

Optional Stopping Theorem (continued).

The OST is important in many areas, such as sequential analysis in statistics. We turn in the next section to related ideas specific to the gambling/financial context.

Write $X_n^T := X_{n \wedge T}$ for the sequence (X_n) *stopped* at time T .

Proposition. (i) If (X_n) is adapted and T is a stopping-time, the stopped sequence $(X_{n \wedge T})$ is adapted.

(ii) If (X_n) is a martingale [supermartingale] and T is a stopping time, (X_n^T) is a martingale [supermartingale].

Proof. If $\phi_j := I\{j \leq T\}$,

$$X_{T \wedge n} = X_0 + \sum_{j=1}^n \phi_j (X_j - X_{j-1}).$$

Since $\{j \leq T\}$ is the complement of $\{T < j\} = \{T \leq j-1\} \in \mathcal{F}_{j-1}$, $\phi_j = I\{j \leq T\} \in \mathcal{F}_{j-1}$, so (ϕ_n) is previsible. So (X_n^T) is adapted.

If (X_n) is a martingale, so is (X_n^T) as it is the martingale transform of (X_n) by (ϕ_n) . Since by previsibility of (ϕ_n)

$$E[X_{T \wedge n} | \mathcal{F}_{n-1}] = X_0 + \sum_{j=1}^{n-1} \phi_j (X_j - X_{j-1}) + \phi_n (E[X_n | \mathcal{F}_{n-1}] - X_{n-1}),$$

i.e.

$$E[X_{T \wedge n} | \mathcal{F}_{n-1}] - X_{T \wedge n} = \phi_n (E[X_n | \mathcal{F}_{n-1}] - X_{n-1}),$$

$\phi_n \geq 0$ shows that if (X_n) is a supermg [submg], so is $(X_{T \wedge n})$. //

§7. The Snell Envelope and Optimal Stopping.

Definition. If $Z = (Z_n)_{n=0}^N$ is a sequence adapted to a filtration (\mathcal{F}_n) , the sequence $U = (U_n)_{n=0}^N$ defined by

$$\begin{cases} U_N := Z_N, \\ U_n := \max(Z_n, E[U_{n+1} | \mathcal{F}_n]) \end{cases} \quad (n \leq N-1)$$

is called the *Snell envelope* of Z (J. L. Snell in 1952; [N] Ch. 6). U is adapted, i.e. $U_n \in \mathcal{F}_n$ for all n . For, Z is adapted, so $Z_n \in \mathcal{F}_n$. Also $E[U_{n+1}|\mathcal{F}_n] \in \mathcal{F}_n$ (definition of conditional expectation). Combining, $U_n \in \mathcal{F}_n$, as required.

We shall see in IV.8 [L20] that the Snell envelope is exactly the tool needed in pricing American options. It is the *least supermg majorant* (also called the *réduite* or *reduced function* – crucial in the mathematics of gambling):

Theorem. The Snell envelope (U_n) of (Z_n) is a supermartingale, and is the smallest supermartingale dominating (Z_n) (that is, with $U_n \geq Z_n$ for all n).

Proof. First, $U_n \geq E[U_{n+1}|\mathcal{F}_n]$, so U is a supermartingale, and $U_n \geq Z_n$, so U dominates Z .

Next, let $T = (T_n)$ be any other supermartingale dominating Z ; we must show T dominates U also. First, since $U_N = Z_N$ and T dominates Z , $T_N \geq U_N$. Assume inductively that $T_n \geq U_n$. Then

$$\begin{aligned} T_{n-1} &\geq E[T_n|\mathcal{F}_{n-1}] \quad (\text{as } T \text{ is a supermartingale}) \\ &\geq E[U_n|\mathcal{F}_{n-1}] \quad (\text{by the induction hypothesis}) \end{aligned}$$

and

$$T_{n-1} \geq Z_{n-1} \quad (\text{as } T \text{ dominates } Z).$$

Combining,

$$T_{n-1} \geq \max(Z_{n-1}, E[U_n|\mathcal{F}_{n-1}]) = U_{n-1}.$$

By backward induction, $T_n \geq U_n$ for all n , as required. //

Note. It is no accident that we are using induction here *backwards in time*. We will use the same method – also known as *dynamic programming (DP)* – in Ch. IV below when we come to pricing American options.

