m3f33chIV

Chapter IV: STOCHASTIC PROCESSES IN DISCRETE TIME

§1. Filtrations.

The Kolmogorov triples $(\Omega, \mathcal{F}, P)$, and the Kolmogorov conditional expectations $E(X|\mathcal{B})$, give us all the machinery we need to handle *static* situations involving randomness. To handle *dynamic* situations, involving randomness which unfolds with *time*, we need further structure.

We may take the initial, or starting, time as $t = 0$. Time may evolve discretely, or continuously. We postpone the continuous case to Ch. VI; in the discrete case, we may suppose time evolves in integer steps, $t = 0, 1, 2, \cdots$ (say, stock-market quotations daily, or tick data by the second). There may be a final time T, or *time horizon*, or we may have an infinite time horizon (in the context of option pricing, the time horizon T is the expiry time).

We wish to model a situation involving randomness unfolding with time. We suppose, for simplicity, that information is never lost (or forgotten): thus, as time increases we learn more. Recall that σ-fields represent information or knowledge. We thus need a sequence of σ-fields $\{\mathcal{F}_n : n = 0, 1, 2, \cdots\}$, which are increasing:

$$\mathcal{F}_n \subset \mathcal{F}_{n+1} \qquad (n = 0, 1, 2, \cdots),$$

with $\mathcal{F}_n$ representing the information, or knowledge, available to us at time n. We shall always suppose all σ-fields to be *complete* (this can be avoided, and is not always appropriate, but it simplifies matters and suffices for our purposes). Thus $\mathcal{F}_0$ represents the initial information (if there is none, $\mathcal{F}_0 = \{\emptyset, \Omega\}$, the trivial σ-field). On the other hand,

$$\mathcal{F}_\infty := lim_{n\to\infty}\mathcal{F}_n$$

represents all we ever will know (the 'Doomsday σ-field'). Often, $\mathcal{F}_\infty$ will be $\mathcal{F}$ (the σ-field from Ch. II, representing 'knowing everything'. But this will not always be so; see e.g. [W], §15.8 for an interesting example.

Such a family $\{\mathcal{F}_n : n = 0, 1, 2, \cdots\}$ is called a *filtration*; a probability space endowed with such a filtration, $\{\Omega, \{\mathcal{F}_n\}, \mathcal{F}, P\}$ is called a *filtered probability space*. (These definitions are due to P.- A. MEYER of Strasbourg; Meyer and the Strasbourg (and more generally, French) school of probabilists have been responsible for the 'general theory of [stochastic] processes', and for much of the progress in stochastic integration, since the 1960s.) Since the filtration is so basic to the definition of a stochastic process, the more modern term for a filtered probability space is a *stochastic basis*.

§2. Discrete-Parameter Stochastic Processes.

A *stochastic process* $X = \{X_t : t \in I\}$ is a family of random variables, defined on some common probability space, indexed by an index-set I. Usually (always in this course), I represents *time* (sometimes I represents *space*, and one calls X a spatial process). Here, $I = \{0, 1, 2, \cdots, T\}$ (finite horizon) or $I = \{0, 1, 2, \cdots\}$ (infinite horizon – as in VII.6, Real/Investment options).

The (stochastic) process $X = (X_n)_{n=0}^{\infty}$ is said to be *adapted* to the filtration $(\mathcal{F}_n)_{n=0}^{\infty}$ if

$$X_n \text{ is } \mathcal{F}_n - \text{measurable}.$$

So if X is adapted, we will know the value of X_n at time n. If

$$\mathcal{F}_n = \sigma(X_0, X_1, \cdots, X_n)$$

we call $(\mathcal{F}_n)$ the *natural filtration* of X. Thus a process is always adapted to its natural filtration. A typical situation is that

$$\mathcal{F}_n = \sigma(W_0, W_1, \cdots, W_n)$$

is the natural filtration of some process $W = (W_n)$. Then X is adapted to $(\mathcal{F}_n)$, i.e. each X_n is $\mathcal{F}_n$- (or $\sigma(W_0, \cdots, W_n)$-) measurable, iff

$$X_n = f_n(W_0, W_1, \cdots, W_n)$$

for some measurable function f_n (non-random) of $n + 1$ variables.
Notation.

For a random variable X on $(\Omega, \mathcal{F}, P)$, $X(\omega)$ is the value X takes on ω (ω represents the randomness). Often, to simplify notation, ω is suppressed - e.g., we may write $E[X] := \int_\Omega X dP$ instead of $E[X] := \int_\Omega X(\omega) dP(\omega)$.

