Lecture notes by:
Matthias Kredler
Universidad Carlos III de Madrid
This version: May 2018

1 HJB: The stochastic case

1.1 Brownian Motion

First, we should have a look at how the law of motion for a process is defined in continuous time in a stochastic environment. In order to do this, we first need to define a shock process that is equivalent to the i.i.d. shocks that we normally use in discrete time. The first thing that comes to mind is the following:

$$z_t \sim \text{i.i.d.},\ E(z_t) = 0$$

One has in mind that the random variables z_t and z_s should be independent for any $t \neq s$, however close together s and t are. However, it turns out that this construct is not useful mathematically. For example, if we want to define the equivalent to a random walk, we already run into difficulties, although this is probably the simplest operation one can think of:

$$Z_t = \int_0^t z_s ds = ?$$

It turns out that Z_t cannot be well-defined, not even when we use the Lebesgue integral instead of the Riemann integral. The problem is that z_t is not a measurable function of time. This means that if we want to determine the size of certain sets of type $A_t = \{t : z_t \leq \alpha, \alpha \in \mathbb{R}\}$, these sets can never be approximated well by taking countable unions and intersections of intervals (why?), which is in the end the only things that we're sure how to deal with.[1] But integrating always requires that we put some simple functions that are constant on sets like A_t below the function we want to integrate (in this case Z_t). If we cannot even find sets of type A_t, definitely the integral cannot be well-defined.

[1]Mathematicians in the 19th century have spend a lot of effort on coming up with the most general class of sets on the real line that we can assign a measure to, and it turned out that measurable sets basically have to be the result of an algorithm as described above.

The next idea is not to work with the shock process z_t itself, but with its integral Z_t. If we can impose certain criteria on it to ensure that its increments behave like the white noise we have in mind and it is mathematically handier, then we should be happy to work with this kind of process. As you can imagine, the Brownian Motion B_t will be the result of this effort.

But let's first motivate the construction a bit more. Let's say that we want to create a process W_t whose increments over a unit of time have standardized variance:

$$Var(W_{t+1} - W_t) = 1$$

Then definitely the increments of length $1/n$ have to have variance $1/n$:

$$Var(W_{t+1} - W_t) = \sum_{i=1}^{n} Var(W_{t+\frac{i+1}{n}} - W_{t+\frac{i}{n}}) = nVar(W_{t+\frac{1}{n}} - W_t),$$

where the covariance terms in the first step are zero because we definitely want the increments to be independent, and where for the second step we require that the variance of increments of the same size be equal.

We can also impose another requirement on our process by just making the chopping of $[t, t+1]$ always finer: Since the increments on the chopped pieces are i.i.d. with finite variance, the Central Limit Theorem applies and $W_{t+1} - W_t$ has to be normally distributed — if we make the chopping number n higher and higher, then definitely this must be true.

Also, note that there is nothing peculiar about the size of our chosen interval $[t, t+1]$; our chopping argument is true for an interval of whatever tiny length. This leads us to our next requirement on B_t, which is that its increments should be normally distributed with variance equal to the size of the interval under consideration. Of course, non-overlapping increments should be independent.

It turns out that we can also require any realization of the process B_t to be a continuous function of time without losing anything, so we definitely want to do this to make our life easier. Note that for a continuous function, the sets $A_t = \{t : z_t \leq \alpha, \alpha \in \mathbb{R}\}$ are well-behaved; mathematically speaking, any continuous function is measurable, which will make it possible to work with B_t.

1.2 Ito-Processes/Stochastic Differential Equations

Now we can turn to defining more general stochastic processes by using our shock process B_t. We write this in differential form for the process X_t:

$$dX_t = a(X_t)dt + s(X_t)dB_t, \qquad \text{given } X_0 \text{ and } B_0 = 0 \tag{1}$$

(1) is called a "stochastic differential equation". But what does it actually mean? The dB_t takes the role of the ε_t in discrete time, $s(X_t)$ is a function that specifies how volatile the process is at a certain point in the state space, and $a(X_t)$ gives us a deterministic trend component depending on the state X_t. Often we will determine functions $a(\cdot)$ and $s(\cdot)$ that are inspired by some idea we have in discrete time, and then we will try to start working from that.