Proposition. $T_0 := \min\{n \geq 0 : U_n = Z_n\}$ is a stopping time, and the stopped sequence $(U_n^{T_0})$ is a martingale.

Proof (not examinable). Since $U_N = Z_N$, $T_0 \in \{0, 1, \dots, N\}$ is well-defined (and we can use minimum rather than infimum). For $k = 0$, $\{T_0 = 0\} = \{U_0 = Z_0\} \in \mathcal{F}_0$; for $k \geq 1$,

$$\{T_0 = k\} = \{U_0 > Z_0\} \cap \dots \cap \{U_{k-1} > Z_{k-1}\} \cap \{U_k = Z_k\} \in \mathcal{F}_k.$$

So T_0 is a stopping-time.

As in the proof of the Proposition in §6,

$$U_n^{T_0} = U_{n \wedge T_0} = U_0 + \sum_1^n \phi_j \Delta U_j,$$

where $\phi_j = I\{T_0 \geq j\}$ is adapted. For $n \leq N - 1$,

$$U_{n+1}^{T_0} - U_n^{T_0} = \phi_{n+1}(U_{n+1} - U_n) = I\{n+1 \leq T_0\}(U_{n+1} - U_n).$$

Now $U_n := \max(Z_n, E[U_{n+1}|\mathcal{F}_n])$, and

$$U_n > Z_n \quad \text{on } \{n+1 \leq T_0\}.$$

So from the definition of U_n ,

$$U_n = E[U_{n+1}|\mathcal{F}_n] \quad \text{on } \{n+1 \leq T_0\}.$$

We next prove

$$U_{n+1}^{T_0} - U_n^{T_0} = I\{n+1 \leq T_0\}(U_{n+1} - E[U_{n+1}|\mathcal{F}_n]). \quad (1)$$

For, suppose first that $T_0 \geq n+1$. Then the left of (1) is $U_{n+1} - U_n$, the right is $U_{n+1} - E[U_{n+1}|\mathcal{F}_n]$, and these agree on $\{n+1 \leq T_0\}$ by above. The other possibility is that $T_0 < n+1$, i.e. $T_0 \leq n$. Then the left of (1) is $U_{T_0} - U_{T_0} = 0$, while the right is zero because the indicator is zero. Combining, this proves (1) as required. Apply $E[\cdot|\mathcal{F}_n]$ to (1): since $\{n+1 \leq T_0\} = \{T_0 \leq n\}^c \in \mathcal{F}_n$,

$$\begin{aligned} E[(U_{n+1}^{T_0} - U_n^{T_0})|\mathcal{F}_n] &= I\{n+1 \leq T_0\}E[(U_{n+1} - E[U_{n+1}|\mathcal{F}_n])|\mathcal{F}_n] \\ &= I\{n+1 \leq T_0\}(E[U_{n+1}|\mathcal{F}_n] - E[U_{n+1}|\mathcal{F}_n]) = 0. \end{aligned}$$

So $E[U_{n+1}^{T_0}|\mathcal{F}_n] = U_n^{T_0}$. This says that $U_n^{T_0}$ is a martingale, as required. //

Note. Just because U is a supermartingale, we knew that stopping it would give a supermartingale, by the Proposition of §6. The point is that, using the special properties of the Snell envelope, we actually get a *martingale*.

Write $\mathcal{T}_{n,N}$ for the set of stopping times taking values in $\{n, n+1, \dots, N\}$ (a finite set, as Ω is finite). We next see that the Snell envelope solves the *optimal stopping problem*: it *maximises* the expectation of our final value of Z – the value when we choose to quit – conditional on our present (publicly

available) information. This is the best we can hope to do in practice (without cheating – insider trading, etc.)

