Background Notes

R. J. Williams, Copyright 2005

October 3, 2005

Contents

Appendix A

Conditional Expectation and L^p-Spaces

In this appendix (A), we briefly review the definition and some basic properties of conditional expectation. We also establish some notation for L^p spaces. For more details we refer the reader to a graduate text in probability such as Chung [9] or D. Williams [38].

Consider a fixed probability space $(\Omega, \mathcal{F}, P)$, where Ω represents a sample space of possible outcomes, $\mathcal{F}$ is a σ-algebra of subsets of Ω representing the events to which probabilities can be assigned, and P is probability measure on $(\Omega, \mathcal{F})$. Expectation under P will be denoted by $E[\,\cdot\,]$. The term a.s. will mean almost surely with respect to the probability measure P. Given a random variable X defined on this space satisfying $E[|X|] < \infty$, and a sub-σ-algebra $\mathcal{G} \subset \mathcal{F}$, the *conditional expectation* of X given $\mathcal{G}$ is denoted by $E[X|\mathcal{G}]$. This represents a member of the equivalence class of $\mathcal{G}$-measurable random variables satisfying

$$E[E[X|\mathcal{G}]1_A] = E[X1_A] \quad \text{for all } A \in \mathcal{G}. \tag{A.1}$$

Because conditional expectations are only determined uniquely up to almost sure equivalence, (in)equalities between conditional expectations are always interpreted to hold almost surely. If Y is a random variable defined on the probability space and $\sigma(Y)$ denotes the σ-algebra generated by Y, then the conditional expectation $E[X|\sigma(Y)]$ is often simply denoted by $E[X|Y]$. We list several common properties of conditional expectation here. In the following, X and Y are integrable random variables defined on $(\Omega, \mathcal{F}, P)$, and $\mathcal{G}, \mathcal{H}$ are sub-σ-algebras of $\mathcal{F}$.

(i) If X is $\mathcal{G}$-measurable, then $E[X|\mathcal{G}] = X$.

(ii) If X is independent of $\mathcal{G}$, then $E[X|\mathcal{G}] = E[X]$.

(iii) If Z is $\mathcal{G}$-measurable and XZ is integrable (in addition to X), then

$$E[XZ|\mathcal{G}] = ZE[X|\mathcal{G}].$$

(iv) For $c \in \mathbb{R}$,

$$E[cX + Y|\mathcal{G}] = cE[X|\mathcal{G}] + E[Y|\mathcal{G}].$$

(v) If $X \leq Y$ a.s., then

$$E[X|\mathcal{G}] \leq E[Y|\mathcal{G}].$$

(vi) $|E[X|\mathcal{G}]| \leq E[|X||\mathcal{G}]$.

(vii) (Jensen's inequality) If $\phi : \mathbb{R} \to \mathbb{R}$ is convex and $\phi(X)$ is integrable (in addition to X), then

$$\phi(E[X|\mathcal{G}]) \leq E[\phi(X)|\mathcal{G}].$$

(viii) (Tower property) If $\mathcal{G} \subset \mathcal{H}$, then

$$E[E[X|\mathcal{H}]|\mathcal{G}] = E[E[X|\mathcal{G}]|\mathcal{H}] = E[X|\mathcal{G}].$$

(ix) (Monotone convergence) If, for each $n = 1, 2, \ldots$, X_n is a random variable such that $0 \leq X_n \leq X$ a.s., and $X_n \uparrow X$ a.s. as $n \to \infty$ (where X is integrable), then

$$E[X_n|\mathcal{G}] \uparrow E[X|\mathcal{G}] \quad \text{a.s. as } n \to \infty.$$

(x) (Dominated convergence) If, for each $n = 1, 2, \ldots$, X_n is a random variable such that $|X_n| \leq Y$ a.s., and $X_n \to X$ a.s. as $n \to \infty$ (where X, Y are integrable), then

$$E[X_n|\mathcal{G}] \to E[X|\mathcal{G}] \quad \text{a.s. as } n \to \infty.$$

The following special case is relevant to the probability models treated in Chapters 2 and 3. When the sample space Ω is a finite set, a sub-σ-algebra $\mathcal{G}$ of $\mathcal{F}$ can be specified by giving a finite partition $\mathcal{P} = \{A_1, \ldots, A_n\}$ of Ω that generates $\mathcal{G}$, i.e., A_i, $i = 1, \ldots, n$ are disjoint sets in $\mathcal{F}$, $\cup_{i=1}^n A_i = \Omega$ and $\mathcal{G} = \sigma(\mathcal{P})$, the smallest σ-algebra containing $\mathcal{P}$. In this case the random variable $Z = E[X|\mathcal{G}]$ must be constant on each of the sets in the partition $\mathcal{P}$ and its value on any set A_i of positive probability is given by

