to the market being *arbitrage-free*.
- ▷ **Uniqueness** of the EMM is equivalent to *completeness* of the market.

For continuous or infinite discrete time, such statements become more subtle to formulate and more difficult to prove.

Remarks.

1) We can see from the proof of the theorem where the assumption $\mathcal{F}_T = \mathcal{F}$ is used. But it is also clear from looking at the statement why it is needed: *completeness is only an assertion about $\mathcal{F}_T$-measurable quantities.*

Remarks.

1) We can see from the proof of the theorem where the assumption $\mathcal{F}_T = \mathcal{F}$ is used. But it is also clear from looking at the statement why it is needed: *completeness is only an assertion about $\mathcal{F}_T$-measurable quantities.*

2) One can show that if a financial market in finite discrete time is complete, then $\mathcal{F}_T$ must be *finite*; see Föllmer/Schied (2011), Theorem 5.38. In effect, finiteness of $\mathcal{F}_T$ means that Ω can also be taken finite. This shows that while it makes the theory nice and simple, *completeness is a very restrictive property*.

Example. (Multinomial model) The multinomial model is incomplete whenever $m > 2$, that is, as soon as there is some node in the tree which allows more than two possible stock price evolutions. This follows since in this situation, there are infinitely many EMMs.

Example. (Multinomial model) The multinomial model is incomplete whenever $m > 2$, that is, as soon as there is some node in the tree which allows more than two possible stock price evolutions. This follows since in this situation, there are infinitely many EMMs.

Example. Consider any model with $d = 1$ (one risky asset) and i.i.d. returns $Y_1, \ldots, Y_T$ under P. If Y_1 has a density (e.g. if we have lognormal returns), then S is incomplete. This is because $\mathcal{F}_1$ (and hence also $\mathcal{F}_T$) must be infinite for Y_1 to have a density.

Example: The Binomial Model

Recall:

- The *binomial* or *Cox–Ross–Rubinstein model* is described by parameters $p \in (0,1)$ and $u > r > d > -1$, via

$$\widetilde{S}^0_k = (1+r)^k \quad \text{and} \quad \widetilde{S}^1_k = S^1_0 \prod_{j=1}^{k} Y_j\,,$$

  where $S^1_0 > 0$ and $Y_1, \dots, Y_T$ are i.i.d. under P taking values $1+u$ or $1+d$ with probability p or $1-p$, respectively.

- The filtration $\mathbb{F}$ is generated by $\widetilde{S} = (\widetilde{S}^0, \widetilde{S}^1)$ or, equivalently, by $\widetilde{S}^1$ or by Y. $\mathcal{F}_0$ is then trivial because $\widetilde{S}^0_0 = 1$ and $\widetilde{S}^1_0 = S^1_0$ is a constant.

- We also take $\mathcal{F} = \mathcal{F}_T$. This is even an automatic conclusion if we construct the model on the canonical path space as in Chapter 1.

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

- We already know from Chapter 2 that this model is *arbitrage-free* and has a *unique EMM* for $S^1 = \widetilde{S}^1/\widetilde{S}^0$.
- Therefore, S^1 is *complete*, and so every $H \in L^0_+(\mathcal{F}_T)$ is *attainable*, with a *price process* given by

$$V_k^H = E_{Q^*}[H \mid \mathcal{F}_k] \qquad \text{for } k = 0, 1, \dots, T\,,$$

  where Q^* is the unique EMM for S^1.
- We also recall that the Y_j are under Q^* again i.i.d., but with

$$Q^*[Y_1 = 1 + u] = q^* := \frac{r - d}{u - d} \in (0, 1).$$

The binomial model - corollary

- All the above quantities S^1, H, V^H are discounted with $\widetilde{S}^0$, that is, they are expressed in units of asset 0.
- The undiscounted quantities are the stock price $\widetilde{S}^1 = S^1 \widetilde{S}^0$, the payoff $\widetilde{H} := H \widetilde{S}^0_T$ and its price process $\widetilde{V}^{\widetilde{H}}$ with $\widetilde{V}^{\widetilde{H}}_k := V^H_k \widetilde{S}^0_k$ for $k = 0, 1, \dots, T$.