Theorem. T_0 solves the optimal stopping problem for Z :

$$U_0 = E[Z_{T_0}|\mathcal{F}_0] = \max\{E[Z_T|\mathcal{F}_0] : T \in \mathcal{T}_{0,N}\}.$$

Proof. As $(U_n^{T_0})$ is a martingale (above),

$$\begin{aligned} U_0 &= U_0^{T_0} && (\text{since } 0 = 0 \wedge T_0) \\ &= E[U_N^{T_0}|\mathcal{F}_0] && (\text{by the martingale property}) \\ &= E[U_{T_0}|\mathcal{F}_0] && (\text{since } T_0 = T_0 \wedge N) \\ &= E[Z_{T_0}|\mathcal{F}_0] && (\text{since } U_{T_0} = Z_{T_0}), \end{aligned}$$

proving the first statement. Now for any stopping time $T \in \mathcal{T}_{0,N}$, since U is a supermartingale (above), so is the stopped process (U_n^T) (§6). So

$$\begin{aligned} U_0 &= U_0^T && (0 = 0 \wedge T, \text{ as above}) \\ &\geq E[U_N^T|\mathcal{F}_0] && ((U_n^T) \text{ a supermartingale}) \\ &= E[U_T|\mathcal{F}_0] && (T = T \wedge N) \\ &\geq E[Z_T|\mathcal{F}_0] && ((U_n) \text{ dominates } (Z_n)), \end{aligned}$$

and this completes the proof. //

The same argument, starting at time n rather than time 0, gives an apparently more general version:

Theorem. If $T_n := \min\{j \geq n : U_j = Z_j\}$,

$$U_n = E[Z_{T_n}|\mathcal{F}_n] = \sup\{E[Z_T|\mathcal{F}_n] : T \in \mathcal{T}_{n,N}\}.$$

To recapitulate: as we are attempting to maximise our payoff by stopping $Z = (Z_n)$ at the most advantageous time, the Theorem shows that T_n gives the best stopping-time that is realistic: it maximises our *expected payoff* given only information *currently available* (it is easy, but irrelevant, to maximise things with hindsight!). We thus call T_0 (or T_n , starting from time n) the *optimal* stopping time for the problem.

§8. Doob Decomposition.

Theorem. Let $X = (X_n)$ be an adapted process with each $X_n \in L_1$. Then X has an (essentially unique) Doob decomposition

$$X = X_0 + M + A : \quad X_n = X_0 + M_n + A_n \quad \forall n \quad (D)$$

with M a martingale null at zero, A a previsible process null at zero. If also X is a submartingale (‘increasing on average’), A is increasing: $A_n \leq A_{n+1}$ for all n , a.s.

Proof. If X has a Doob decomposition (D),

$$E[X_n - X_{n-1} | \mathcal{F}_{n-1}] = E[M_n - M_{n-1} | \mathcal{F}_{n-1}] + E[A_n - A_{n-1} | \mathcal{F}_{n-1}].$$

The first term on the right is zero, as M is a martingale. The second is $A_n - A_{n-1}$, since A_n (and A_{n-1}) is \mathcal{F}_{n-1} -measurable by previsibility. So

$$E[X_n - X_{n-1} | \mathcal{F}_{n-1}] = A_n - A_{n-1}, \quad (1)$$

and summation gives

$$A_n = \sum_{k=1}^n E[X_k - X_{k-1} | \mathcal{F}_{k-1}], \quad a.s.$$

We use this formula to *define* (A_n) , clearly previsible. We then use (D) to *define* (M_n) , then a martingale, giving the Doob decomposition (D).

If X is a submartingale, the LHS of (1) is ≥ 0 , so the RHS of (1) is ≥ 0 , i.e. (A_n) is increasing. //

Note. 1. Although the Doob decomposition in discrete time is simple, the continuous-time analogue is deep (Ch. V). This illustrates the contrasts between the theories of stochastic processes in discrete and continuous time.

2. Previsible processes are ‘easy’ (trading is easy if you can foresee price movements!). So the Doob Decomposition splits any (adapted) process X into two bits, the ‘easy’ (previsible) bit A and the ‘hard’ (martingale) bit M . Moral: martingales are everywhere!

3. Submartingales model favourable games, so are *increasing on average*. It ‘ought’ to be possible to split such a process into an *increasing* bit, and a

remaining ‘trendless’ bit. It is – the trendless bit is the martingale.

4. This situation resembles that in Statistics, specifically Regression (see e.g. [BF]), where one has a decomposition

$$\text{Data} = \text{Signal} + \text{noise} = \text{fitted value} + \text{residual}.$$

5. The Doob decomposition is linked to the *Riesz decomposition*, which (for us) splits the value of an American option into its intrinsic (European option) part and its ‘early-exercise premium’.