The problem is that we still don't know how to actually compute the value of X_t for $t > 0$. The goal should be that we can calculate X_t when somebody gives us a realization B_t of the shock process over an interval $[0, T]$. Since we know well how to think about discrete pieces of time, we will have a look at the following approximation of what we mean by equation (1)[2]

$$X_{t+\Delta t} - X_t = a(X_t)\Delta t + s(X_t)\varepsilon_t, \quad \varepsilon_t \sim \text{i.i.d. } \mathcal{N}(0, \Delta t) \tag{2}$$

This is meant recursively: We compute $a(X_0)$ and $s(X_0)$ with the X_0 from our initial condition. Then, we read off $B_{\Delta t}$ from the realization of B_t what we're given, and compute $X_{\Delta t}$. With this value $X_{\Delta t}$ in hand, we can again evaluate a and s, read off the next increment and compute $X_{2\Delta t}$, and so forth. Then we repeat the same, but this time for a smaller Δt.

In the end, we will — under some technical assumptions — arrive at some limit values X_t for all $t \in [0, T]$; we're lucky, since Prof. Ito has proved this for us. This limit is called the *Ito Integral*, which involves integrating against the Brownian Motion B_t and looks somewhat intimidating at first:

$$X_t = \int_0^t a(X_t)dt + \int_0^t s(X_t)dB_t$$

Everybody who says this is an easy equation is a fool: As we have seen, the values for X_t have to be evaluated underway as a function of the shocks to B_t up to t, and it is not clear at all that we can have something like a closed-form solution to solve for X_t given the previous evolution of B_t .

[2] We have in mind that $a(\cdot)$ and $s(\cdot)$ are sufficiently smooth functions in order to do this.

Actually, if we *have* something like this, we call it a "solution to a stochastic differential equation".[3] Specifically, a solution is a function $g(t, X_0, \{B_s\}_{s\in[0,t]})$ that tells us how to compute X_t given the ingredients X_0 and B_t up to t.

1.3 An example: Geometric Brownian Motion

An example is in order to illustrate what it means to obtain a solution for a stochastic differential equation. Consider the following law of motion:

$$dX_t = \mu X_t dt + \sigma X_t dB_t, \qquad \text{given } X_0 = 1, \tag{3}$$

where μ and σ are constants. This is the stochastic version of a growth process: X_t tends to grow at a certain rate over time, but there are shocks (whose size is proportional to the size of X_t) that disturb this growth. Note that if σ was zero, then the exponential function would be the solution to the resulting deterministic differential equation. It turns out that the following solution exists for (3) in the stochastic case:

$$X_t = X_0 \exp\left[(\mu - \sigma^2/2)t + \sigma B_t\right] \tag{4}$$

Notice that a solution like this makes our life a lot easier: Given B_t, we don't have to do the approximation procedure from equation (2) anymore, but we can just substitute the single number B_t into (4) to obtain a 100%-correct value for X_t. Also, with this solution in hand, we see immediately that X_t is log-normally distributed; we are given the mean and variances for free, we can easily compute quantiles, and even the joint distribution properties for any collection $\{X_1, \ldots, X_n\}$ are relatively easy to obtain.

Look again at equation (4). It is important to be surprised by the term $-\sigma^2/2$ that shows up in front of t: This is the famous Ito-term, which has to do with Jensen's inequality.

