Stochastic Processes

A very simple introduction

Péter Medvegyev

2009, January

Summary from measure theory

Definition

$(X, \mathcal{A})$ is a measurable space if $\mathcal{A} \neq \varnothing$ is a set of subsets of X such that

1. if $A \in \mathcal{A}$, then $A^c \in \mathcal{A}$,
2. if (A_n) is a countable sequence of sets of $\mathcal{A}$, then $\cup_n A_n \in \mathcal{A}$,
3. if (A_n) is a countable sequence of sets of $\mathcal{A}$, then $\cap_n A_n \in \mathcal{A}$.

The set of sets $\mathcal{A}$ is called σ-algebra, the sets in $\mathcal{A}$ are the measurable sets.

Measurable functions and random variables

Definition

If $(X, \mathcal{A})$ and $(Y, \mathcal{B})$ are measurable spaces a mapping $f : X \to Y$ is called measurable if for every $B \in \mathcal{B}$

$$f^{-1}(B) \stackrel{\circ}{=} \{x \in X \mid f(x) \in B\} \in \mathcal{A}.$$

If $Y = \mathbb{R}$ then $\mathcal{B}$ is always the set of Borel measurable sets which is the smallest σ-algebra containing the intervals. If the image space is the set of real numbers we are talking about measurable functions. If $(\Omega, \mathcal{A}, \mathbf{P})$ is a probability space then the random variables are the measurable (extended real valued) functions. We say that two random variables ξ_1 and ξ_2 are equivalent or indistinguishable if

$$\mathbf{P}(\xi_1 \neq \xi_2) = 0.$$

Measurable functions and random variables

Definition

If $s = \sum_k c_k \chi_{A_k}$ then s is a step function. If the sets A_k are measurable then s is a measurable step function.

Theorem

If $f \geq 0$ is a measurable function then there is a sequence of measurable step functions (s_n) such that $0 \leq s_n \nearrow f$.

Measurable functions and random variables

Theorem

Let $(X, \mathcal{A})$ *be a measurable space.*

1. *If* f *and* g *are measurable functions then* $f + g$, $f \cdot g$, f/g *are measurable. (Of course only when they are well-defined.) If* λ *is a real number and* f *is a measurable function then* λf *is also measurable.*
2. *If* (f_n) *is a sequence of measurable functions then*

$$\sup_n f_n, \inf_n f_n, \limsup_{n\to\infty} f_n, \liminf_{n\to\infty} f_n$$

   are measurable.

Measurable functions and random variables

Theorem

If (f_n) *is a sequence of measurable functions and the limit*

$$f(x) \circeq \lim_{n\to\infty} f_n(x)$$

exists on a set A *then*

$$g(x) \circeq \begin{cases} f(x) & \text{if} \quad x \in A \\ 0 & \text{if} \quad x \notin A \end{cases}$$

is also measurable.

Integral and expected value

Definition

The integral of a step function $s \geq 0$ is

$$\int_X s d\mu \mathrel{\mathring{=}} \sum_k c_k \mu\left(A_k\right).$$

The integral of a measurable function $f \geq 0$ is

$$\int_X f d\mu \mathrel{\mathring{=}} \sup_{0 \leq s \leq f} \int_X s d\mu$$

where $s \geq 0$ is a measurable step function. If f is arbitrary then

$$\int_X f d\mu \mathrel{\mathring{=}} \int_X f^+ d\mu - \int_X f^- d\mu$$

assuming that the sum is not of the form $\infty - \infty$. A function f is integrable if the integral is finite.

Integral and expected value

Definition

If μ is a probability measure then

$$\mathbf{E}(\xi) \doteq \int_{\Omega} \xi \, d\mathbf{P}.$$

Integral and expected value

Theorem

If f and g are measurable functions then

$$\int_X (f+g)\, d\mu = \int_X f d\mu + \int_X g d\mu$$

assuming that all the three integrals exists. If two exist then one can prove that the third also exits.

Integral and expected value

Theorem (Monotone convergence)

If $0 \leq f_n \nearrow f$ *are measurable functions then*

$$\int_X f_n d\mu \nearrow \int_X f d\mu$$

Theorem (Dominated convergence)

If (f_n) *are measurable functions and* $f_n \to f$ *and there is an integrable function* g *such that* $|f_n| \leq g$ *the*

$$\int_X f_n d\mu \to \int_X f d\mu$$

and

$$\int_X |f_n - f| \, d\mu \to 0$$

Banach space of integrable functions

Definition

If $p \geq 1$ and $|f|^p$ is integrable then we say that $f \in L^p$.

Theorem

The equivalence classes of integrable functions L^p with the norm

$$\|f\|_p \circeq \sqrt[p]{\int_X |f|^p \, d\mu}$$

is a Banach space, that is a complete normed space.