$$E[X 1_{A_i}]/P(A_i).$$

For any σ-algebra $\mathcal{G} \subset \mathcal{F}$ and $p \in [1, \infty)$, $L^p(\Omega, \mathcal{G}, P)$ will denote the space of random variables $X : \Omega \to \mathbb{R}$ that are $\mathcal{G}$-measurable, i.e., $X^{-1}(B) \equiv \{\omega \in \Omega : X(\omega) \in B\} \in \mathcal{G}$ for each Borel set B in $\mathbb{R}$, and that satisfy $E[|X|^p] < \infty$. The norm of $X \in L^p(\Omega, \mathcal{G}, P)$ is given by

$$\|X\|_p = E[|X|^p]^{1/p}.$$

Furthermore, $L^\infty(\Omega, \mathcal{G}, P)$ denotes the space of $\mathcal{G}$-measurable random variables $X : \Omega \to \mathbb{R}$ such that X is bounded P-a.s. The norm on this space is the so-called essential supremum norm given for $X \in L^\infty(\Omega, \mathcal{G}, P)$ by

$$\|X\|_\infty = \inf\{a \geq 0 : P(|X| > a) = 0\}.$$

The following interpretation of $E[X|\mathcal{G}]$ is useful when X is square integrable. If the random variable X is such that $E[X^2] < \infty$, then the conditional expectation $Z = E[X|\mathcal{G}]$ is a version of the orthogonal projection of X onto the L^2 space $L^2(\Omega, \mathcal{G}, P)$, i.e., Z is the $\mathcal{G}$-measurable, square integrable random variable (unique up to a.s. equivalence) that minimizes the squared distance

$$E[(Z - X)^2].$$

Appendix B

Discrete Time Stochastic Processes

In this appendix (B), we briefly review some concepts and results for the types of discrete time stochastic processes used in this book. For more details and further background we refer the reader to a graduate text in probability such as Chung [9] or D. Williams [38].

Only discrete time stochastic processes indexed by a finite set of times are considered in this appendix, and so, without loss of generality, here we assume that the time index set is the set of consecutive non-negative integers $\{0, 1, \ldots, T\}$ for some finite non-negative integer T.

Throughout this appendix (B), we assume that $(\Omega, \mathcal{F}, P)$ is a fixed probability space, where Ω is a sample space representing the set of possible outcomes, $\mathcal{F}$ is a σ-algebra of subsets of Ω representing the events to which we can assign probabilities, and P is a probability measure on $(\Omega, \mathcal{F})$. The expectation with respect to P will be denoted by $E[\,\cdot\,]$.

A *filtration* is a family $\{\mathcal{F}_t, t = 0, 1, \ldots, T\}$ of sub-σ-algebras of $\mathcal{F}$ indexed by $t = 0, 1, \ldots, T$ such that

$$\mathcal{F}_0 \subset \mathcal{F}_1 \subset \mathcal{F}_2 \ldots \subset \mathcal{F}_T,$$

i.e., the family is increasing with time. If the sample space Ω is a finite set, often the σ-algebra $\mathcal{F}_0$ is trivial, consisting simply of the empty set $\emptyset$ and the whole sample space Ω. Also, since we shall only be considering random variables that are $\mathcal{F}_T$-measurable, without loss of generality (by redefining $\mathcal{F}$ to equal $\mathcal{F}_T$), we can assume that $\mathcal{F}_T = \mathcal{F}$. We often write $\{\mathcal{F}_t\}$ instead of the more cumbersome $\{\mathcal{F}_t, t = 0, 1, \ldots, T\}$. Intuitively, the filtration keeps track of what information is known at each of the times $t = 0, 1, \ldots, T$, where information only increases with time. More precisely, for each $t = 0, 1, \ldots, T$, the σ-algebra $\mathcal{F}_t$ tells us which events may be observed by time t. We call the quadruple $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ a filtered probability space.

For the remainder of this appendix (B) we suppose that in addition to

$(\Omega, \mathcal{F}, P)$ being fixed, $\{\mathcal{F}_t\}$ is a fixed filtration, and so $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ is a given filtered probability space. All random variables and processes considered in this appendix are assumed to be defined on this space.