To see that we need this term, suppose that $\mu = 0$ and that the term $-\sigma^2/2$ didn't appear, i.e. we suggest a different solution $\hat{X}_t = \hat{g}(B_t) = \exp(\sigma B_t)$ in this particular case. We will now show that this solution $\hat{X}_t$

[3] These solutions can be found using Ito's Lemma, which is something like a rule for taking derivatives of stochastic processes: With this rule, we can take of a function $f(X_t)$ of some stochastic process with respect to time, and then see if the result resembles some stochastic differential equations that we have written down before, having in mind some discrete-time model.

would give us values that are fundamentally at odds with some properties of the stochastic differential equation we wrote down in (3).

If we take any discrete chopping of $[0, t]$ and approximate X_t à la (2), we see that all increments must have expectation zero:

$$E[X_{t+\frac{i+1}{n}} - X_{t+\frac{i}{n}}] = \sigma(X_t)E[B_{t+\frac{i+1}{n}} - B_{t+\frac{i}{n}}] = 0.$$

Hence we have $E[X_t] = \sum_n E[X_{t+(i+1)/n} - X_{t+i/n}] = 0$ for any approximation, so it is impossible that $E[X_t] > 0$ in the limit when $n \to \infty$. However, our candidate solution $\hat{X}_t = \exp(\sigma B_t)$ claims exactly this: The noise term σB_t would make the expectation of $\exp(\sigma B_t)$ grow over time, as a simple result of Jensen's inequality. Note that the correct solution $X_t = -\sigma^2 t/2 + \exp(\sigma B_t)$ corrects for this effect by subtracting the mean of the log-normal distribution of $\exp(\sigma B_t)$.

1.4 The Ito Rule

Solutions to stochastic differential equations are obtained in a way similar to those for deterministic ones: When you have taken the derivatives of many functions in your life, you start to see patterns and you can guess the solution to certain differential equations with some practice: The exponential function is the one whose integral is always equal to its derivative, linear functions have the same slope everywhere and so forth. When people get some solutions that are not obvious at first glance, they put them into books so that other people don't have to go through the ordeal of looking for the solution again.

For the stochastic case, the analogon of taking derivatives is the *Ito Rule*. Suppose we already know a lot about some stochastic process, for example a simple one like the *Brownian Motion with drift*:

$$dX_t = \mu dt + \sigma dB_t,$$

where μ and σ are constants. The solution of this stochastic differential equation is of course $X_t = \mu t + \sigma B_t$. Now, we ask how some newly defined random variable $Y_t \equiv f(X_T)$ which is a function of X_t would behave over time. Suppose we take again $f(\cdot) = \exp(\cdot)$. Then, by looking at $exp(\mu t + \sigma B_t)$, we can deduce a lot of properties of the stochastic process $f(X_t)$, since $\mu t + \sigma B_t$ is normally distributed with mean μt and variance σ^2. For example, we can infer that $\exp(X_t)$ is log-normally distributed, calculate the moments and quantiles, and find out a lot about joint distributions (as described in the previous section).

However, we do not know yet how the process $f(X_t)$ behaves in the sense of a stochastic differential equation as given in (1). It turns out that for every well-behaved (to be precise: twice differentiable) function $f(\cdot)$ and any Ito process as defined in (1), the behavior of $f(X_t)$ on tiny pieces of time can again be described by a stochastic differential equation like (1).[4] We will now derive *Ito's Lemma* (or the *Ito Rule*) which is a simple recipe how to do this. The following derivation is somewhat heuristic; see the book by Oksendahl for a more rigorous treatment.

The basic idea is to study how the function $f(X_t)$ changes over a small time interval $[t, t+r]$ by taking Taylor approximations. Before we start, first observe that if we take r small enough, then X will not move far away from its initial value X_t and we can approximate the stochastic process X_t by

$$dX_t = \bar{a}dt + \bar{s}dB_t$$

for the interval $[t, t+r]$, where $\bar{a} \equiv a(X_t)$ and $\bar{s} \equiv s(X_t)$. For this, it is enough that the functions a and s are continuous.