Conditional expectation

Definition

Let ξ be a random variable $\mathcal{F} \subseteq \mathcal{A}$ is a sub σ-algebra. Assume that $\mathbf{E}(\xi)$ exists. The conditional expectation $\mathbf{E}(\xi \mid \mathcal{F})$ is an $\mathcal{F}$-measurable random variable such that

$$\int_F \xi d\mathbf{P} = \int_F \mathbf{E}(\xi \mid \mathcal{F})\, d\mathbf{P}, \quad \forall F \in \mathcal{F}.$$

$$\mathbf{P}(A \mid \mathcal{F}) \circeq \mathbf{E}(\chi_A \mid \mathcal{F}).$$

Example

If $\mathcal{F} = \{\emptyset, \Omega\}$ then

$$\mathbf{E}(\xi \mid \mathcal{F}) = \mathbf{E}(\xi).$$

Theorem

Assume that all variables below are integrable.

- *If* ξ *and* $\mathcal{F}$ *are independent, then* $\mathbf{E}(\xi \mid \mathcal{F}) = \mathbf{E}(\xi)$. *Specially if* $\mathcal{F} = \{\emptyset, \Omega\}$, *then* $\mathbf{E}(\xi \mid \mathcal{F}) = \mathbf{E}(\xi)$.
- *If* ξ *is* $\mathcal{F}$*-measurable and the conditional expectation* $\mathbf{E}(\xi \mid \mathcal{F})$ *exists, then* $\mathbf{E}(\xi \mid \mathcal{F}) = \xi$.
- *If* $\xi \leq \eta$, *then* $\mathbf{E}(\xi \mid \mathcal{F}) \leq \mathbf{E}(\eta \mid \mathcal{F})$ *,specially*

$$0 \leq \mathbf{P}(A \mid \mathcal{F}) \leq 1,$$
$$|\mathbf{E}(\xi \mid \mathcal{F})| \leq \mathbf{E}(|\xi| \mid \mathcal{F}).$$

- $\mathbf{E}(\xi + \eta \mid \mathcal{F}) = \mathbf{E}(\xi \mid \mathcal{F}) + \mathbf{E}(\eta \mid \mathcal{F})$. *If* A *and* B *are disjoint sets, then*

$$\begin{aligned} \mathbf{P}(A \cup B \mid \mathcal{F}) &= \mathbf{P}(A \mid \mathcal{F}) + \mathbf{P}(B \mid \mathcal{F}), \\ \mathbf{P}(B \setminus A \mid \mathcal{F}) &= \mathbf{P}(B \mid \mathcal{F}) - \mathbf{P}(A \mid \mathcal{F}), \quad A \subseteq B. \end{aligned}$$

Tower laws

Theorem

- *If* $\mathcal{G} \subseteq \mathcal{F}$, *then*

$$\mathbf{E}\left(\mathbf{E}\left(\xi \mid \mathcal{F}\right) \mid \mathcal{G}\right) = \mathbf{E}\left(\xi \mid \mathcal{G}\right).$$

Specially

$$\mathbf{E}\left(\mathbf{E}\left(\xi \mid \mathcal{F}\right)\right) = \mathbf{E}\left(\xi\right).$$

- *If* $\mathcal{G} \subseteq \mathcal{F}$, *then*

$$\mathbf{E}\left(\mathbf{E}\left(\xi \mid \mathcal{G}\right) \mid \mathcal{F}\right) = \mathbf{E}\left(\xi \mid \mathcal{G}\right).$$

Monotone convergence

Theorem

The monotone convergence theorem is satisfied, that is, if $\xi_n \geq 0$, *and* $\xi_n \nearrow \xi$,*then*

$$\mathbf{E}(\xi_n \mid \mathcal{F}) \overset{\cdot}{\nearrow} \mathbf{E}(\xi \mid \mathcal{F}).$$

Specially, if $A_n \nearrow A$, *then*

$$\mathbf{P}(A_n \mid \mathcal{F}) \overset{\cdot}{\nearrow} \mathbf{P}(A \mid \mathcal{F}).$$

If $\xi_n \geq 0$, *then*

$$\sum_{n=1}^{\infty} \mathbf{E}(\xi_n \mid \mathcal{F}) = \mathbf{E}\left(\sum_{n=1}^{\infty} \xi_n \mid \mathcal{F}\right).$$

If the sets (A_n) *are disjoint and* $A = \cup_n A_n$, *then*

$$\sum_{n=1}^{\infty} \mathbf{P}(A_n \mid \mathcal{F}) = \mathbf{P}(A \mid \mathcal{F}).$$