Given a positive integer d, a *d-dimensional (stochastic) process* with time index set $\{0, 1, \dots, T\}$, defined on the given filtered probability space, is a collection $X = \{X_t, t = 0, 1, \dots, T\}$ where each X_t is a d-dimensional random vector, i.e., a function $X_t : \Omega \to \mathbb{R}^d$ such that $X_t^{-1}(B) \equiv \{\omega \in \Omega : X_t(\omega) \in B\} \in \mathcal{F}$ for each Borel subset B of $\mathbb{R}^d$. The process X is called *adapted* if $X_t^{-1}(B) \in \mathcal{F}_t$ for each Borel set B in $\mathbb{R}^d$ and for each $t = 0, 1, \dots, T$. We often write $X_t \in \mathcal{F}_t$, as shorthand for $X_t^{-1}(B) \in \mathcal{F}_t$ for all Borel sets B in $\mathbb{R}^d$.

Two d-dimensional processes $\{Y_t, t = 0, 1, \dots, T\}$ and $\{Z_t, t = 0, 1, \dots, T\}$ defined on $(\Omega, \mathcal{F}, P)$ are *modifications of one another* if $P(Y_t = Z_t) = 1$ for $t = 0, 1, \dots, T$. Since the time index set is finite, this is equivalent to Y and Z being *indistinguishable,* i.e., $P(Y_t = Z_t \text{ for } t = 0, 1, \dots, T) = 1$. We shall regard two such indistinguishable processes as being equal as stochastic processes.

A (discrete) *stopping time* (or optional time) is a function $\tau : \Omega \to \{0, 1, \dots, T\} \cup \{\infty\}$ such that

$$\{\tau = t\} \in \mathcal{F}_t \quad \text{for} \quad t = 0, 1, \dots, T. \tag{B.1}$$

Note that for such a stopping time τ, we have

$$\{\tau = \infty\} = \Omega \setminus (\cup_{t=0}^{T} \{\tau = t\}) \in \mathcal{F}_T.$$

For convenience we define $\mathcal{F}_\infty = \mathcal{F}_T$ and then (B.1) also holds with $t = \infty$. A deterministic time in $\{0, 1, \dots, T\} \cup \{\infty\}$ is a stopping time. Furthermore, if τ and σ are two stopping times, then it is straightforward to verify that $\tau \wedge \sigma = \min(\tau, \sigma)$ and $\tau \vee \sigma = \max(\tau, \sigma)$ are also stopping times. There is a σ-algebra associated with any (discrete) stopping time τ, defined by

$$\mathcal{F}_\tau = \{A \in \mathcal{F}_\infty : A \cap \{\tau = t\} \in \mathcal{F}_t \ \text{ for } \ t = 0, 1, \dots, T\}.$$

Note that, for $A \in \mathcal{F}_\tau$, $A \cap \{\tau = \infty\} \in \mathcal{F}_\infty$.

A collection $M = \{M_t, \mathcal{F}_t, t = 0, 1, \dots, T\}$, where each M_t is a real-valued random variable, is called a *martingale* if the following three properties hold:

(i) $E[|M_t|] < \infty$ for $t = 0, 1, \dots, T$,

(ii) M_t is $\mathcal{F}_t$-measurable for $t = 0, 1, \dots, T$,

(iii) $E[M_t | \mathcal{F}_{t-1}] = M_{t-1}$ for $t = 1, \dots, T$.

In this discrete time setting, the condition (iii) can be equivalently replaced by

(iii)' $E[M_t | \mathcal{F}_s] = M_s$ for all $s < t$ in $\{0, 1, \dots, T\}$.

We call M a *submartingale* if the "$=$" in (iii) (or (iii)') is replaced by "$\geq$", and we call M a *supermartingale* if the "$=$" in (iii) (or (iii)') is replaced by "$\leq$". The condition (iii) (or (iii)') is often referred to as the (sub/super) martingale property. Note that M is a submartingale if and only if $-M$ is a supermartingale.

If M is a (sub/super) martingale and in addition there is $p \in (1, \infty)$ such that $M_t \in L^p(\Omega, \mathcal{F}, P)$ for each $t = 0, 1, \ldots, T$, then we call M an L^p-(sub/super) martingale. In describing (sub/super) martingales, we shall sometimes omit the filtration $\{\mathcal{F}_t\}$ from the notation for M when it is understood.

We adopt the convention that a d-dimensional process $M = \{M_t, \mathcal{F}_t, t = 0, 1, \ldots, T\}$ is called a (sub/super) martingale if and only if each of its one-dimensional components $\{M_t^i, \mathcal{F}_t, t = 0, 1, \ldots, T\}$, $i = 1, \ldots, d$, is a (sub/super) martingale.