§9. Examples.

1. *Simple random walk.*

Recall the simple random walk: $S_n := \sum_1^n X_k$, where the X_n are independent tosses of a fair coin, taking values ± 1 with equal probability $1/2$. Suppose we decide to bet until our net gain is first $+1$, then quit. Let T be the time we quit; T is a stopping time.

The stopping-time T has been analysed in detail; see e.g.

[GS] GRIMMETT, G. R. & STIRZAKER, D.: *Probability and random processes*, OUP, 3rd ed., 2001 [2nd ed. 1992, 1st ed. 1982], §5.2:

- (i) $T < \infty$ a.s.: the gambler will certainly achieve a net gain of $+1$ eventually;
- (ii) $E[T] = +\infty$: the mean waiting-time till this happens is infinity. So:
- (iii) No bound can be imposed on the gambler’s maximum net loss before his net gain first becomes $+1$.

At first sight, this looks like a foolproof way to make money out of nothing: just bet till you get ahead (which happens eventually, by (i)), then quit. However, as a gambling strategy, this is hopelessly impractical: because of (ii), you need unlimited time, and because of (iii), you need unlimited capital – neither of which is realistic.

Notice that the Optional Stopping Theorem fails here: we start at zero, so $S_0 = 0$, $E[S_0] = 0$; but $S_T = 1$, so $E[S_T] = 1$. This shows two things:

- (a) The Optional Stopping Theorem does indeed need conditions, as the conclusion may fail otherwise [none of the conditions (i) - (iii) in the OST are satisfied in the example above],
- (b) Any practical gambling (or trading) strategy needs to have some integrability or boundedness restrictions to eliminate such theoretically possible but practically ridiculous cases.

2. *The doubling strategy.*

The strategy of doubling when losing – *the martingale*, according to the Oxford English Dictionary (§3) has similar properties – and would be suicidal in practice as a result.

Chapter IV. MATHEMATICAL FINANCE IN DISCRETE TIME.

We follow [BK], Ch. 4.

§1. The Model.

It suffices, to illustrate the ideas, to work with a *finite* probability space (Ω, \mathcal{F}, P) , with a finite number $|\Omega|$ of points ω , each with positive probability: $P(\{\omega\}) > 0$. We will use a finite time-horizon N , which will correspond to the expiry date of the options.

As before, we use a filtration $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \cdots \subset \mathcal{F}_N$: we may (and shall) take $\mathcal{F}_0 = \{\emptyset, \Omega\}$, the trivial σ -field, $\mathcal{F}_N = \mathcal{F} = \mathcal{P}(\Omega)$ (here $\mathcal{P}(\Omega)$ is the *power-set* of Ω , the class of all $2^{|\Omega|}$ subsets of Ω : we need every possible subset, as they all (apart from the empty set) carry positive probability).

The financial market contains $d+1$ financial assets: a riskless asset (bank account) labelled 0, and d risky assets (stocks, say) labelled 1 to d . The prices of the assets at time n are random variables, $S_n^0, S_n^1, \dots, S_n^d$ say [note that we use superscripts here as labels, *not* powers, and suppress ω for brevity], non-negative and \mathcal{F}_n -measurable [at time n , we know the prices S_n^i].

We take $S_0^0 = 1$ (that is, we reckon in units of our initial bank holding). We assume for convenience a constant interest rate $r > 0$ in the bank, so 1 unit in the bank at time 0 grows to $(1+r)^n$ at time n . So $1/(1+r)^n$ is the *discount factor* at time n .

Definition. A *trading strategy* H is a vector stochastic process $H = (H_n)_{n=0}^N = ((H_n^0, H_n^1, \dots, H_n^d))_{n=0}^N$ which is *predictable* (or *previsible*): each H_n^i is \mathcal{F}_{n-1} -measurable for $n \geq 1$.

Here H_n^i denotes the number of shares of asset i held in the portfolio at time n – to be determined on the basis of information available *before* time n ; the vector $H_n = (H_n^0, H_n^1, \dots, H_n^d)$ is the *portfolio* at time n . Writing $S_n = (S_n^0, S_n^1, \dots, S_n^d)$ for the vector of prices at time n , the *value* of the portfolio at time n is the scalar product

$$V_n(H) = H_n \cdot S_n := \sum_{i=0}^d H_n^i S_n^i.$$

The *discounted value* is

$$\tilde{V}_n(H) = \beta_n(H_n \cdot S_n) = H_n \cdot \tilde{S}_n,$$

where $\beta_n := 1/S_n^0$ and $\tilde{S}_n = (1, \beta_n S_n^1, \dots, \beta_n S_n^d)$ is the vector of discounted prices.

