Stochastic Analysis in Continuous Time[1]

Stochastic basis with the usual assumptions: $\mathbb{F} = (\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$, where Ω is the probability space (a *set* of elementary outcomes ω); $\mathcal{F}$ is the *sigma-algebra of events* (sub-sets of Ω that can be *measured* by $\mathbb{P}$); $\{\mathcal{F}_t\}_{t\geq 0}$ is the `filtration`: $\mathcal{F}_t$ is the sigma-algebra of events that happened by time t so that $\mathcal{F}_s \subseteq \mathcal{F}_t$ for $s < t$; $\mathbb{P}$ is `probability`: finite non-negative countably additive measure on the $(\Omega, \mathcal{F})$ normalized to one ($\mathbb{P}(\Omega) = 1$). The `usual assumptions` are about the filtration: $\mathcal{F}_0$ is $\mathbb{P}$-complete, that is, if $A \in \mathcal{F}_0$ and $\mathbb{P}(A) = 0$, then every sub-set of A is in $\mathcal{F}_0$ [this minimises the technical difficulties related to null sets and is not hard to achieve]; and $\{\mathcal{F}_t\}_{t\geq 0}$ is *right-continuous*, that is, $\mathcal{F}_t = \bigcap_{\varepsilon>0} \mathcal{F}_{t+\varepsilon}$, $t \geq 0$ [this simplifies some further technical developments, like stopping time vs. optional time and might be hard to achieve, but is assumed nonetheless.[2]]

Stopping time τ on $\mathbb{F}$ is a random variable with values in $[0, +\infty]$ and such that, for every $t \geq 0$, the set $\{\omega : \tau(\omega) \leq t\}$ is in $\mathcal{F}_t$.[3] The sigma algebra $\mathcal{F}_\tau$ consists of all the events A from $\mathcal{F}$ such that, for every $t \geq 0$, $A \bigcap \{\omega : \tau(\omega) \leq t\} \in \mathcal{F}_t$. In particular, non-random $\tau = T \geq 0$ is a stopping time and $\mathcal{F}_\tau = \mathcal{F}_T$. For a stopping time τ, intervals such as $(0, \tau)$ denote a *random set* $\{(\omega, t) : 0 < t < \tau(\omega)\}$.

A random/stochastic process $X = X(t) = X(\omega, t)$, $t \geq 0$, is a collection of random variables[4]. Equivalently, X is a measurable mapping from Ω to $\mathbb{R}^{[0,+\infty)}$. The process X is called `measurable` if the mapping $(\omega, t) \mapsto X(\omega, t)$ is measurable as a function from $\Omega \times [0, +\infty)$ to $\mathbb{R}$. The process X is called `adapted` if $X(t) \in \mathcal{F}_t$, $t \geq 0$, that is, the random variable $X(t)$ is $\mathcal{F}_t$-measurable for every $t \geq 0$. *The standing assumption is that every process is measurable and adapted.*

The stopped process X^τ is defined for a stopping time τ as follows:

$$X^\tau(t) = X(t \wedge \tau) = \begin{cases} X(\omega, t), & t \leq \tau(\omega), \\ X(\omega, \tau(\omega)), & \tau(\omega) \leq t. \end{cases}$$

The filtration $\{\mathcal{F}_t^X\}_{t\geq 0}$ generated by the process X is $\mathcal{F}_t^X = \sigma\big(X(s),\ s \leq t\big)$.

A typical stopping time is $\inf\{t \geq 0 : X(t) \in A\}$ for a Borel set $A \in \mathbb{R}$, with convention $\inf\{\emptyset\} = +\infty$. In particular, setting $\tau_n = \inf\{t > 0 : |X(t)| > n\}$, we call $\tau_X^* = \lim_{n\to\infty} \tau_n$ the `explosion time` of X and say that X does not explode if $\mathbb{P}(\tau_X^* = +\infty) = 1$.