Similarly, the function $f(X)$ can be approximated well by a second-order Taylor approximation around X_t as long as the state X does not move away too far from its initial level (we will later see why we need to keep terms up to second order):

$$f(X) \simeq f(X_t) + \bar{f}'(X - X_t) + \tfrac{1}{2}\bar{f}''(X - X_t)^2 \qquad \text{for } X \in [X_t - \epsilon, X_t + \epsilon],$$

where we define the constants $\bar{f}' \equiv f'(X_t)$ and $\bar{f}'' \equiv f''(X_t)$.

Now let's chop $[t, t+r]$ into N smaller intervals of size $\Delta t \equiv r/N$ and approximate how $f(X)$ changes over the interval Δt when starting from some level X_τ at time $\tau \in [t, t+r)$:

$$f(X_{\tau+\Delta t}) = f(X_\tau) + \bar{f}'(\bar{a}\Delta t + \bar{s}\Delta B_\tau) + \bar{f}''(\bar{a}\Delta t + \bar{s}\Delta B_\tau)^2,$$

where $\Delta B_\tau \equiv B_{\tau+\Delta t} - B_\tau$ is the increment of Brownian motion over the small time interval $[\tau, \tau + \Delta t]$.[5] Now, concatenate the changes in $f(X)$ over

[4]Mathematically speaking, Ito processes are a family of integrals that are "stable under smooth maps".

[5]Notice that actually we would have to take $f'(X_\tau) = \bar{f}' + \bar{f}''(X_\tau - X_t)$ when taking the second-order approximation of $f(X)$ seriously – follow the rest of the proof why it is not necessary to be this precise once the terms $X_\tau - X_t$ become small!

all the small Δt-intervals in $[t, t+r]$:

$$f(X_{t+s}) = f(X_t) + \sum_{j=1}^{N} \bar{f}'(\bar{a}\Delta t + \bar{s}\Delta B_{t_j}) + \sum_{j=1}^{N} \tfrac{1}{2}\bar{f}''(\bar{s}^2 \Delta B_{t_j}^2 + \underbrace{2\bar{a}\bar{s}\Delta t \Delta B_{t_j}}_{\to 0} + \underbrace{\bar{a}^2 \Delta t^2}_{\to 0}).$$

Now, as we let $N \to \infty$, we see that the last term is of order Δt^2 and vanishes when summed up. Similarly, the second-to-last term is of order $\Delta t^{3/2}$ and vanishes (see your homework). As for the remaining term multiplying $\bar{f}''$, the random variable $\Delta B_{t_j}^2$ is i.i.d. χ^2-distributed with mean Δt and finite variance. By the law of large numbers, we will have $\sum_{j=1}^{N} \Delta B_{t_j}^2 \to N\Delta t = s$ as $N \to \infty$ (see your homework again).

Replacing the bar-notation ($\bar{a}$,$\bar{s}$,$\bar{f}'$,$\bar{f}''$) again by functions evaluated at X_t, we conclude that

$$f(X_{t+s}) - f(X_t) = f'(X_t)\big[a(X_t)r + s(X_t)(B_{t+r} - B_t)\big] + \tfrac{1}{2}f''(X_t)s(X_t)^2 r. \quad (5)$$

We see that the expected change in $f(X)$ over the interval $[t, t+r]$ is given by

$$E_t\big[f(X_{t+r}) - f(X_t)\big] = f'(X_t)a(X_t)r + \tfrac{1}{2}f''(X_t)s(X_t)^2 r.$$

So the expected change in f is partly due to the drift a in X_t (the first term), and partly due to convexity/concavity of f in combination with shocks (the second term – look again at the derivation of this term if the intuition is not clear to your). Additional to the expected change in f, there is the *martingale term* $f'(X_t)s(X_t)(B_{t+r} - B_t)$, which captures the effect of shocks on f.