Two basic results from discrete time martingale theory are stated here for ease of reference. We state them for one-dimensional processes.

Theorem B.0.1 *(Doob's L^p-inequality)*
For $p \in (1, \infty)$, let $M = \{M_t, \mathcal{F}_t, t = 0, 1, \ldots, T\}$ be a real-valued L^p-martingale. Then, for each $t = 0, 1, \ldots, T$, $N_t \equiv \sup\{|M_s|, s = 0, 1, \ldots, t\} \in L^p(\Omega, \mathcal{F}_t, P)$ and

$$E\left[(N_t)^p\right] \leq q^p E\left[|M_t|^p\right], \tag{B.2}$$

where $\frac{1}{p} + \frac{1}{q} = 1$.

Theorem B.0.2 *(Doob's stopping theorem)*
Let $M = \{M_t, \mathcal{F}_t, t = 0, 1, \ldots, T\}$ be a real-valued (sub)martingale and suppose that τ is a stopping time. Then $\{M_{t\wedge\tau}, \mathcal{F}_{t\wedge\tau}, t = 0, 1, \ldots, T\}$ is a (sub)martingale. Furthermore, if σ is another stopping time and $\sigma \leq \tau \leq T$, then M_τ, M_σ are integrable random variables and

$$E[M_\tau] \geq E[M_\sigma], \tag{B.3}$$

where equality holds if M is a martingale.

We shall use the last result above in particular with $\sigma = 0$.

Appendix C

Continuous Time Stochastic Processes

In this appendix (C), we briefly review some concepts and results from the theory of continuous time stochastic processes. Many of these concepts and results are similar to those introduced for discrete time stochastic processes in the preceding appendix. Some notable exceptions are the definitions of stopping time and martingale, and the fact that filtrations and stochastic processes in continuous time are frequently assumed to have some regularity, e.g., right continuity, due to the fact that the time index set is now a continuum. In this appendix (C), we restrict to compact time intervals.

In the following, we consider the compact time interval $[0, T]$ where T is a fixed value in $[0, \infty)$. We denote the Borel σ-algebra on $[0, T]$ by $\mathcal{B}_T$.

A *filtered probability space* is a quadruple $(\Omega, \mathcal{F}, \{\mathcal{F}_t, t \in [0, T]\}, P)$, where Ω is a set representing the sample space, $\mathcal{F}$ is a σ-algebra of subsets of Ω, P is a probability measure on $(\Omega, \mathcal{F})$, and $\{\mathcal{F}_t, t \in [0, T]\}$ is a filtration, i.e., a family of sub-σ-algebras of $\mathcal{F}$ indexed by $t \in [0, T]$ that is increasing: for each $s, t \in [0, T]$ satisfying $s < t$ we have $\mathcal{F}_s \subset \mathcal{F}_t$. We often write $\{\mathcal{F}_t\}$ instead of the more cumbersome $\{\mathcal{F}_t, t \in [0, T]\}$. Intuitively, we may regard $\mathcal{F}$ as containing all events which might ever be observed or to which we can assign probabilities. The filtration keeps track of what information is known by each of the times $t \in [0, T]$, where information only increases with time. More precisely, for each $t \in [0, T]$, the σ-algebra $\mathcal{F}_t$ tells us which events might be observed by time t. For such a filtered probability space, a set $A \in \mathcal{F}$ whose probability is zero is called a *P-null set.* The expectation with respect to P will be denoted by $E[\,\cdot\,]$. Since we shall only be considering random variables that are $\mathcal{F}_T$-measurable, without loss of generality (by redefining $\mathcal{F}$ to equal $\mathcal{F}_T$), we can assume that $\mathcal{F}_T = \mathcal{F}$.

In continuous time, some additional regularity assumptions are typically imposed on a filtered probability space. In particular, a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ is said to satisfy the *usual conditions* if the following three conditions hold:

(i) $(\Omega, \mathcal{F}, P)$ is complete, i.e., if $A \subset B$ where $B \in \mathcal{F}$ and $P(B) = 0$, then $A \in \mathcal{F}$ and $P(A) = 0$,

(ii) $\mathcal{F}_0$ contains all of the P-null sets,

(iii) $\{\mathcal{F}_t\}$ is right continuous, i.e., $\mathcal{F}_t = \mathcal{F}_{t+} \equiv \bigcap_{s \in (t,T]} \mathcal{F}_s$, for all $t < T$.