For a stochastic process $X = (X_n)$, it is convenient (e.g., if using suffices, n_i say) to use X_n, $X(n)$ interchangeably, and we shall feel free to do this. With ω displayed, these become $X_n(\omega)$, $X(n, \omega)$, etc.

§3. Discrete-Parameter Martingales.

We summarise what we need; for details, see [W], or e.g. [N]

Definition.

A process $X = (X_n)$ is called a *martingale* (mg for short) relative to $((\mathcal{F}_n), P)$ if
(i) X is adapted (to $(\mathcal{F}_n)$),
(ii) $E[|X_n|] < \infty$ for all n,
(iii) $E[X_n|\mathcal{F}_{n-1}] = X_{n-1} \qquad P-a.s. \qquad (n \geq 1)$;
X is a *supermartingale* if in place of (iii)

$$E[X_n|\mathcal{F}_{n-1}] \leq X_{n-1} \qquad P-a.s. \qquad (n \geq 1);$$

X is a *submartingale* if in place of (iii)

$$E[X_n|\mathcal{F}_{n-1}] \geq X_{n-1} \qquad P-a.s. \qquad (n \geq 1).$$

Thus: a mg is 'constant on average', and models a *fair* game;
a supermg is 'decreasing on average', and models an *unfavourable game*;
a submg is 'increasing on average', and models a *favourable* game.

Note. 1. Martingales have many connections with harmonic functions in probabilistic potential theory. The terminology in the inequalities above comes from this: supermartingales correspond to superharmonic functions, submartingales to subharmonic functions.

2. X is a submg [supermg] iff $-X$ is a supermg [submg]; X is a mg iff it is both a submg and a supermg.

3. (X_n) is a mg iff $(X_n - X_0)$ is a mg. So we may without loss of generality take $X_0 = 0$ when convenient.

4. If X is a mg, then for $m < n$

$$\begin{aligned} E[X_n|\mathcal{F}_m] &= E[E(X_n|\mathcal{F}_{n-1})|\mathcal{F}_m] \qquad \text{(iterated conditional expectations)} \\ &= E[X_{n-1}|\mathcal{F}_m] \qquad a.s. \qquad \text{(martingale property)} \\ &= \cdots = E[X_m|\mathcal{F}_m] \qquad a.s. \qquad \text{(induction on } n\text{)}, \\ &= X_m \qquad (X_m \text{ is } \mathcal{F}_m\text{-measurable}) \end{aligned}$$

and similarly for submartingales, supermartingales.

5. Examples of a mg include: sums of independent, integrable zero-mean random variables [submg: positive mean; supermg: negative mean].

From the OED: martingale (etymology unknown)
1. 1589. An article of harness, to control a horse's head.
2. Naut. A rope for guying down the jib-boom to the dolphin-striker.
3. A system of gambling which consists in doubling the stake when losing in order to recoup oneself (1815).
Thackeray: 'You have not played as yet? Do not do so; above all avoid a martingale if you do.'
Problem. Analyse this strategy.

Gambling games have been studied since time immemorial - indeed, the Pascal-Fermat correspondence of 1654 which started the subject was on a problem (de Méré's problem) related to gambling.

The doubling strategy above has been known at least since 1815.

The term 'mg' in our sense is due to J. VILLE (1939). Martingales were studied by Paul LÉVY (1886-1971) from 1934 on [see obituary, *Annals of Probability* **1** (1973), 5-6] and by J. L. DOOB (1910-2004) from 1940 on. The first systematic exposition was Doob's book [D], Ch. VII.
Example: Accumulating data about a random variable ([W], 96, 166-167).
If $\xi \in L_1(\Omega, \mathcal{F}, P)$, $M_n := E(\xi|\mathcal{F}_n)$ (so M_n represents our best estimate of ξ based on knowledge at time n), then

$$\begin{aligned} E[M_n|\mathcal{F}_{n-1}] &= E[E(\xi|\mathcal{F}_n)|\mathcal{F}_{n-1}] \\ &= E[\xi|\mathcal{F}_{n-1}] \qquad \text{(iterated conditional expectations)} \\ &= M_{n-1}, \end{aligned}$$

so (M_n) is a mg. One has the convergence (see IV.4 below)

$$M_n \to M_\infty := E[\xi|\mathcal{F}_\infty] \qquad a.s. \qquad \text{and in } L_1.$$

§4. Martingale Convergence.

A supermartingale is 'decreasing on average'. Recall that a decreasing sequence [of real numbers] that is bounded below converges (decreases to its greatest lower bound or infimum). This suggests that a supermartingale which is bounded below converges a.s. This is so [Doob's Forward Convergence Theorem: [W], §§11.5, 11.7].