As $r \to 0$, we re-write (5) as *Ito's Lemma* (or the *Ito Rule*) in differential notation:

$$df(X_t) = \underbrace{f'(X_t)a(X_t)dt}_{\text{drift-induced term}} + \underbrace{f'(X_t)s(X_t)dB_t}_{\text{martingale term}} + \underbrace{\tfrac{1}{2}f''(X_t)s(X_t)^2 dt}_{\text{Ito term}} \quad (6)$$

If you're not sure about how to read this type of notation in any stochastic-calculus problem, always go back to the discrete version (5) that we derived to gain intuition.

The term $f'(X_t)a(X_t)dt$ in (6) clearly captures the effect of the deterministic drift component, and $f'(X_t)s(X_t)dB_t$ represents random motions. The *Ito term* $\frac{1}{2}f''(X_t)s(X_t)^2 dt$ captures the effects of Jensen's inequality: If the function $f(\cdot)$ is convex, then there is an upward movement of $f(X_t)$ in expectation additional to the one induced by the drift $a(X_t)dt$; positive changes in B_t will have a larger impact on $f(X_t)$ then negative movements of the same size, since the slope of $f(X_t)$ is increasing in X_t.

It is a good exercise to apply Ito's Rule to several important functions on the process $dX_t = adt + \sigma dB_t$. Think about the intuition for the resulting stochastic differential equations and the meaning of the single terms in it. If you start with $\exp(X_t)$, you will obtain the example from the previous section, for example.

Finally, we can integrate up (6) over time (that is: for a given realization B_t of Brownian motion) to obtain

$$f(X_t) = \int_0^t f'(X_t)a(X_t)dt + \int_0^t \frac{1}{2}f''(X_t)s(X_t)^2 dt + \int_0^t f'(X_t)s(X_t)dB_t.$$

We can now come back to a problem we faced before: How can we find closed-form solutions to SDEs? It works as follows: We can try and apply Ito's Lemma to different functions $f(\cdot)$ and obtain stochastic differential equations as in (6). Note that then, of course $f(X_t)$ is a solution to the stochastic differential equation (6). If X_t is something we're easily able to evaluate as a function of B_t, as is the case for $X_t = at + \sigma B_t$ and $f(X) = \exp(X)$, we have found a solution to an SDE. This is analogous to taking derivatives of conventional functions and using the result as a solution to an ordinary differential equation (ODE): For $g(X) = \exp(X)$, we have $g'(X) = \exp(X) = g(X)$, so $g(X) = \exp(X)$ solves the ODE $g'(x) = g(x)$ with initial condition $g(0) = 1$.

1.5 The stochastic HJB

Now we have everything in place to look at the HJB for the stochastic case. We just need to give the planner some way to influence the stochastic process; we do this by giving him the chance to choose a function $h(X_t)$ that specifies the policy rule for each given state of the system X_t, and by letting this control influence the law of motion for X_t:

$$dX_t = a[X_t, h(X_t)]dt + s[X_t, h(X_t)]dB_t \tag{7}$$

This means that the drift as well as the volatility of X_t may be influenced by the control $h(X_t)$. Of course this general case encompasses the easier case where the planner cannot choose his exposure to certain shocks, in which case $s[\cdot]$ is only a function of X_t but not of $h(X_t)$.

Look at the following example to understand what this means exactly: Suppose that an agent can either put his assets A into a safe investment that increases his wealth at a certain rate r or invest them into a risky asset which increases wealth at a stochastic rate, whose average is higher than r:

$$dA_t = h(A_t)A_t r dt + [1 - h(A_t)]A_t[\mu dt + \sigma dB_t],$$

where A_t is the level of asset holdings at t and the function $0 \le h(A_t) \le 1$ specifies which fraction of his assets the agents wants put into the riskless bond at a certain level of A_t. Notice that when we fix a certain policy $h(\cdot)$, equation (7) gives us a stochastic differential equation that determines the behavior of A_t.