Remark. If we have a filtered probability space that does not satisfy the usual conditions, by a standard procedure of completing the probability space, augmenting the members of the filtration using all of the P-null sets, and replacing $\mathcal{F}_t$ by $\mathcal{F}_{t+}$ for $t < T$, we can ensure that the usual conditions are satisfied (cf. Chung [9], page 29ff.).

For the remainder of this appendix, we suppose that $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ is a given filtered probability space satisfying the usual conditions. All random variables and processes considered in this section are assumed to be defined on this space.

Given a positive integer d, a *d-dimensional (stochastic) process* on the given filtered probability space is a collection $X = \{X_t, t \in [0, T]\}$ where each X_t is a d-dimensional random vector, i.e., a function $X_t : \Omega \to \mathbb{R}^d$ such that $X_t^{-1}(B) \equiv \{\omega \in \Omega : X_t(\omega) \in B\} \in \mathcal{F}$ for each Borel subset B of $\mathbb{R}^d$. (In fact, it is enough to check that this last property holds for all open balls B in $\mathbb{R}^d$.) The process X is called *adapted* if $X_t^{-1}(B) \in \mathcal{F}_t$ for each Borel set B in $\mathbb{R}^d$ and for all $t \in [0, T]$. We often write $X_t \in \mathcal{F}_t$, as shorthand for $X_t^{-1}(B) \in \mathcal{F}_t$ for all Borel sets B in $\mathbb{R}^d$. By setting $X(t, \omega) = X_t(\omega)$ for all $t \in [0, T]$, $\omega \in \Omega$, we may sometimes view X as a function from $[0, T] \times \Omega$ into $\mathbb{R}^d$. The process X is said to be (left/right) continuous if all of its sample paths are (left/right) continuous. (Some authors only require that (left/right) continuous processes have (left/right) continuous paths almost surely. Frequently a process with such a property can be redefined on a P-null set to have (left/right) continuous paths surely. For most practical purposes such a redefined process can be considered to be the same as the original process. Accordingly, we have chosen to require (left/right) continuous processes to have the desired path regularity surely.)

Two d-dimensional processes $\{Y_t, t \in [0, T]\}$ and $\{Z_t, t \in [0, T]\}$ defined on $(\Omega, \mathcal{F}, P)$ are

(i) *modifications of one another* if $P(Y_t = Z_t) = 1$ for all $t \in [0, T]$,

(ii) *indistinguishable* if $P(Y_t = Z_t \text{ for all } t \in [0, T]) = 1$.

(Note that for (ii) it is implicit that the event $\{Y_t = Z_t \text{ for all } t \in [0, T]\}$ is measurable.) Two right continuous processes that are modifications of one another are indistinguishable and we shall regard them as being equal as stochastic processes.

A *stopping time* (or optional time) is a function $\tau : \Omega \to [0, T] \cup \{\infty\}$ such that

$$\{\tau \leq t\} \in \mathcal{F}_t \quad \text{for all} \quad t \in [0, T]. \tag{C.1}$$

Note that for such a stopping time τ, we have $\{\tau = \infty\} = \Omega \setminus \{\tau \leq T\} \in \mathcal{F}_T$, and so we define $\mathcal{F}_\infty = \mathcal{F}_T$. Note that then (C.1) also holds with $t = \infty$. There is a σ-algebra associated with any stopping time τ, defined by

$$\mathcal{F}_\tau = \{A \in \mathcal{F}_\infty : A \cap \{\tau \leq t\} \in \mathcal{F}_t \text{ for all } t \in [0,T]\}\}.$$

It is easily verified that a stopping time τ is measurable as a mapping from $(\Omega, \mathcal{F}_\tau)$ into $([0,T] \cup \{\infty\}, \mathcal{B}_\infty)$ where $\mathcal{B}_\infty$ is the σ-algebra on $[0,T] \cup \{\infty\}$ generated by the Borel sets in $[0,T]$ together with the singleton set $\{\infty\}$. If τ and σ are two stopping times, then it is straightforward to verify that $\tau \wedge \sigma = \min(\tau, \sigma)$ and $\tau \vee \sigma = \max(\tau, \sigma)$ are also stopping times.

Example. If X is a continuous, adapted d-dimensional process, then for any open or closed set $A \subset \mathbb{R}^d$, $\tau_A = \inf\{t \in [0,T] : X_t \in A\}$ is a stopping time. (As usual, the infimum of the empty set is defined to equal ∞.) For a proof of this, which uses the fact that the filtration satisfies the usual conditions, see Chung [9], Section 2.4.