- All the above quantities S^1, H, V^H are discounted with $\widetilde{S}^0$, that is, they are expressed in units of asset 0.
- The undiscounted quantities are the stock price $\widetilde{S}^1 = S^1 \widetilde{S}^0$, the payoff $\widetilde{H} := H \widetilde{S}^0_T$ and its price process $\widetilde{V}^{\widetilde{H}}$ with $\widetilde{V}^{\widetilde{H}}_k := V^H_k \widetilde{S}^0_k$ for $k = 0, 1, \dots, T$.

Corollary. *In the binomial model, the undiscounted arbitrage-free price process of any undiscounted payoff $\widetilde{H} \in L^0_+(\mathcal{F}_T)$ is given by*

$$\widetilde{V}^{\widetilde{H}}_k = \widetilde{S}^0_k E_{Q^*}\left[\frac{\widetilde{H}}{\widetilde{S}^0_T} \,\middle|\, \mathcal{F}_k \right] = E_{Q^*}\left[\widetilde{H} \frac{\widetilde{S}^0_k}{\widetilde{S}^0_T} \,\middle|\, \mathcal{F}_k \right] = \frac{\widetilde{S}^0_k}{\widetilde{S}^0_T} E_{Q^*}\left[\widetilde{H} \,\middle|\, \mathcal{F}_k \right]$$

for $k = 0, 1, \dots, T$.

- For a *European call option* on $\widetilde{S}^1$ with maturity T and undiscounted strike $\widetilde{K}$, we have

$$\widetilde{H} = \left(\widetilde{S}^1_T - \widetilde{K}\right)^+ = \left(\widetilde{S}^1_T - \widetilde{K}\right) I_{\{\widetilde{S}^1_T > \widetilde{K}\}}.$$

- For every $k = 0, 1, \dots, T-1$, we can rewrite the set as

$$\{\widetilde{S}^1_T > \widetilde{K}\} = \left\{\widetilde{S}^1_k \prod_{j=k+1}^{T} Y_j > \widetilde{K}\right\} = \left\{\sum_{j=k+1}^{T} \log Y_j > \log\left(\widetilde{K}/\widetilde{S}^1_k\right)\right\}.$$

- For a *European call option* on $\widetilde{S}^1$ with maturity T and undiscounted strike $\widetilde{K}$, we have

$$\widetilde{H} = \left(\widetilde{S}^1_T - \widetilde{K}\right)^+ = \left(\widetilde{S}^1_T - \widetilde{K}\right) I_{\{\widetilde{S}^1_T > \widetilde{K}\}} .$$

- For every $k = 0, 1, ..., T-1$, we can rewrite the set as

$$\{\widetilde{S}^1_T > \widetilde{K}\} = \left\{ \widetilde{S}^1_k \prod_{j=k+1}^{T} Y_j > \widetilde{K} \right\} = \left\{ \sum_{j=k+1}^{T} \log Y_j > \log\left(\widetilde{K}/\widetilde{S}^1_k\right) \right\}.$$

- If we define

$$W_j := I_{\{Y_j = 1+u\}} = \begin{cases} 1 & \text{if } Y_j = 1+u \\ 0 & \text{if } Y_j = 1+d, \end{cases}$$

then $W_1, \ldots, W_T$ are under Q^* independent 0-1 experiments with success parameter q^*, so that their sum has under Q^* a *binomial distribution*.

The binomial model - consequence

- We can now write $\log Y_j = W_j \log(1+u) + (1-W_j)\log(1+d)$, and this gives us

$$\sum_{j=k+1}^{T} \log Y_j = W_{k,T} \log(1+u) + (T-k-W_{k,T})\log(1+d)$$
$$= W_{k,T} \log \frac{1+u}{1+d} + (T-k)\log(1+d),$$

where $W_{k,T} := \sum_{j=k+1}^{T} W_j \sim \mathrm{Bin}(T-k, q^*)$ is independent of $\mathcal{F}_k$ under Q^*.