More is true. Call X L_1-bounded if

$$\sup_n E[|X_n|] < \infty.$$

Theorem (Doob). An L_1-bounded supermartingale is a.s. convergent: there exists X_∞ finite such that

$$X_n \to X_\infty \qquad (n \to \infty) \qquad a.s.$$

In particular, we have

Doob's Martingale Convergence Theorem [W, §11.5]. An L_1-bounded martingale converges a.s.

We say that

$$X_n \to X_\infty \text{ in } L_1$$

if

$$E[\ |X_n - X_\infty|\ ] \to 0 \qquad (n \to \infty).$$

For a class of martingales, one gets convergence in L_1 as well as almost surely [= with probability one]. Such martingales are called *uniformly integrable* (UI) [W], or *regular* [N], or *closed* (see below), They are "the nice ones". Fortunately, they are the ones we need.

The following result is in [N], IV.2, [W], Ch. 14; cf. SP L18-19, SA L6.

Theorem (UI Martingale Convergence Theorem). The following are equivalent for martingales $X = (X_n)$:
(i) X_n converges in L_1,
(ii) X_n is L_1-bounded, and its a.s. limit X_∞ (which exists, by above) satisfies

$$X_n = E[X_\infty | \mathcal{F}_n],$$

(iii) There exists an integrable random variable X with

$$X_n = E[X | \mathcal{F}_n].$$

The random variable X_∞ above serves to "close" the martingale, by giving X_n a value at "$n = \infty$"; then $\{X_n : n = 1, 2, \ldots, \infty\}$ is again a martingale – which we may accordingly call a closed mg. The terms closed, regular and UI are used interchangeably here.

Notice that *all the randomness in a closed mg is in the closing value* X_∞ (so, although a stochastic process is an infinite-dimensional object, the randomness in a closed mg is one-dimensional). As time progresses, more is revealed, by "progressive revelation" – as in (choose your metaphor) a striptease, or the "Day of Judgement" (when all will be revealed).

As we shall see (Risk-Neutral Valuation Formula): closed mgs are vital in mathematical finance, and the closing value corresponds to the payoff of an option.

§5. Martingale Transforms.

Now think of a gambling game, or series of speculative investments, in discrete time. There is no play at time 0; there are plays at times $n = 1, 2, \cdots$, and

$$\Delta X_n := X_n - X_{n-1}$$

represents our net winnings per unit stake at play n. Thus if X_n is a martingale, the game is 'fair on average'.

Call a process $C = (C_n)_{n=1}^{\infty}$ *previsible* (or *predictable*) if

$$C_n \text{ is } \mathcal{F}_{n-1} - \text{measurable for all } n \geq 1.$$

Think of C_n as your stake on play n (C_0 is not defined, as there is no play at time 0). Previsibility says that you have to decide how much to stake on play n based on the history *before* time n (i.e., up to and including play $n-1$). Your winnings on game n are $C_n \Delta X_n = C_n(X_n - X_{n-1})$. Your total (net) winnings up to time n are

$$Y_n = \sum_{k=1}^{n} C_k \Delta X_k = \sum_{k=1}^{n} C_k (X_k - X_{k-1}).$$

We write

$$Y = C \bullet X, \qquad Y_n = (C \bullet X)_n, \qquad \Delta Y_n = C_n \Delta X_n$$

($(C \bullet X)_0 = 0$ as $\sum_1^0$ is empty), and call $C \bullet X$ the *martingale transform* of X by C.

Theorem. (i) If C is a bounded non-negative previsible process and X is a supermartingale, $C \bullet X$ is a supermartingale null at zero.
(ii) If C is bounded and previsible and X is a martingale, $C \bullet X$ is a martingale null at zero.

Proof.

With $Y = C \bullet X$ as above,

$$E[Y_n - Y_{n-1}|\mathcal{F}_{n-1}] = E[C_n(X_n - X_{n-1})|\mathcal{F}_{n-1}]$$

$$= C_n E[(X_n - X_{n-1})|\mathcal{F}_{n-1}]$$

(as C_n is bounded, so integrable, and $\mathcal{F}_{n-1}$-measurable, so can be taken out)

$$\leq 0$$

in case (i), as $C \geq 0$ and X is a supermartingale,

$$= 0$$

in case (ii), as X is a martingale. //

Interpretation. You can't beat the system!
In the martingale case, previsibility of C means we can't foresee the future (which is realistic and fair). So we expect to gain nothing – as we should.
Note. 1. Martingale transforms were introduced and studied by Donald L. BURKHOLDER (1927 - 2013) in 1966 [*Ann. Math. Statist.* **37**, 1494-1504]. For a textbook account, see e.g. [N], VIII.4.
2. Martingale transforms are the discrete analogues of stochastic integrals. They dominate the mathematical theory of finance in discrete time, just as stochastic integrals dominate the theory in continuous time.
3. In mathematical finance, X plays the role of a price process, C plays the role of our trading strategy, and the mg transform $C \bullet X$ plays the role of our gains (or losses!) from trading.