Dominated convergence

Theorem

The dominated convergence theorem is satisfied, that is if

$$\lim_{n\to\infty} \xi_n = \xi, \quad \text{and} \quad |\xi_n| \leq \eta,$$

where $\mathbf{E}(\eta) < \infty$, *then*

$$\lim_{n\to\infty} \mathbf{E}\left(|\xi_n - \xi| \mid \mathcal{F}\right) = 0, \quad \lim_{n\to\infty} \mathbf{E}\left(\xi_n \mid \mathcal{F}\right) = \mathbf{E}\left(\xi \mid \mathcal{F}\right).$$

Specially, if $A_n \searrow A$, *then*

$$\lim_{n\to\infty} \mathbf{P}\left(A_n \mid \mathcal{F}\right) = \mathbf{P}\left(A \mid \mathcal{F}\right).$$

Homogeneity

Theorem

If variable η is $\mathcal{F}$-measurable, and $\mathbf{E}(\xi)$ and $\mathbf{E}(\eta \cdot \xi)$ are finite, then

$$\mathbf{E}(\eta \cdot \xi \mid \mathcal{F}) = \eta \cdot \mathbf{E}(\xi \mid \mathcal{F}).$$

It is also true, if η and ξ are nonnegative.

Definition of Stochastic Processes

Let us fix a probability space $(\Omega, \mathcal{A}, \mathbf{P})$. As in probability theory we refer to the set of real-valued $(\Omega, \mathcal{A})$-measurable functions as *random variables*. In the theory of stochastic processes random variables very often have infinite value. Hence the image space of the measurable functions is not $\mathbb{R}$ but the set of extended real numbers $\overline{\mathbb{R}} \circeq [-\infty, \infty]$. The most important examples of random variables with infinite value are stopping times. Stopping times give the random time of the occurrence of observable events. If for a certain outcome ω the event never occurs it is reasonable to say that the value of the stopping time for this ω is $+\infty$. In the most general sense *stochastic processes* are such functions $X(t, \omega)$ that for any fixed parameter t the mappings $\omega \mapsto X(t, \omega)$ are random variables on $(\Omega, \mathcal{A}, \mathbf{P})$.

Definition of Stochastic Processes

The set of possible time parameters Θ is some subset of the extended real numbers. In the theory of continuous-time stochastic processes Θ is an interval, generally $\Theta = \mathbb{R}_+ \overset{\circ}{=} [0, \infty)$, but sometimes $\Theta = [0, \infty]$ and $\Theta = (0, \infty)$ is also possible. If we do not say explicitly what the domain of definition of the stochastic process is, then Θ is $\mathbb{R}_+$.

Definition of Stochastic Processes

It is very important to append some remarks to this definition. In probability theory the random variables are equivalence classes which means that the random variables $X(t)$ are defined up to measure zero sets. This means that in general $X(t,\omega)$ is meaningless for a fixed ω. If the possible values of the time parameter t are countable then we can select from the equivalence classes $X(t)$ one element and we can fix a measure zero set and outside of this set the expressions $X(t,\omega)$ are meaningful. But this is impossible if Θ is not countable. Therefore we shall always assume that $X(t)$ is a function already carefully selected from its equivalence class. To put it in another way: when one defines a stochastic process, one should fix the space of possible trajectories and the stochastic processes are function-valued random variables which are defined on the space $(\Omega, \mathcal{A}, \mathbf{P})$.

Definition of Stochastic Processes

Definition

Let us fix the probability space $(\Omega, \mathcal{A}, \mathbf{P})$ and the set of possible time parameters Θ. The function X defined on $\Theta \times \Omega$ is a *stochastic process* over $\Theta \times \Omega$ if for every $t \in \Theta$ it is measurable on $(\Omega, \mathcal{A}, \mathbf{P})$ in its second variable.

Definition

If we fix an outcome $\omega \in \Omega$ then the function $t \mapsto X(t, \omega)$ defined over Θ is the *trajectory* or *realization* of X corresponding to the outcome ω. If all the trajectories of the process X have a certain property then we say that the process itself has this property. For example if all the trajectories of X are continuous then we say that X is continuous, if all the trajectories of X have finite variation then we say that X has finite variation, etc.

When do stochastic process be equal?

Definition

Let X and Y be two stochastic processes on the probability space $(\Omega, \mathcal{A}, \mathbf{P})$.

1. The process X is a *modification* of the process Y if for all $t \in \Theta$ the variables $X(t)$ and $Y(t)$ are almost surely equal, that is for all $t \in \Theta$

$$\mathbf{P}\left(X(t) = Y(t)\right) \circeq \mathbf{P}\left(\{\omega : X(t,\omega) = Y(t,\omega)\}\right) = 1.$$

   By this definition, the set of outcomes ω where $X(t,\omega) \neq Y(t,\omega)$, can depend on $t \in \Theta$.