- So we get

$$\{\widetilde{S}^1_T > \widetilde{K}\} = \left\{ W_{k,T} \log \frac{1+u}{1+d} > \log \frac{\widetilde{K}}{\widetilde{S}^1_k} - (T-k)\log(1+d) \right\}.$$

- Since $W_{k,T}$ is independent of $\mathcal{F}_k$ under Q^* and $\widetilde{S}_k^1$ is $\mathcal{F}_k$-measurable, we get

$$Q^*\left[\widetilde{S}_T^1 > \widetilde{K} \,\middle|\, \mathcal{F}_k\right] = Q^*\left[W_{k,T} > \frac{\log \frac{\widetilde{K}}{s} - (T-k)\log(1+d)}{\log \frac{1+u}{1+d}}\right]\Bigg|_{s=\widetilde{S}_k^1}.$$

- Since $W_{k,T}$ is independent of $\mathcal{F}_k$ under Q^* and $\widetilde{S}^1_k$ is $\mathcal{F}_k$-measurable, we get

$$Q^*\big[\widetilde{S}^1_T > \widetilde{K} \,\big|\, \mathcal{F}_k\big] = Q^*\left[W_{k,T} > \frac{\log \frac{\widetilde{K}}{s} - (T-k)\log(1+d)}{\log \frac{1+u}{1+d}}\right]\Bigg|_{s=\widetilde{S}^1_k}.$$

- The above probability can be computed explicitly because $W_{k,T}$ has a binomial distribution. And since

$$E_{Q^*}\big[\widetilde{H} \,\big|\, \mathcal{F}_k\big] = E_{Q^*}\left[\widetilde{S}^1_T I_{\{\widetilde{S}^1_T > \widetilde{K}\}} \,\middle|\, \mathcal{F}_k\right] - \widetilde{K} Q^*\big[\widetilde{S}^1_T > \widetilde{K} \,\big|\, \mathcal{F}_k\big],$$

we already have the second half of the *binomial call pricing formula*.

For the first term, one can either use explicit (and lengthy) computations or more elegantly a so-called *change of numéraire*: since $\widetilde{S}^1/\widetilde{S}^0 = S^1$ is under Q^* a positive martingale, one can use it to define a probability measure $\widetilde{Q}^* \approx Q^*$ on $\mathcal{F}_T$ via

$$d\widetilde{Q}^*/dQ^* := S^1_T/S^1_0 .$$

Then S^1/S^1_0 is by construction the density process of $\widetilde{Q}^*$ w.r.t. Q^*.

Then Bayes formula gives

$$\begin{aligned} E_{Q^*}\left[\widetilde{S}^1_T I_{\{\widetilde{S}^1_T > \widetilde{K}\}} \,\middle|\, \mathcal{F}_k\right] &= \widetilde{S}^1_k \frac{\widetilde{S}^0_T}{\widetilde{S}^0_k} \frac{\widetilde{S}^0_k}{\widetilde{S}^1_k} E_{Q^*}\left[\frac{\widetilde{S}^1_T}{\widetilde{S}^0_T} I_{\{\widetilde{S}^1_T > \widetilde{K}\}} \,\middle|\, \mathcal{F}_k\right] \\ &= \widetilde{S}^1_k \frac{\widetilde{S}^0_T}{\widetilde{S}^0_k} \widetilde{Q}^*\left[\widetilde{S}^1_T > \widetilde{K} \,\middle|\, \mathcal{F}_k\right] \\ &= \widetilde{S}^1_k \frac{\widetilde{S}^0_T}{\widetilde{S}^0_k} \widetilde{Q}^*\left[W_{k,T} > \frac{\log\frac{\widetilde{K}}{s} - (T-k)\log(1+d)}{\log\frac{1+u}{1+d}}\right]\Bigg|_{s=\widetilde{S}^1_k} \end{aligned}$$

Note that the above quantity can be computed since $W_{k,T}$ under $\widetilde{Q}^*$ is $\mathrm{Bin}(T-k, \widetilde{q}^*)$-distributed with

$$\widetilde{q}^* := q^* \frac{1+u}{1+r}, \qquad \text{hence } 1 - \widetilde{q}^* = (1-q^*)\frac{1+d}{1+r}.$$