A (sample) path/trajectory of the process X is the mapping $t \mapsto X(\omega, t)$ for fixed ω. The two main spaces for the sample paths are $\mathcal{C}$ (continuous functions with the sup norm) and $\mathcal{D}$ (the Skorokhod space of càdlàg [right-continuous, with limits from the left] functions with a suitable metric). On a finite time interval, both $\mathcal{C}$ and $\mathcal{D}$ are complete separable metric spaces.[5]

Equality of random processes. In general $X = Y$ for two random processes can mean

(1) Equality of finite-dimensional distributions: random vectors $\{X(t_k), k = 1, \ldots, n\}$ and $\{Y(t_k), k = 1, \ldots, n\}$ have the same distributions for every finite collection of t_k.
(2) Equality in law: $\mathbb{P}(X \in \mathcal{A}) = \mathbb{P}(Y \in \mathcal{A})$ for all measurable sets $\mathcal{A}$ in some function space [typically $\mathcal{C}$ or $\mathcal{D}$].
(3) Equality as modifications: $\mathbb{P}\big(X(t) = Y(t)\big) = 1$ for all $t \geq 0$.
(4) X and Y are indistinguishable: $\mathbb{P}\big(X(t) = Y(t) \text{ for all } t \geq 0\big) = 1$.

If X, Y have sample paths in $\mathcal{D}$, then equality as modifications implies that X and Y are indistinguishable.

The Kolmogorov continuity criterion: if $\mathbb{E}|X(t) - X(s)|^p \leq C|t - s|^{1+q}$, $C, p, q > 0$, then X has a continuous modification [in fact, the sample paths are Hölder continuous of every order less than q/p].[6]

Special types of random processes.

(1) X has `independent increments` if $X(t) - X(s)$ is independent of $\mathcal{F}_s$ for all $t > s \geq 0$.

---

[1] Sergey Lototsky, USC; updated on June 25, 2022

[2] Note that any filtration can be made right-continuous by re-defining $\mathcal{F}_t$ to be $\bigcap_{\varepsilon>0} \mathcal{F}_{t+\varepsilon}$, but this "cheap trick" enlarges the filtration and can potentially ruin some useful properties (Markov, martingale) available under the original filtration.

[3] Optional time has $\{\tau < t\} \in \mathcal{F}_t$, $t \geq 0$; every stopping time is optional, and, for a right-continuous filtration, every optional time is stopping.

[4] more generally, $X(t)$ can take values in any measurable space

[5] In fact, both $\mathcal{C}$ and $\mathcal{D}$ are Banach spaces with respect to the sup norm, but $\mathcal{D}$ is not separable with the corresponding metric; this is why a special (Skorokhod) metric is necessary

[6] In the case of random fields, that is, $t, s \in \mathbb{R}^n$, the same conclusion requires the inequality to hold with power $n + q$ instead of $1 + q$.

(2) X is `Markov` if $\mathbb{P}\big(X(t)\in A|\mathcal{F}_s\big)=\mathbb{P}\big(X(t)\in A|X(s)\big)$, $t>s\geq 0$.
(3) X is a `(sub/super) martingale` if $\mathbb{E}|X(t)|<\infty$ for all $t>0$ and $\mathbb{E}\big(X(t)|\mathcal{F}_s\big)\ (\geq/\leq)=X(s)$.
(4) X is a `square-integrable martingale` if X is a martingale and $\mathbb{E}|X(t)|^2<\infty$ for all $t>0$.
(5) X is a `local (square-integrable) martingale` if there is a sequence $\tau_n,\ n\geq 1$, of stopping times such that, for each n, the process X^{τ_n} is a (square-integrable) martingale and also, with probability one, $\tau_{n+1}\geq\tau_n$ and $\lim_{n\to+\infty}\tau_n=+\infty$.
(6) X is a `strict local martingale` if it is a local martingale but not a martingale.
(7) X is a `semimartingale if` $X=M+A$ for a local martingale M and a process of bounded variation A.
(8) X is `predictable` if it is measurable with respect to the sigma-algebra on $\Omega\times[0,+\infty)$ generated by continuous processes [random processes with continuous sample paths]; in particular, a continuous process is predictable.
(9) X is a `Wiener process` if $X(0)=0$ and the processes $t\mapsto X(t)$ and $t\mapsto X^2(t)-t$ are continuous martingales.

Basic facts.