2. The processes X and Y are *indistinguishable* if there is a set $N \subseteq \Omega$ which has probability zero, and whenever $\omega \notin N$ then $X(\omega) = Y(\omega)$, that is $X(t,\omega) = Y(t,\omega)$ for all $t \in \Theta$ and $\omega \notin N$.

When do stochastic process be equal?

Definition

Let X and Y be two stochastic processes on the probability space $(\Omega, \mathcal{A}, \mathbf{P})$. $X = Y$ means that X and Y are *indistinguishable, that is the trajectories are almost surely equal.*

Filtration, usual conditions

Definition

If $(\mathcal{F}_t)_t$ is a monotone increasing set of σ-algebras, then $\mathcal{F} = (\mathcal{F}_t)_t$ is called filtration, that is $\mathcal{F}$ is a filtration if for $s < t$ $\mathcal{F}_s \subseteq \mathcal{F}_t$. The stochastic process X is called adapted, if $X(t)$ is $\mathcal{F}_t$ measurable for all t.

Definition

The filtration $(\Omega, \mathcal{A}, \mathbf{P}, \mathcal{F})$satisfies the usual conditions if

1. $\mathcal{A}$ is complete, that is if $N \subseteq M$, and $\mathbf{P}(M) = 0$, then $N \in \mathcal{A}$.
2. $\mathcal{F}_t$ for all t contains the measure zero sets,
3. the filtration $\mathcal{F}$ is right-continuous, that is for all t $\mathcal{F}_t = \mathcal{F}_{t+} \doteq \cap_{s>t}\mathcal{F}_s$.

We will always assume, that the usual conditions are satisfied. We will also assume that every process is adapted to the given filtration.

Filtration, usual conditions

The usual interpretation of the σ-algebra $\mathcal{F}_t$ is that it contains the events which occurred up to time t, that is $\mathcal{F}_t$ contains the information which is available at moment t.

As $\mathcal{F}_t$ is the information at moment t one can interpret $\mathcal{F}_{t+}$ as the information available infinitesimally just after t. If a process has "speed", then one can foresee the future infinitesimally as by definition in this case the right and left derivatives are equal and one knows the left derivative at time t. The stochastic processes are very often non-differentiable. Even in this case one should assume this "infinitesimal wisdom". The trick is that one should add the measure zero sets to the filtration anyway, and surprisingly it makes the filtrations right-continuous. We will return to this later.

Generated filtration

It is a quite natural question how one can define a filtration $\mathcal{F}$?

Definition

Let X be a stochastic process that is let X be a function of two variables. .Let us define the σ-algebras $\mathcal{F}_t^X \subseteq \mathcal{A}$ generated by the sets

$$\{X(t_1) \in I_1, \ldots, X(t_n) \in I_n\}$$

where $t_1, \ldots, t_n \leq t$ are arbitrary elements in Θ and $I_1, \ldots, I_n$ are arbitrary intervals. Obviously if $s < t$ then $\mathcal{F}_s^X \subseteq \mathcal{F}_t^X$, hence $\mathcal{F}^X$ is really a filtration. $\mathcal{F}^X$ is called *the filtration generated by* X.

There is no guarantee that $\mathcal{F}^X$ satisfies the usual conditions. Very often if we add to $\mathcal{F}^X$ the measure zero sets the new filtration is right-continuous and the usual conditions hold.

Martingales

Definition

The stochastic process X is a martingale relative to a filtration $\mathcal{F}$, if

1. $\mathbf{E}(|X(t)|) < \infty$, for all t,
2. if $s < t$, then $\mathbf{E}(X(t) \mid \mathcal{F}_s) = X(s)$,
3. the trajectories are right-continuous, and have limits from the left.

Definition

The L^2-bounded martingales are denoted by $\mathcal{H}^2$.

It is trivial from the assumptions that every martingale is adapted.

Wiener processes

Definition

The process w called Wiener process, if

1. $w(0) = 0$,
2. w has stationary and independent increments,
3. w has continuous trajectories,
4. $w(t) \simeq N(0, t) \left(= N\left(0, \sigma^2\right)\right)$.
5. w is adapted.

Wiener processes on finite intervals

Example

On any finite interval Wiener processes are in $\mathcal{H}^2$, on infinite intervals the Wiener processes are just martingales.

Only the second condition is not trivial.

$$\begin{aligned}
\mathbf{E}(w(t) \mid \mathcal{F}_s) &= \mathbf{E}(w(s) + w(t) - w(s) \mid \mathcal{F}_s) = \\
&= \mathbf{E}(w(s) \mid \mathcal{F}_s) + \mathbf{E}(w(t) - w(s) \mid \mathcal{F}_s) = \\
&= w(s) + \mathbf{E}(w(t) - w(s)) = w(s).
\end{aligned}$$

Generated filtration of Wiener processes

Let

$$A \doteq \{\omega : w(\omega) \text{ is zero for some interval } [0, \delta(\omega)]\}.$$

As $\mathcal{F}_0^w = \{\emptyset, \Omega\}$ obviously $A \notin \mathcal{F}_0^w$. On the other hand $\mathcal{A} \in \mathcal{F}_{1/n}^w$ for every $n > 0$ as for every $n > 0$ we know that A holds or not. Hence $A \in \cap_n \mathcal{F}_{1/n}^w = \mathcal{F}_{0+}^w$. Hence $\mathcal{F}^w$ is not right-continuous.