Note. The *previsibility* of H reflects that there is no *insider trading*.

Definition. The strategy H is *self-financing*, $H \in SF$, if

$$H_n \cdot S_n = H_{n+1} \cdot S_n \quad (n = 0, 1, \dots, N-1).$$

Interpretation. When new prices S_n are quoted at time n , the investor adjusts his portfolio from H_n to H_{n+1} , without bringing in or consuming any wealth.

Note.

$$\begin{aligned} V_{n+1}(H) - V_n(H) &= H_{n+1} \cdot S_{n+1} - H_n \cdot S_n \\ &= H_{n+1} \cdot (S_{n+1} - S_n) + (H_{n+1} \cdot S_n - H_n \cdot S_n). \end{aligned}$$

For a self-financing strategy, the second term on the right is zero. Then the LHS, the net increase in the value of the portfolio, is shown as due only to the price changes $S_{n+1} - S_n$. So for $H \in SF$,

$$V_n(H) - V_{n-1}(H) = H_n \cdot (S_n - S_{n-1}),$$

$$\Delta V_n(H) = H_n \cdot \Delta S_n, \quad V_n(H) = V_0(H) + \sum_1^n H_j \cdot \Delta S_j$$

and $V_n(H)$ is the *martingale transform* of S by H (III.6). Similarly with discounting:

$$\Delta \tilde{V}_n(H) = H_n \cdot \Delta \tilde{S}_n, \quad \tilde{V}_n(H) = V_0(H) + \sum_1^n H_j \cdot \Delta \tilde{S}_j$$

$$(\Delta \tilde{S}_n := \tilde{S}_n - \tilde{S}_{n-1} = \beta_n S_n - \beta_{n-1} S_{n-1}).$$

As in I, we are allowed to borrow (so S_n^0 may be negative) and sell short (so S_n^i may be negative for $i = 1, \dots, d$). So it is hardly surprising that if we decide what to do about the risky assets, the bank account will take care of itself, in the following sense.

Proposition. If $((H_n^1, \dots, H_n^d))_{n=0}^N$ is predictable and V_0 is \mathcal{F}_0 -measurable, there is a unique predictable process $(H_n^0)_{n=0}^N$ such that $H = (H^0, H^1, \dots, H^d)$ is self-financing with initial value V_0 .

Proof. If H is self-financing, then as above

$$\tilde{V}_n(H) = H_n \cdot \tilde{S}_n = H_n^0 + H_n^1 \tilde{S}_n^1 + \dots + H_n^d \tilde{S}_n^d,$$

while as $\tilde{V}_n = H \cdot \tilde{S}_n$,

$$\tilde{V}_n(H) = V_0 + \sum_1^n (H_j^1 \Delta \tilde{S}_j^1 + \cdots + H_j^d \Delta \tilde{S}_j^d)$$

($\tilde{S}_n = (1, \beta_n S_n^1, \dots, \beta_n S_n^d)$, so $\tilde{S}_n^0 \equiv 1$, $\Delta \tilde{S}_n^0 = 0$). Equate these:

$$H_n^0 = V_0 + \sum_1^n (H_j^1 \Delta \tilde{S}_j^1 + \cdots + H_j^d \Delta \tilde{S}_j^d) - (H_n^1 \tilde{S}_n^1 + \cdots + H_n^d \tilde{S}_n^d),$$

which defines H_n^0 uniquely. The terms in \tilde{S}_n^i are $H_n^i \Delta \tilde{S}_n^i - H_n^i \tilde{S}_n^i = -H_n^i \tilde{S}_{n-1}^i$, which is \mathcal{F}_{n-1} -measurable. So

$$H_n^0 = V_0 + \sum_1^{n-1} (H_j^1 \Delta \tilde{S}_j^1 + \cdots + H_j^d \Delta \tilde{S}_j^d) - (H_n^1 \tilde{S}_{n-1}^1 + \cdots + H_n^d \tilde{S}_{n-1}^d),$$

where as H^1, \dots, H^d are predictable, all terms on the RHS are \mathcal{F}_{n-1} -measurable, so H^0 is predictable. //