A collection $M = \{M_t, \mathcal{F}_t, t \in [0,T]\}$, where each M_t is a real-valued random variable, is called a *martingale* if the following three properties hold:

(i) $E[|M_t|] < \infty$ for all $t \in [0,T]$,

(ii) M_t is $\mathcal{F}_t$-measurable for each $t \in [0,T]$,

(iii) $E[M_t|\mathcal{F}_s] = M_s$ for all $s \leq t$ in $[0,T]$.

We call M a *submartingale* if the "=" in (iii) is replaced by "$\geq$", and we call M a *supermartingale* if the "=" in (iii) is replaced by "$\leq$". This condition (iii) is referred to as the (sub/super) martingale property whenever M is a (sub/super) martingale. If M is a (sub/super) martingale and in addition there is $p \in (1, \infty)$ such that $M_t \in L^p(\Omega, \mathcal{F}, P)$ for all $t \in [0,T]$, then we call M an L^p-(sub/super) martingale.

It is well known that under the usual conditions that we have assumed to hold for the filtered probability space, every martingale has a modification whose paths are all right continuous with finite left limits, cf. Chung [9], Theorem 3, page 29 and Corollary 1, page 26.

The stochastic integrals that we consider will often define martingales; however, sometimes they will only locally define martingales. This leads us to the notion of a local martingale. A collection $M = \{M_t, \mathcal{F}_t, t \in [0,T]\}$, where each M_t is a real-valued random variable, is a called a *local martingale* if M is an adapted process and there is a sequence of stopping times $\{\tau_k\}_{k=1}^\infty$ such that $\tau_k \leq \tau_{k+1}$ for each k, $\lim_{k\to\infty} \tau_k = \infty$ P-a.s., and for each k,

$$M^k = \{M_{t \wedge \tau_k}, \mathcal{F}_t, t \in [0,T]\}$$

is a martingale. We call such a sequence $\{\tau_k\}$ a localizing sequence for M. Here we have used a slightly stronger notion of local martingale than that introduced in [11] and [35]. In the case when the initial value M_0 is a constant, our definition is equivalent to that in [11] and [35]. As that is the only circumstance considered here, we have chosen to give the stronger and simpler definition. Indeed, for stochastic integrals, the initial value is not only constant but zero.

In describing (sub/super/local) martingales, we shall sometimes omit the filtration $\{\mathcal{F}_t\}$ from the notation for M when it is understood. We also adopt the convention that a d-dimensional process $M = \{M_t, t \in [0,T]\}$ is called a (sub/super/local) martingale if and only if each of its one-dimensional components $\{M_t^i, t \in [0,T]\}$, $i = 1, \ldots, d$, is a (sub/super/local) martingale. In the case of a local martingale, one can choose a common localizing sequence for all components, since the minimum of finitely many stopping times is again a stopping time. We shall always assume that such a common sequence is being used when we consider multi-dimensional local martingales.

Two basic results from continuous time martingale theory are stated here for ease of reference. We state them for one-dimensional processes.

Theorem C.0.3 *(Doob's L^p-inequality)*
For $p \in (1,\infty)$, let $M = \{M_t, \mathcal{F}_t, t \in [0,T]\}$ be a real-valued right continuous L^p-martingale. Then, for each $t \in [0,T]$, $N_t \equiv \sup_{0 \le s \le t} |M_s| \in L^p(\Omega, \mathcal{F}_t, P)$ and

$$E\left[(N_t)^p\right] \le q^p E\left[|M_t|^p\right], \tag{C.2}$$

where $\frac{1}{p} + \frac{1}{q} = 1$.

Theorem C.0.4 *(Doob's Stopping Theorem)*
Let $M = \{M_t, \mathcal{F}_t, t \in [0,T]\}$ be a real-valued right continuous (sub)martingale and suppose that τ is a stopping time. Then $\{M_{t\wedge\tau}, \mathcal{F}_{t\wedge\tau}, t \in [0,T]\}$ is a right continuous (sub)martingale. Furthermore, if σ is another stopping time and $\sigma \le \tau \le T$, then M_τ and M_σ are integrable random variables and

$$E[M_\tau] \ge E[M_\sigma], \tag{C.3}$$

where equality holds if M is a martingale.

We shall use the last result above in particular with $\sigma = 0$.

Appendix D

Brownian Motion

In this appendix (D), we briefly review some definitions and basic results for Brownian motion. For more details and further background we refer the reader to Chung [9] and Chung and Williams [11].