One can prove that if $\mathcal{F}_t \doteq \sigma(\mathcal{F}_t^w, \mathcal{N})$ where $\mathcal{N}$ is the set of events with probability zero, then $\mathcal{F}$ satisfies the usual conditions. Very often $\mathcal{F}$ is generated by some finite number of Wiener processes, that is

$$\mathcal{F}_t \doteq \sigma\left(\mathcal{F}_t^{w_1}, \ldots, \mathcal{F}_t^{w_n}, \mathcal{N}\right).$$

Poisson processes

Definition

The process π called Poisson process, if

1. $\pi(0) = 0$,
2. the trajectories are continuous from the right and have limits from the left and have pure jumps of height one,
3. $\pi(t) \simeq P(\lambda t)$, that is $\mathbf{P}(\pi(t) = n) = \frac{(\lambda t)^n}{n!} \exp(-\lambda t)$.

Lévy processes

Definition

An adapted process X is called Lévy process (with respect to $\mathcal{F}$) if

1. $X(0) = 0$,
2. X has independent increments with respect to $\mathcal{F}$, that is if $t > s$ then $X(t) - X(s)$ is independent of $\mathcal{F}_s$,
3. X has stationary increments,
4. the trajectories are continuous from the right and have limits from the left.

Example

If the increments of a Lévy process X has zero expected value, then X is a martingale.

$$\begin{aligned}
\mathbf{E}\left(X\left(t\right) \mid \mathcal{F}_s\right) &= \mathbf{E}\left(X\left(s\right)+X\left(t\right)-X\left(s\right) \mid \mathcal{F}_s\right) = \\
&= \mathbf{E}\left(X\left(s\right) \mid \mathcal{F}_s\right)+\mathbf{E}\left(X\left(t\right)-X\left(s\right) \mid \mathcal{F}_s\right) = \\
&= X\left(s\right)+\mathbf{E}\left(X\left(t\right)-X\left(s\right)\right) = X\left(s\right).
\end{aligned}$$

Independence and Fourier transformation

Definition

If $\boldsymbol{\xi}$ is a vector of random variables then $\varphi(\mathbf{u}) \triangleq \mathbf{E}(\exp(i\langle\mathbf{u},\boldsymbol{\xi}\rangle))$ is called the Fourier transform or characteristic function of vector $\boldsymbol{\xi}$ and $L(\mathbf{u}) \triangleq \mathbf{E}(\exp(\langle\mathbf{u},\boldsymbol{\xi}\rangle))$ is called the moment generating function of ξ.

Theorem

The Fourier transform determines the distribution of $\boldsymbol{\xi}$.

Theorem

ξ *and* η *are independent if and only if*

$$\begin{aligned}\varphi_{(\xi,\eta)}(u,v) &= \mathbf{E}(\exp(iu\xi+iv\eta)) = \mathbf{E}(\exp(iu\xi))\,\mathbf{E}(\exp(iv\eta)) = \\ &= \varphi_\xi(u)\,\varphi_\eta(v).\end{aligned}$$

Independence and Fourier transformation

The easy part is:that if ξ and η are independent then

$$\varphi_{(\xi,\eta)}(u,v) = \varphi_{\xi}(u)\,\varphi_{\eta}(v)\,.$$

By definition ξ and η are independent if the generated σ-algebras are independent, therefore $\exp(iu\xi)$ and $\exp(iv\eta)$ are independent and the independent variables are uncorrelated.

Remarks

1. The moment generating function generally does not determine the distribution.

Remarks

1. The moment generating function generally does not determine the distribution.
2. $\varphi(u)$ is well-defined and finite for every u, but $L(u)$ is not necessarily finite for every u.

Remarks

1. The moment generating function generally does not determine the distribution.
2. $\varphi(u)$ is well-defined and finite for every u, but $L(u)$ is not necessarily finite for every u.
3. If L is well-defined and finite in a neighborhood of the origin, then L determines the distribution.

Remarks

1. The moment generating function generally does not determine the distribution.
2. $\varphi(u)$ is well-defined and finite for every u, but $L(u)$ is not necessarily finite for every u.
3. If L is well-defined and finite in a neighborhood of the origin, then L determines the distribution.
4. If $\xi \geq 0$ and
$$L(s) \circeq \mathbf{E}(\exp(-s\xi)), s \geq 0$$
then L determines the distribution. (In this case L is the Laplace transform of ξ.)