Note that the above quantity can be computed since $W_{k,T}$ under $\widetilde{Q}^*$ is $\mathrm{Bin}(T-k, \widetilde{q}^*)$-distributed with

$$\widetilde{q}^* := q^* \frac{1+u}{1+r}, \qquad \text{hence } 1-\widetilde{q}^* = (1-q^*)\frac{1+d}{1+r}.$$

- One can verify [$\rightarrow$ *exercise*] that $\widetilde{Q}^*$ is the unique probability measure equivalent to P on $\mathcal{F}_T$ such that $1/S^1 = \widetilde{S}^0/\widetilde{S}^1$ becomes a $\widetilde{Q}^*$-martingale, and one can also check that $Y_1, \dots, Y_T$ are i.i.d. under $\widetilde{Q}^*$ with $\widetilde{Q}^*[Y_1 = 1+u] = \widetilde{q}^*$.
- The measure $\widetilde{Q}^*$ is also called *dual martingale measure*.

All in all: We obtain the fairly simple formula

$$\widetilde{V}_k^{\widetilde{H}} = \widetilde{S}_k^1 \widetilde{Q}^*[W_{k,T} > x] - \widetilde{K} \frac{\widetilde{S}_k^0}{\widetilde{S}_T^0} Q^*[W_{k,T} > x]$$

with

$$x = \frac{\log \frac{\widetilde{K}}{s} - (T-k)\log(1+d)}{\log \frac{1+u}{1+d}} \quad \text{for } s = \widetilde{S}_k^1.$$

This *binomial call pricing formula* is the *discrete analogue of* the famous *Black–Scholes formula*.

The binomial model - general payoffs

- For a *general payoff* $\widetilde{H}$, the discounted price process V^H is, by its construction, a Q^*-martingale with final value H, so that $V^H_T = H$ and

$$V^H_{k-1} = E_{Q^*}\left[V^H_k \,\middle|\, \mathcal{F}_{k-1}\right] \qquad \text{for } k = 1, \ldots, T.$$

- For a *general payoff* $\widetilde{H}$, the discounted price process V^H is, by its construction, a Q^*-martingale with final value H, so that $V_T^H = H$ and

$$V_{k-1}^H = E_{Q^*}\left[V_k^H \,\middle|\, \mathcal{F}_{k-1}\right] \qquad \text{for } k = 1, \dots, T.$$

- This provides a *simple recursive algorithm*, since the filtration $\mathbb{F}$ in the binomial model has the structure of a *binary tree*:

  Indeed, if we fix some node at time $k-1$ (corresponding to some atom of $\mathcal{F}_{k-1}$) and denote by v_{k-1} the value of V_{k-1}^H there, then there are only two possible successor nodes and V_k^H can only take two values there, say v_k^u and v_k^d.

The binomial model - general payoffs

- For a *general payoff* $\widetilde{H}$, the discounted price process V^H is, by its construction, a Q^*-martingale with final value H, so that $V_T^H = H$ and

$$V_{k-1}^H = E_{Q^*}\left[V_k^H \,\middle|\, \mathcal{F}_{k-1}\right] \qquad \text{for } k = 1, \dots, T.$$

- This provides a *simple recursive algorithm*, since the filtration $\mathbb{F}$ in the binomial model has the structure of a *binary tree*:

  Indeed, if we fix some node at time $k-1$ (corresponding to some atom of $\mathcal{F}_{k-1}$) and denote by v_{k-1} the value of V_{k-1}^H there, then there are only two possible successor nodes and V_k^H can only take two values there, say v_k^u and v_k^d.

- The Q^*-martingale property then says that

$$v_{k-1} = q^* v_k^u + (1-q^*) v_k^d ,$$

  because the one-step transition probabilities of Q^* are the same throughout the tree, namely q^* and $1-q^*$.