Numéraire. What units do we reckon value in? All that is really necessary is that our chosen unit of account should always be *positive* (as we then reckon our holdings by dividing by it, and one cannot divide by zero). Common choices are pounds sterling (UK), dollars (US), euros etc. Gold is also possible (now priced in sterling etc. – but the pound sterling represented an amount of gold, till the UK ‘went off the gold standard’). By contrast, risky stocks *can* have value 0 (if the company goes bankrupt). We call such an always-positive asset, used to reckon values in, a *numéraire*.

Of course, one has to be able to change numéraire – e.g. when going from UK to the US or eurozone. As one would expect, this changes nothing important. In particular, we quote (*numéraire invariance theorem* – see e.g. [BK] Prop. 4.1.1) that the set SF of self-financing strategies is invariant under change of numéraire.

Note. 1. This alerts us to what is meant by ‘risky’. To the owner of a goldmine, sterling is risky. The danger is not that the UK government might go bankrupt, but that sterling might depreciate against the dollar, or euro, etc. 2. With this understood, we shall feel free to refer to our numéraire as ‘bank account’. The point is that we don’t trade in it (why would a goldmine owner trade in gold?); it is the other – ‘risky’ – assets that we trade in. 3. Most large businesses need to trade in more than one currency – sterling (UK), dollars (US), euros (eurozone: check which nations are currently in!), yen, etc. As currencies may and do appreciate or depreciate against each other, questions of foreign exchange – ‘forex’ – are very important and all

around us – and involve large amounts of option trading!

§2. Viability (NA): Existence of Equivalent Martingale Measures.

Although we are allowed to borrow (from the bank), and sell (stocks) short, we are – naturally – required to stay solvent (recall that trading while insolvent is an offence under the Companies Act!).

Definition. A strategy H is *admissible* if it is self-financing, and $V_n(H) \geq 0$ for each time $n = 0, 1, \dots, N$.

Recall that arbitrage is riskless profit – making ‘something out of nothing’. Formally:

Definition. An *arbitrage strategy* is an admissible strategy with zero initial value and positive probability of a positive final value.

Definition. A market is *viable* if no arbitrage is possible, i.e. if the market is arbitrage-free (no-arbitrage, NA).

This leads to the first of two fundamental results:

Theorem (No-arbitrage Theorem: NA iff EMMs exist). The market is viable (is arbitrage-free, is NA) iff there exists a probability measure P^* equivalent to P (i.e., having the same null sets) under which the discounted asset prices are P^* -martingales – that is, iff there exists an equivalent martingale measure (EMM).

Proof. \Leftarrow . Assume such a P^* exists. For any self-financing strategy H , we have as before

$$\tilde{V}_n(H) = V_0(H) + \sum_1^n H_j \cdot \Delta \tilde{S}_j.$$

By the Martingale Transform Lemma, \tilde{S}_j a (vector) P^* -martingale implies $\tilde{V}_n(H)$ is a P^* -martingale. So the initial and final P^* -expectations are the same: using E^* for P^* -expectation,

$$E^*[\tilde{V}_N(H)] = E^*[\tilde{V}_0(H)].$$

If the strategy is admissible and its initial value – the RHS above – is zero, the LHS $E^*[\tilde{V}_N(H)]$ is zero, but $\tilde{V}_N(H) \geq 0$ (by admissibility). Since each $P(\{\omega\}) > 0$ (by assumption), each $P^*(\{\omega\}) > 0$ (by equivalence). This and

$\tilde{V}_N(H) \geq 0$ force $\tilde{V}_N(H) = 0$ (sum of non-negatives can only be 0 if each term is 0). In words: if we start with nothing, we end with nothing – as we should. So no arbitrage is possible. //

The above is the converse part of the NA Theorem – we proved it first as it is easier. The direct part is true, but harder, and needs a preparatory result – which is interesting and important in its own right.

Separating Hyperplane Theorem (SHT).