All stochastic processes considered in this appendix are indexed by time lying in the compact time interval $[0,T]$ for a fixed value T in $[0,\infty)$ and all processes are assumed to be defined on a given complete probability space $(\Omega, \mathcal{F}, P)$.

A standard *one-dimensional Brownian motion* (on the time interval $[0,T]$) is a one-dimensional process $W = \{W_t, t \in [0,T]\}$ such that

(i) $W_0 = 0$ P-a.s.,

(ii) all of the sample paths of W are continuous,

(iii) W has independent increments: $W_{t_1} - W_{t_0}, \ldots, W_{t_n} - W_{t_{n-1}}$ are independent for any $0 \le t_0 < t_1, < \cdots < t_n < \infty$, $n = 1, 2, \ldots,$

(iv) for $0 \le s < t < \infty$, $W_t - W_s$ is a normally distributed random variable with mean zero and variance $t - s$.

One can think of standard Brownian motion as a limit in distribution of a sequence of simple symmetric random walks where the time step and walk step size are scaled to zero in a suitable way (see e.g., [27]).

For a positive integer $n \ge 1$, a standard *n-dimensional Brownian motion* is an n-dimensional process $W = \{W_t, t \in [0,T]\}$ such that the coordinate processes $W^1, W^2, \ldots, W^n$ are mutually independent and for each $i \in \{1, \ldots, n\}$, $W^i = \{W^i_t, t \in [0,T]\}$ is a standard one-dimensional Brownian motion.

Assuming that W is a standard n-dimensional Brownian motion defined on the complete probability space $(\Omega, \mathcal{F}, P)$, for each $t \in [0,T]$, let $\mathcal{F}^o_t = \sigma\{W_s : 0 \le s \le t\}$, the smallest σ-algebra on Ω with respect to which W_s is measurable for each $0 \le s \le t$. For each $t \in [0,T]$, let $\mathcal{F}_t$ denote the smallest σ-algebra containing $\mathcal{F}^o_t$ and all of the P-null sets in $\mathcal{F}$. This is called the augmentation of $\mathcal{F}_t$ by the P-null sets, and $\{\mathcal{F}_t, t \in [0,T]\}$ is called the filtration generated by the Brownian motion W under P. The latter is sometimes referred to as the *standard filtration* associated with the Brownian motion or simply the Brownian filtration. It is well known that this filtration is right continuous, i.e., for each $t \in [0,T)$, $\mathcal{F}_t = \mathcal{F}_{t+} \equiv \cap_{s\in(t,T]}\mathcal{F}_s$, cf. Chung [9], Section 2.3, Theorem 4. Hence the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ satisfies the usual conditions. It is well known that in this case when the filtration is the standard Brownian filtration, every martingale relative to $\{\mathcal{F}_t\}$ has a continuous modification, cf. Theorem V.3.5 in Revuz and Yor [35]. It is straightforward to check that $\{W_t, \mathcal{F}_t, t \in [0,T]\}$ is a continuous, n-dimensional martingale.

Bibliography

[1] Artzner, P., and Heath, D. (1995). Approximate completeness with multiple martingale measures, *Mathematical Finance,* **5,** 1–11.

[2] Bank for International Settlements (2005). *OTC derivatives market activity in the second half of 2004 (2005),* report of the Monetary and Economic Department, Basel, Switzerland.

[3] Bertsimas, D., and Tsitsiklis, J. N. (1997). *Introduction to Linear Optimization,* Athena Scientific, Belmont, MA.

[4] Bingham, N. H., and Kiesel, R. (1998). *Risk-Neutral Valuation,* Springer-Verlag, New York.

[5] Black, F., and Scholes, M. (1973). The pricing of options and corporate liabilities, *J. Political Economy*, **81,** 637–659.

[6] Cameron, R. H., and Martin, W. T. (1944). Transformations of Wiener integrals under translations, *Annals of Mathematics,* **45,** 386–396.

[7] Cameron, R. H., and Martin, W. T. (1945). Transformations of Wiener integrals under a general class of linear transformations, *Trans. Amer. Math. Soc.,* **58,** 184–219.

[8] Cameron, R. H., and Martin, W. T. (1949). The transformation of Wiener integrals by nonlinear transformations, *Trans. Amer. Math. Soc.,* **66,** 253–283.

[9] Chung, K. L. (1982). *Lectures from Markov Processes to Brownian Motion,* Springer-Verlag, New York.

[10] Chung, K. L. (2000). *A Course in Probability Theory Revised,* Academic Press, San Diego.