Exponential martingales

Example

If X is a Lévy process, then for every u

$$M(t,u) \circeq \frac{\exp(iuX(t))}{\varphi_t(u)}$$

is a martingale, where $\varphi_t(u) \circeq \mathbf{E}(\exp(iu(X(t))))$. Alternatively if $L_t(u)$ exists for some u then one can define

$$M(t,u) \circeq \frac{\exp(uX(t))}{L_t(u)}.$$

Exponential martingales

$$
\begin{aligned}
\mathbf{E}\left(M\left(t,u\right)\mid\mathcal{F}_s\right) &= \mathbf{E}\left(\frac{\exp\left(iuX\left(t\right)\right)}{\varphi_t\left(u\right)}\mid\mathcal{F}_s\right) = \\
&= \mathbf{E}\left(\frac{\exp\left(iu\left[X\left(s\right)+X\left(t\right)-X\left(s\right)\right]\right)}{\varphi_t\left(u\right)}\mid\mathcal{F}_s\right) = \\
&= \mathbf{E}\left(\frac{\exp\left(iu\left[X\left(s\right)+X\left(t\right)-X\left(s\right)\right]\right)}{\varphi_s\left(u\right)\varphi_{t-s}\left(u\right)}\mid\mathcal{F}_s\right) = \\
&= \mathbf{E}\left(M\left(s,u\right)\frac{\exp\left(iu\left[X\left(t\right)-X\left(s\right)\right]\right)}{\varphi_{t-s}\left(u\right)}\mid\mathcal{F}_s\right) = \\
&= \frac{M\left(s,u\right)}{\varphi_{t-s}\left(u\right)}\mathbf{E}\left(\exp\left(iu\left[X\left(t\right)-X\left(s\right)\right]\right)\mid\mathcal{F}_s\right) = \\
&= \frac{M\left(s,u\right)}{\varphi_{t-s}\left(u\right)}\mathbf{E}\left(\exp\left(iu\left[X\left(t\right)-X\left(s\right)\right]\right)\right) = M\left(s,u\right).
\end{aligned}
$$

Exponential martingale of Wiener processes

Example

If w is a Wiener process, then for every u

$$M(t,u) \doteq \frac{\exp(iuw(t))}{\exp(-tu^2/2)} = \exp\left(iuw(t) + \frac{tu^2}{2}\right)$$

$$M(t,u) \doteq \exp\left(uw(t) - \frac{tu^2}{2}\right)$$

are martingales, called, the exponential martingales of w.

Exponential martingale of Wiener processes

The second one is easy

$$\mathbf{E}\left(\exp\left(uw\left(t\right)\right)\right)=\mathbf{E}\left(\exp\left(u\sqrt{t}N\left(0,1\right)\right)\right)=L_{N(0,1)}\left(u\sqrt{t}\right).$$

$$\begin{aligned}L_{N(0,1)}\left(u\right) &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\exp\left(-\frac{1}{2}x^{2}\right)\exp\left(ux\right)dx=\\ &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\exp\left(-\frac{x^{2}-2ux}{2}\right)dx=\\ &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\exp\left(-\frac{x^{2}-2ux+u^{2}}{2}+\frac{u^{2}}{2}\right)dx=\\ &= \exp\left(\frac{u^{2}}{2}\right)\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\exp\left(-\frac{\left(x-u\right)^{2}}{2}\right)dx=\exp\left(\frac{u^{2}}{2}\right).\end{aligned}$$

Exponential martingale of Wiener processes

Also

$$\mathbf{E}\left(\exp\left(iuw\left(t\right)\right)\right) = \mathbf{E}\left(\exp\left(iu\sqrt{t}N\left(0,1\right)\right)\right) = \varphi_{N(0,1)}\left(u\sqrt{t}\right).$$

$$\begin{aligned}
\varphi_{N(0,1)}\left(u\right) &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\exp\left(-\frac{1}{2}x^2\right)\exp\left(iux\right)dx = \\
&= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\exp\left(-\frac{1}{2}x^2\right)\left(\cos ux + i\sin ux\right)dx = \\
&= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\exp\left(-\frac{1}{2}x^2\right)\cos ux dx.
\end{aligned}$$

Exponential martingale of Wiener processes

Differentiating and integrating by parts

$$\frac{d}{du}\varphi_{N(0,1)}(u) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\exp\left(-\frac{1}{2}x^2\right)\frac{d}{du}\cos uxdx =$$

$$= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}(-x)\exp\left(-\frac{1}{2}x^2\right)\sin uxdx =$$

$$= \left[\exp\left(-\frac{1}{2}x^2\right)\sin ux\right]_{-\infty}^{\infty} - \frac{1}{\sqrt{2\pi}}u\int_{-\infty}^{\infty}\exp\left(-\frac{1}{2}x^2\right)\cos uxdx =$$

$$= 0 - \frac{1}{\sqrt{2\pi}}u\int_{-\infty}^{\infty}\exp\left(-\frac{1}{2}x^2\right)\cos uxdx = -u\varphi(u).$$

As $\varphi(0) = 1$ the solution of this differential equation is $\varphi(u) = \exp\left(-u^2/2\right)$.