In a vector space V , a *hyperplane* is a translate of a (vector) subspace U of codimension 1 – that is, U and some one-dimensional subspace, say \mathbb{R} , together span V : V is the direct sum $V = U \oplus \mathbb{R}$ (e.g., $\mathbb{R}^3 = \mathbb{R}^2 \oplus \mathbb{R}$). Then

$$H = [f, \alpha] := \{x : f(x) = \alpha\}$$

for some α and linear functional f . In the finite-dimensional case, of dimension n , say, one can think of $f(x)$ as an inner product,

$$f(x) = f.x = f_1x_1 + \dots + f_nx_n.$$

The hyperplane $H = [f, \alpha]$ *separates* sets $A, B \subset V$ if

$$f(x) \geq \alpha \quad \forall x \in A, \quad f(x) \leq \alpha \quad \forall x \in B$$

(or the same inequalities with A, B , or \geq, \leq , interchanged).

Call a set A in a vector space V *convex* if

$$x, y \in A, \quad 0 \leq \lambda \leq 1 \quad \Rightarrow \quad \lambda x + (1 - \lambda)y \in A$$

– that is, A contains the line-segment joining any pair of its points.

We can now state (without proof) the SHT (see e.g, [BK] App. C).

SHT. Any two non-empty disjoint convex sets in a vector space can be separated by a hyperplane.

A *cone* is a subset of a vector space closed under vector addition and multiplication by *positive* constants (so: like a vector subspace, but with a sign-restriction in scalar multiplication).

We turn now to the proof of the direct part.

Proof of the direct part (not examinable). \Rightarrow : Write Γ for the cone of strictly positive random variables. Viability (NA) says that for any admissible strategy H ,

$$V_0(H) = 0 \quad \Rightarrow \quad \tilde{V}_N(H) \notin \Gamma. \quad (*)$$

To any admissible process (H_n^1, \dots, H_n^d) , we associate its discounted cumulative *gain* process

$$\tilde{G}_n(H) := \sum_1^n (H_j^1 \Delta \tilde{S}_j^1 + \dots + H_j^d \Delta \tilde{S}_j^d).$$

By the Proposition, we can extend (H_1, \dots, H_d) to a unique predictable process (H_n^0) such that the strategy $H = ((H_n^0, H_n^1, \dots, H_n^d))$ is self-financing with initial value zero. By NA, $\tilde{G}_N(H) = 0$ – that is, $\tilde{G}_N(H) \notin \Gamma$.

We now form the set \mathcal{V} of random variables $\tilde{G}_N(H)$, with $H = (H^1, \dots, H^d)$ a previsible process. This is a vector subspace of the vector space \mathbb{R}^Ω of random variables on Ω , by linearity of the gain process $G(H)$ in H . By (*), this subspace \mathcal{V} does not meet Γ . So \mathcal{V} does not meet the subset

$$K := \{X \in \Gamma : \sum_\omega X(\omega) = 1\}.$$

Now K is a convex set not meeting the origin. By the Separating Hyperplane Theorem, there is a vector $\lambda = (\lambda(\omega) : \omega \in \Omega)$ such that for all $X \in K$

$$\lambda.X := \sum_\omega \lambda(\omega) X(\omega) > 0, \tag{1}$$

but for all $\tilde{G}_N(H)$ in \mathcal{V} ,

$$\lambda.\tilde{G}_N(H) = \sum_\omega \lambda(\omega) \tilde{G}_N(H)(\omega) = 0. \tag{2}$$

Choosing each $\omega \in \Omega$ successively and taking X to be 1 on this ω and zero elsewhere, (1) tells us that each $\lambda(\omega) > 0$. So

$$P^*(\{\omega\}) := \lambda(\omega) / (\sum_{\omega' \in \Omega} \lambda(\omega'))$$

defines a probability measure equivalent to P (no non-empty null sets). With E^* as P^* -expectation, (2) says that

$$E^*[\tilde{G}_N(H)] = 0,$$

i.e.

$$E^*[\sum_1^N H_j \cdot \Delta \tilde{S}_j] = 0.$$

In particular, choosing for each i to hold only stock i ,

$$E^*[\sum_1^N H_j^i \Delta \tilde{S}_j^i] = 0 \quad (i = 1, \dots, d).$$

By the Martingale Transform Lemma, this says that the discounted price processes (\tilde{S}_n^i) are P^* -martingales. //