- In undiscounted terms, we have

$$\frac{\widetilde{V}^{\widetilde{H}}_{k-1}}{\widetilde{S}^0_{k-1}} = E_{Q^*}\left[\frac{\widetilde{V}^{\widetilde{H}}_k}{\widetilde{S}^0_k} \,\middle|\, \mathcal{F}_{k-1}\right]$$

  or

$$\widetilde{V}^{\widetilde{H}}_{k-1} = \frac{1}{1+r} E_{Q^*}\big[\widetilde{V}^{\widetilde{H}}_k \,\big|\, \mathcal{F}_{k-1}\big],$$

  which translates at the level of node values to the *recursion*

$$\tilde{v}_{k-1} = \frac{1}{1+r}\big(q^* \tilde{v}^u_k + (1-q^*)\tilde{v}^d_k\big). \tag{1}$$

- In undiscounted terms, we have

$$\frac{\widetilde{V}^{\widetilde{H}}_{k-1}}{\widetilde{S}^0_{k-1}} = E_{Q^*}\left[\frac{\widetilde{V}^{\widetilde{H}}_{k}}{\widetilde{S}^0_{k}}\,\middle|\,\mathcal{F}_{k-1}\right]$$

or

$$\widetilde{V}^{\widetilde{H}}_{k-1} = \frac{1}{1+r}E_{Q^*}\left[\widetilde{V}^{\widetilde{H}}_{k}\,\middle|\,\mathcal{F}_{k-1}\right],$$

which translates at the level of node values to the *recursion*

$$\tilde{v}_{k-1} = \frac{1}{1+r}\left(q^*\tilde{v}^u_k + (1-q^*)\tilde{v}^d_k\right). \tag{1}$$

- The *terminal condition* $\widetilde{V}^{\widetilde{H}}_T = \widetilde{H}$ means that the values $\tilde{v}_T$ at the terminal nodes are given by the values of $\widetilde{H}$ there.

- To work out the *replicating strategy*, also for a general payoff $\widetilde{H}$, we recall that

$$V_k^H = V_k(V_0, \vartheta) = V_0 + \sum_{j=1}^{k} \vartheta_j \Delta S_j^1 \qquad \text{for } k = 0, 1, \ldots, T.$$

- To work out the *replicating strategy*, also for a general payoff $\widetilde{H}$, we recall that

$$V_k^H = V_k(V_0, \vartheta) = V_0 + \sum_{j=1}^{k} \vartheta_j \Delta S_j^1 \qquad \text{for } k = 0, 1, \dots, T.$$

- For the *increments*, this means that

$$\Delta V_k^H = V_k^H - V_{k-1}^H = \vartheta_k \Delta S_k^1 = \vartheta_k (S_k^1 - S_{k-1}^1). \tag{2}$$

- To work out the *replicating strategy*, also for a general payoff $\widetilde{H}$, we recall that

$$V_k^H = V_k(V_0, \vartheta) = V_0 + \sum_{j=1}^{k} \vartheta_j \Delta S_j^1 \qquad \text{for } k = 0, 1, \dots, T.$$

- For the *increments*, this means that

$$\Delta V_k^H = V_k^H - V_{k-1}^H = \vartheta_k \Delta S_k^1 = \vartheta_k (S_k^1 - S_{k-1}^1). \tag{2}$$

- As ϑ_k is $\mathcal{F}_{k-1}$-measurable (because ϑ is predictable), the value of ϑ_k is already known at time $k-1$, hence in that node, and it cannot change as we move forward to time k.

- If we denote as before by v_{k-1} the value of V^H_{k-1} in the chosen node at time $k-1$ and by s_{k-1} the value of S^1_{k-1} there, we know that, in the next step
  - v_{k-1} evolves to either v^u_k or v^d_k, and
  - s_{k-1} evolves to $s^u_k = s_{k-1}\frac{1+u}{1+r}$ or $s^d_k = s_{k-1}\frac{1+d}{1+r}$.

- If we denote as before by v_{k-1} the value of V^H_{k-1} in the chosen node at time $k-1$ and by s_{k-1} the value of S^1_{k-1} there, we know that, in the next step
  - v_{k-1} evolves to either v^u_k or v^d_k, and
  - s_{k-1} evolves to $s^u_k = s_{k-1}\frac{1+u}{1+r}$ or $s^d_k = s_{k-1}\frac{1+d}{1+r}$.
- The relation (2) between increments must hold in all nodes and at all times. So if ξ_k denotes the value of ϑ_k in the chosen node at time $k-1$, we obtain the *two equations*

$$v^u_k - v_{k-1} = \xi_k(s^u_k - s_{k-1}),$$
$$v^d_k - v_{k-1} = \xi_k(s^d_k - s_{k-1}).$$