Compensated Poisson processes

Example

The compensated Poisson process $\pi(t) - \lambda t$ is a martingale.

$\pi(t) - \lambda t$ is a Lévy process with expected value zero.

Compensated square of Wiener processes

Example

If w is Wiener process, then $N(t) \circeq w^2(t) - t$ is a martingale.

$$\begin{aligned}\mathbf{E}(N(t) \mid \mathcal{F}_s) &= \mathbf{E}\left(w^2(t) - t \mid \mathcal{F}_s\right) = \\ &= \mathbf{E}\left((w(s) + w(t) - w(s))^2 - t \mid \mathcal{F}_s\right).\end{aligned}$$

$$\begin{aligned}\mathbf{E}\left((w(t) - w(s))^2 \mid \mathcal{F}_s\right) &= \mathbf{E}\left((w(t) - w(s))^2\right) = \\ &= \mathbf{E}\left(N^2(0, t-s)\right) = t - s.\end{aligned}$$

$$\mathbf{E}(w(s)(w(t) - w(s)) \mid \mathcal{F}_s) = w(s)\,\mathbf{E}((w(t) - w(s)) \mid \mathcal{F}_s) = 0.$$

Stopping times

Definition

A random variable $0 \leq \tau \leq \infty$ is called stopping time, if $\{\tau \leq t\} \in \mathcal{F}_t$ for all t This is the same as $\tau \wedge t$ is $\mathcal{F}_t$-measurable for every t that is $\tau \wedge t$ is "observable" for every t.

Theorem

If the trajectories of the adapted stochastic process X are right- or left-continuous, or the process is measurable with respect to the σ-algebra generated by the class of the adapted right-continuous processes, then for any Borel set B the hitting time

$$\tau \circeq \inf \{t : X(t) \in B\}$$

is a stopping time.

Stopped process, stopped variable

Definition

Let X be a stochastic process, and let τ be a stopping time.

1. By a *stopped* or *truncated process* we mean the process
$$X^{\tau}(t,\omega) \circeq X\left(\tau\left(\omega\right) \wedge t, \omega\right).$$
2. We shall call the random variable
$$X_{\tau}\left(\omega\right) \circeq X\left(\tau\left(\omega\right), \omega\right)$$
a *stopped variable*.
3. The *stopped σ-algebra* $\mathcal{F}_{\tau}$ is the set of events $A \in \mathcal{A}$ for which for all t
$$A \cap \{\tau \leq t\} \in \mathcal{F}_t.$$

Stopped process, stopped variable

Instead of X_τ we shall very often use the more readable notation $X(\tau)$. Observe that the definition of stopped variable is not entirely correct as X is generally not defined on the set $\{\tau = \infty\}$ and it is not clear what the definition of X_τ on this set is. If $\tau(\omega) \notin \Theta$ then one can use the definition

$$X_\tau(\omega) \circeq 0.$$

If one uses the convention that the product of an undefined value with zero is zero, then one can write the definition of the stopped variable X_τ in the following way:

$$X_\tau(\omega) \circeq X(\tau(\omega), \omega)\, \chi(\tau \in \Theta)(\omega).$$

Optional Sampling Theorem

Theorem

If M is a martingale and τ is a stopping time, then the stopped process M^{τ} is also a martingale.

Theorem

An adapted process M is a martingale if and only if, the trajectories of M are right-continuous and they have limits from the left and for any bounded stopping time τ

$$\mathbf{E}\left(M_{\tau}\right)=\mathbf{E}\left(M\left(0\right)\right),$$

where as above $M_{\tau}\left(\omega\right)\stackrel{\circ}{=}M\left(\tau\left(\omega\right),\omega\right).$

Uniformly integrable martingales

When can one drop the condition that τ is bounded? For any T obviously $\tau \wedge T$ is bounded so $\mathbf{E}(M(\tau \wedge T)) = \mathbf{E}(M(0))$. Can we take the limit $T \to \infty$ under the integral sign:

$$\mathbf{E}(M(\tau)) = \mathbf{E}\left(\lim_{T\to\infty} M(\tau \wedge T)\right) = \lim_{T\to\infty} \mathbf{E}(M(\tau \wedge T)) = \mathbf{E}(M(0))?$$

If $\tau < \infty$ and M has an integrable majorant, that is $|M(t)| \leq \xi \in L^1(\Omega)$ then the answer is certainly yes. There is a generalization. There is a special concept, called "uniformly integrable" martingale which generalizes the mentioned situation.