[11] Chung, K. L., and Williams, R. J. (1990). *Introduction to Stochastic Integration,* Second Edition, Birkhäuser, Boston.

[12] Cox, J. C., Ross, S. A., and Rubinstein, M. (1979). Option pricing: a simplified approach, *J. Finan. Econom.*, **7,** 229–263.

[13] Davis, M. H. A. (2004). Valuation, hedging and investment in incomplete financial markets, in *Applied Mathematics Entering the 21st Century,* eds. J. M. Hill and R. Moore, Society for Industrial and Applied Mathematics, 49–70.

[14] Delbaen, F. and Schachermayer, W. (1994). A general version of the fundamental theorem of asset pricing, *Math. Annalen,* **300,** 463–520.

[15] Elliott, R. J., and Kopp, P. E. (2004). *Mathematics of Financial Markets,* Second Edition, Springer-Verlag, New York.

[16] Föllmer, H., and Schied, A. (2002). *Stochastic Finance – an Introduction in Discrete Time,* de Gruyter, Berlin.

[17] Fouque, J.-P., Papanicolaou, G., and Sircar, K. R. (2000). *Derivatives in Financial Markets with Stochastic Volatility,* Cambridge University Press, Cambridge, U.K.

[18] Girsanov, I. V. (1960). On transforming a certain class of stochastic processes by absolutely continuous substitutions of measures, *Theory of Probability and Its Applications,* **5,** 285–301.

[19] Harris, L. (2003). *Trading and Exchanges,* Oxford University Press, Oxford, U.K.

[20] Harrison, J. M., and Kreps, D. M. (1979). Martingales and arbitrage in multiperiod securities markets, *J. Economic Theory*, **20,** 381–408.

[21] Harrison, J. M., and Pliska, S. R. (1981). Martingales and stochastic integrals in the theory of continuous trading, *Stochastic Processes and their Applications*, **11,** 215–260.

[22] Harrison, J. M., and Pliska, S. R. (1983). A stochastic calculus model of continuous trading: complete markets, *Stochastic Processes and their Applications*, **15,** 313–316.

[23] Hull, J. C. (2005). *Options, Futures and other Derivative Securities,* Prentice Hall, Sixth Edition, Upper Saddle River, New Jersey.

[24] Jacka, S. D. (1991). Optimal stopping and the American put, *Mathematical Finance,* **1,** 1–14.

[25] Jacod, J., and Yor, M. (1977). Étude des solutions extrémales et représentation intégrales des solutions pour certains problémes de martingales, *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebeite,* **38,** 83–125.

[26] Jarrow, R. A., Jin, X., and Madan, D. B. (1999). The second fundamental theorem of asset pricing, *Mathematical Finance,* **9,** 255–273.

[27] Karatzas, I., and Shreve, S. E. (1997). *Brownian Motion and Stochastic Calculus,* Second Edition, Springer-Verlag, New York.

[28] Karatzas, I., and Shreve, S. E. (1998). *Methods of Mathematical Finance,* Springer-Verlag, New York.

[29] Kreps, D. M. (1981). Arbitrage and equilibrium in economies with infinitely many commodities, *J. of Mathematical Economics,* **8,** 15–35.

[30] Lamberton, D., and Lapeyre, B. (1996). *Introduction to Stochastic Calculus Applied to Finance,* Chapman and Hall, London.

[31] Merton, R. C. (1973). Theory of rational option pricing, *Bell J. Econ. Mgt. Science*, **4,** 141–183.

[32] Musiela, M., and Rutkowski, M. (1998). *Martingale Methods in Financial Modelling,* Springer, New York.

[33] Pliska, S. (1997). *Introduction to Mathematical Finance: Discrete Time Models,* Blackwell, Oxford, UK.

[34] Protter, P. (2001). A partial introduction to financial asset pricing theory, *Stochastic Processes and their Applications,* **91,** 169–203.

[35] Revuz, D., and Yor, M. (1999). *Continuous Martingales and Brownian Motion,* Third Edition, Springer-Verlag, New York.

[36] Rogers, L. C. G., and Williams, D. (1994). *Diffusions, Markov Processes, and Martingales,* Volume 1, Second Edition, John Wiley and Sons, Chichester, U.K.

[37] Rogers, L. C. G., and Williams, D. (1987). *Diffusions, Markov Processes, and Martingales,* Volume 2, John Wiley and Sons, Chichester, U.K.

[38] Williams, D. (1991). *Probability with Martingales,* Cambridge University Press, Cambridge, U.K.