(1) If $W=W(t)$ is a standard Brownian motion, then $\mathcal{F}_t^W$ is right-continuous. Once $\mathcal{F}_0^W$ is $\mathbb{P}$-completed, W becomes a continuous square-integrable martingale on the stochastic basis $(\Omega,\mathcal{F},\{\mathcal{F}_t^W\}_{t\geq 0},\mathbb{P})$ satisfying the usual assumptions.
(2) A continuous local martingale M is a local square-integrable martingale: replace the original τ_n with $\tau_n\wedge\inf\{t\geq 0:|M(t)|\geq n\}$.
(3) A process X with independent increments is Markov; if also $\mathbb{E}|X(t)|<\infty$, then $t\mapsto X(t)-\mathbb{E}X(t)$ is a martingale.
(4) The process $X=X(t)$ with $X(0)=0$ is a martingale if $\mathbb{E}X(\tau)=0$ for every bounded stopping time τ.
(5) The Optional Stopping Theorem: if X is a martingale with $X(0)=0$ and τ is a stopping time with $\mathbb{P}(\tau<\infty)=1$, then $\mathbb{E}X(\tau)=0$ as long as X and τ "cooperate" with each other [bounded τ or uniformly integrable family $\{X(t),\ t\geq 0\}$ always works].
(6) If X is a submartingale[7] and the function $t\mapsto\mathbb{E}X(t)$ is in $\mathcal{D}$, then X has a modification in $\mathcal{D}$; in particular, every martingale has a càdlàg modification.
(7) Jensen's inequality: If X is a martingale and $f=f(x)$ is convex, with $\mathbb{E}|f(X(t))|<\infty$, then $f(X)$ is a submartingale.
(8) Doob-Meyer decomposition: If X is a submartingale with càdlàg sample paths, then $X=M+A$ for a local martingale M and a predictable non-decreasing process A, and the representation is unique up to a modification.
(9) Lévy characterisation of the Brownian motion: A Wiener process is a standard Brownian motion.
(10) A non-negative local martingale is supermartingale; if the trajectories are càdlàg, then there is no explosion.

Two basic constructions.

(1) Quadratic characteristic $\langle X\rangle$ of a local square-integrable martingales X is the increasing process in the Doob-Meyer decomposition of X^2. To indicate time dependence, notation $\langle X\rangle_t$ is used. For example, if W is Wiener process, then $\langle W\rangle_t=t$. If X is a square-integrable martingale, then $\mathbb{E}X^2(t)=\mathbb{E}\langle X\rangle_t$; if N is Poisson with intensity λ and $M(t)=N(t)-\lambda t$, then $\langle M\rangle_t=\lambda t$. If X is a continuous square-integrable martingale, then $\langle X\rangle_t$ is the `quadratic variation` of X: $\langle X\rangle_t$ is the limit in probability of $\sum_{k=1}^n\big(X(t_k)-X(t_{k-1})\big)^2$, as the size of the partition of $[0,t]$ goes to 0 [Karatzas-Shreve, *Brownian motion and stoch. calc*, Thm. 1.5.8].
(2) Local time $L^a=L^a(t)$, $a\in\mathbb{R}$, of a *continuous* martingale X is the increasing process in the Doob-Meyer decomposition of $|X-a|$.

Burholder-Davis-Gundy (BDG) inequality: Let $M=M(t)$ be a continuous local martingale with $M(0)=0$ and let τ be a stopping time. Define $M^*(\tau)=\sup_{t\leq\tau}|M(t)|$. Then, for every $p>0$, there exist positive numbers c_p and C_p such that $c_p\mathbb{E}\langle M\rangle_\tau^{p/2}\leq\mathbb{E}\big(M^*(\tau)\big)^p\leq C_p\mathbb{E}\langle M\rangle_\tau^{p/2}$. For $0<p<2$, we can take $c_p=\frac{2-p}{4-p}$ and $C_p=\frac{4-p}{2-p}$ [L-Sh-Mart, Thm. 1.9.5] so that $c_1=1/3, C_1=3$.

In general, **the Itô stochastic integral** $\int_0^t Y(s)\,dX(s)$ of a *predictable* process Y with respect to a semimartingale X is defined as a suitable limit of $\sum_k Y(t_k)\big(X(t_{k+1})-X(t_k)\big)$.

In particular, if $X=W$ is a Wiener process, then, as long as Y is adapted and $\int_0^T Y^2(t)dt<\infty$ with probability one for some non-random $T>0$, the stochastic integral $V(t)=\int_0^t Y(s)dW(s)$ defines a continuous local martingale for $t\in[0,T]$ with $\langle V\rangle_t=\int_0^t Y^2(s)\,ds$.

[7] A single word, as opposed to dashed, as in *sub-martingale*, seems to be the standard; constructions such as *localmartingale* can sometime happen too.