An important case when M is uniformly integrable but there is no integrable majorant is when $M \in \mathcal{H}^2$.

Example

If w is a Wiener process and

$$\tau_a \circeq \inf\left(t : w\left(t\right) \geq a\right)$$

then as w is continuous and as w is unbounded $\tau_a < \infty$ and $w\left(\tau_a\right) = a$. Hence if $a \neq 0$

$$\mathbf{E}\left(w\left(\tau_a\right)\right) = \mathbf{E}\left(a\right) = a \neq \mathbf{E}\left(w\left(0\right)\right) = 0.$$

Observe that

$$w\left(\infty\right) \circeq \lim_{t\to\infty} w\left(t\right)$$

is not well-defined. Observe that w is in $\mathcal{H}^2$ on any finite interval but $w \notin \mathcal{H}^2$ on $\mathbb{R}_+$.

Optional Sampling Theorem for uniformly integrable martingales

Theorem

If M is a uniformly integrable martingale, then

1. *one can extend M from $[0,\infty)$ to $[0,\infty]$,*
2. *the extended process remains a martingale and*
3. $\mathbf{E}(M(\tau)) = \mathbf{E}(M(0))$ *holds for any stopping time τ.*

Ruin probabilities

Example

If $a < 0 < b$ and τ_a and τ_b are the respective first passage times of some Wiener process w, then

$$\mathbf{P}\left(\tau_a < \tau_b\right) = \frac{b}{b-a}, \quad \mathbf{P}\left(\tau_b < \tau_a\right) = \frac{-a}{b-a}.$$

With probability one the trajectories of w are unbounded. Therefore as w starts from the origin the trajectories of w finally leave the interval $[a, b]$. So

$$\mathbf{P}\left(\tau_a < \tau_b\right) + \mathbf{P}\left(\tau_b < \tau_a\right) = 1.$$

If $\tau \doteq \tau_a \wedge \tau_b$ then w^τ is a bounded martingale. Hence one can use the Optional Sampling Theorem. Obviously w_τ^τ is ether a or b, hence

$$\mathbf{E}\left(w_\tau^\tau\right) = a\mathbf{P}\left(\tau_a < \tau_b\right) + b\mathbf{P}\left(\tau_b < \tau_a\right) = \mathbf{E}\left(w^\tau(0)\right) = 0.$$

We have two equations with two unknowns. Solving this system of linear equations, one can easily deduce the formulas above.

Energy Equality

Theorem

If $M \in \mathcal{H}^2$, then

$$\mathbf{E}\left((M(t) - M(s))^2\right) = \mathbf{E}\left(M^2(t)\right) - \mathbf{E}\left(M^2(s)\right)$$

By the properties of the conditional expectation for

$$\Delta \overset{\circ}{=} 2 \cdot \mathbf{E}\left(M(s)(M(s) - M(t))\right)$$

the difference of the two sides

$$\begin{aligned} \Delta &\overset{\circ}{=} 2 \cdot \mathbf{E}\left(M(s) \cdot (M(s) - M(t))\right) = \\ &= 2 \cdot \mathbf{E}\left(M(s) \cdot \mathbf{E}\left(M(s) - M(t) \mid \mathcal{F}_s\right)\right) = \\ &= 2 \cdot \mathbf{E}\left(M(s) \cdot 0\right) = 0. \end{aligned}$$

Energy Equality and the Uniform Integrability

Recall that $M \in \mathcal{H}^2$ if $\|M(t)\|_2 \leq K$. By the Energy Equality the function $t \mapsto \|M(t)\|$ is increasing. As it is bounded it is convergent. Hence for any $\varepsilon > 0$

$$\|M(t) - M(s)\| \leq |\|M(t)\| - \|M(s)\|| < \varepsilon$$

it $t, s \geq N(\varepsilon)$. Hence $(M(t))_t$ is a Cauchy sequence in $L^2(\Omega)$. As $L^2(\Omega)$ is complete it is convergent, that is $M(t) \to M(\infty)$ and the convergence holds in $L^2(\Omega)$. As M is a martingale

$$\mathbf{E}(M(t) \mid \mathcal{F}_s) = M(s).$$

If $t \to \infty$, then as convergence in $L^2(\Omega)$ implies convergence in $L^1(\Omega)$

$$\begin{aligned}\mathbf{E}(M(\infty) \mid \mathcal{F}_s) &= \mathbf{E}\left(\lim_{t\to\infty} M(t) \mid \mathcal{F}_s\right) = \lim_{t\to\infty} \mathbf{E}(M(t) \mid \mathcal{F}_s) = \\ &= \lim_{t\to\infty} M(s) = M(s).\end{aligned}$$