- The above two equations are readily solved to give

$$\xi_k = \frac{v_k^u - v_k^d}{s_k^u - s_k^d} = \frac{v_k^u - v_k^d}{\frac{u-d}{1+r} s_{k-1}}. \tag{3}$$

- The above two equations are readily solved to give

$$\xi_k = \frac{v_k^u - v_k^d}{s_k^u - s_k^d} = \frac{v_k^u - v_k^d}{\frac{u-d}{1+r} s_{k-1}}. \tag{3}$$

- At time $k = T$, the right-hand side is known, because $V_T^H = H$. So both the price process V^H and the hedging strategy ϑ can be computed in parallel while working backwards through the tree.

- The above two equations are readily solved to give

$$\xi_k = \frac{v_k^u - v_k^d}{s_k^u - s_k^d} = \frac{v_k^u - v_k^d}{\frac{u-d}{1+r} s_{k-1}}. \tag{3}$$

- At time $k = T$, the right-hand side is known, because $V_T^H = H$. So both the price process V^H and the hedging strategy ϑ can be computed in parallel while working backwards through the tree.
- If the payoff $\widetilde{H}$ is like the call option of the *simple path-independent form* $\widetilde{H} = \tilde{h}(\widetilde{S}_T^1)$ for some function $\tilde{h}$, then the above formulas and computation scheme simplify considerably.

The binomial model - general payoffs

- Indeed one can show by backward induction that

$$\widetilde{V}_k^{\widetilde{H}} = \tilde{v}\big(k, \widetilde{S}_k^1\big) \qquad \text{for } k = 0, 1, \dots, T$$

and

$$\vartheta_k = \tilde{\xi}\big(k, \widetilde{S}_{k-1}^1\big) \qquad \text{for } k = 1, \dots, T$$

▷ Indeed one can show by backward induction that

$$\widetilde{V}_k^{\widetilde{H}} = \tilde{v}\big(k, \widetilde{S}_k^1\big) \qquad \text{for } k = 0, 1, \dots, T$$

and

$$\vartheta_k = \tilde{\xi}\big(k, \widetilde{S}_{k-1}^1\big) \qquad \text{for } k = 1, \dots, T$$

with *functions* $\tilde{v}$ and $\tilde{\xi}$ that are given by the *recursions*:

- $\tilde{v}(k-1, s) = \frac{1}{1+r}\Big(q^* \tilde{v}\big(k, s(1+u)\big) + (1-q^*)\tilde{v}\big(k, s(1+d)\big)\Big)$ (compare (1)), with terminal condition

$$\tilde{v}(T, s) = \tilde{h}(s)$$

- and $\tilde{\xi}(k, s) = \dfrac{\tilde{v}\big(k, s(1+u)\big) - \tilde{v}\big(k, s(1+d)\big)}{(u-d)s}$, from (3) multiplying numerator and denominator by $\widetilde{S}_k^0 = (1+r)^k$.

- In particular, here, it is enough to do all the computations in the *simplified, recombining tree* because neither $\widetilde{V}^{\widetilde{H}}$ nor ϑ have any path-dependence, but they are only depend on current values of $\widetilde{S}^1$.

- In particular, here, it is enough to do all the computations in the *simplified, recombining tree* because neither $\widetilde{V}^{\widetilde{H}}$ nor ϑ have any path-dependence, but they are only depend on current values of $\widetilde{S}^1$.
- So instead of 2^T terminal nodes for all the trajectories ω, we need here only $T + 1$ terminal nodes, for all the possible values of $\widetilde{S}^1_T$.

- In particular, here, it is enough to do all the computations in the *simplified, recombining tree* because neither $\widetilde{V}^{\widetilde{H}}$ nor ϑ have any path-dependence, but they are only depend on current values of $\widetilde{S}^1$.
- So instead of 2^T terminal nodes for all the trajectories ω, we need here only $T+1$ terminal nodes, for all the possible values of $\widetilde{S}^1_T$.
- The corresponding tree is therefore also massively smaller, and so are computation times and storage requirements.