Proposition (Martingale Transform Lemma). An adapted sequence of real integrable random variables (M_n) is a martingale iff for any bounded previsible sequence (H_n),

$$E[\sum_{r=1}^{n} H_r \Delta M_r] = 0 \qquad (n = 1, 2, \cdots).$$

Proof.

If (M_n) is a martingale, X defined by $X_0 = 0$, $X_n = \sum_1^n H_r \Delta M_r \qquad (n \geq 1)$ is the martingale transform $H \bullet M$, so is a martingale.

Conversely, if the condition of the Proposition holds, choose j, and for any $\mathcal{F}_j$-measurable set A write $H_n = 0$ for $n \neq j+1$, $H_{j+1} = I_A$. Then (H_n) is previsible, so the condition of the Proposition, $E[\sum_1^n H_r \Delta M_r] = 0$, becomes

$$E[I_A(M_{j+1} - M_j)] = 0.$$

As this holds for every $A \in \mathcal{F}_j$, the definition of conditional expectation gives

$$E[M_{j+1}|\mathcal{F}_j] = M_j.$$

Since this holds for every j, (M_n) is a martingale. //

§6. Stopping Times and Optional Stopping.

A random variable T taking values in $\{0, 1, 2, \cdots; +\infty\}$ is called a *stopping time* (or *optional time*) if

$$\{T \leq n\} = \{\omega : T(\omega) \leq n\} \in \mathcal{F}_n \qquad \forall n \leq \infty.$$

Equivalently,

$$\{T = n\} \in \mathcal{F}_n \qquad n \leq \infty.$$

Think of T as a time at which you decide to quit a gambling game: whether or not you quit at time n depends only on the history up to and including time n – NOT the future. [Elsewhere, T denotes the expiry time of an option. If we mean T to be a stopping time, we will say so.]

The following important classical theorem is discussed in [W], 10.10.

Theorem (Doob's Optional Stopping Theorem, OST). Let T be a stopping time, $X = (X_n)$ be a supermartingale, and assume that one of the following holds:
(i) T is bounded [$T(\omega) \leq K$ for some constant K and all $\omega \in \Omega$];
(ii) $X = (X_n)$ is bounded [$|X_n(\omega)| \leq K$ for some K and all n, ω];
(iii) $E[T] < \infty$ and $(X_n - X_{n-1})$ is bounded.

Then X_T is integrable, and

$$E[X_T] \leq E[X_0].$$

If here X is a martingale, then

$$E[X_T] = E[X_0].$$

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

$$X_{T\wedge n} = X_0 + \sum_1^n \phi_j (X_j - X_{j-1}).$$

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)

$$E[X_{T\wedge n}|\mathcal{F}_{n-1}] = X_0 + \sum_1^{n-1} \phi_j (X_j - X_{j-1}) + \phi_n (E[X_n|\mathcal{F}_{n-1}] - X_{n-1}), \qquad \text{i.e.}$$

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

$\phi_n \geq 0$ 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 *backward recursion* by

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

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.

The Snell envelope (see IV.8 L20) 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

$$T_{n-1} \geq E[T_n|\mathcal{F}_{n-1}] \qquad \text{(as } T \text{ is a supermartingale)}$$

$$\geq E[U_n|\mathcal{F}_{n-1}] \qquad \text{(by the induction hypothesis)}$$

and

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

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.

We omit the proof (not hard, but fiddly – for details, see e.g. L13, 2014). 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 $\{n, n+1, \cdots, N\}$ (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),

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

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

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

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

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

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

The proof in discrete time is quite easy (see L13, 2014). It is hard in continuous time – but more important there (see Ch. V: quadratic variation (QV) and the Itô integral). This illustrates the contrasts between the theories of stochastic processes in discrete and continuous time.

§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 $1/2$. We decide to bet until our net gain is first $+1$, then quit – at time T, a stopping time. This 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. But as a gambling strategy, this is hopelessly impractical: because of (ii), you need unlimited time, and because of (iii), you need unlimited capital!

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.* Similarly for the doubling strategy (§3).