Now we turn to the optimal choice of $h(\cdot)$: Choosing among all functions $h(\cdot)$ in a certain class, we want to find the one that maximizes a certain reward function $u(\cdot)$ discounted over time, i.e.

$$\begin{aligned} &\max_{h(\cdot)\in\mathcal{H}} \quad E_0 \int_0^\infty e^{-\beta t} u(X_t) dt && (8) \\ &\text{where} \quad X_t = \int_0^t a[X_t, h(X_t)]dt + \int_0^t s[X_t, h(X_t)]dB_t, \end{aligned}$$

where $\mathcal{H}$ is some function space that I am too lazy to specify exactly.

We will take an approach as for the discrete-time HJB and ask which fixed value $\bar{h}$ we should choose given that we are in a certain state X_t. Again, we will argue that in the problems we are interested in, the optimal policy should be a continuous function of the state X_t, the process X_t shouldn't make crazy jumps etc. Then we can look — as in the discrete case — for an optimal $\bar{h}$ to get the functional equation for the value function:

$$v(X_t) \simeq \max_{\bar{h}} \left\{ \int_t^{t+\varepsilon} e^{-\beta(s-t)} u(X_t, \bar{h}) ds + E_t\big[e^{-\beta\varepsilon} v(X_{t+\varepsilon})\big] \right\} \qquad (9)$$

Note that I have fixed X_t in the first term pertaining to the instantaneous reward $u(\cdot)$: As in the deterministic case, we argue that X_t will barely move over a very short time and that we don't loose much accuracy by doing this.

Now we use what we know from the Ito rule (6) to calculate how $v(X_{t+\varepsilon}$

evolves:

$$
\begin{aligned}
v(X_{t+\varepsilon}) =& v(X_t) + \int_t^{t+\varepsilon} v'(X_s)a(X_s,\bar{h})ds + \frac{1}{2}\int_t^{t+\varepsilon} v''(X_s)s(X_s,\bar{h})^2 ds + \\
&+ \int_t^{t+\varepsilon} v'(X_s)s(X_s,\bar{h})dB_s \\
\simeq& v(X_t) + \left[v'(X_t)a(X_t,\bar{h}) + \frac{1}{2}v''(X_t)s(X_t,\bar{h})^2\right]\varepsilon + \\
&+ v'(X_t)s(X_t,\bar{h}) \underbrace{\int_t^{t+\varepsilon} dB_s}_{=B_{t+\varepsilon}-B_t}
\end{aligned}
\tag{10}
$$

For the approximation in the second step we use again that for tiny ε the movements in the state X_t are (almost surely) so small that we get very close to the truth by just putting $X_{t+s} = X_t$ when we evaluate the functions $a(\cdot)$, $s(\cdot)$, $v'(\cdot)$ and $v''(\cdot)$. Now we plug (10) into (9) and observe that $E_t[B_{t+\varepsilon} - B_t] = 0$. As in the deterministic case, we can then take the term $e^{-\beta\varepsilon}v(X_t)$ to the left-hand side, divide everything by ε and take the limit $\varepsilon \to 0$ to obtain the *stochastic Hamilton-Jacobi-Bellman (HJB) equation*:

$$
\beta v(X_t) = \max_{\bar{h}} \left\{ u(X_t,\bar{h}) + v'(X_t)a(X_t,\bar{h}) + \frac{v''(X_t)s(X_t,\bar{h})^2}{2} \right\} \tag{11}
$$

Note that the only point where the stochastic nature of the problem enters into the HJB is in the last term, which involves the second derivative of the value function. The intuition for this term again has to do with Jensen's inequality: If the value function is concave, one will not like the effects from random movements over the state space. If possible, the planner will try to keep volatility low by choosing $\bar{h}$ such that $s(\cdot,\bar{h})$ is small.

For the sake of completeness, we also have a look at the HJB for the generalized problem where we are given a terminal value. Actually, a more general case can be stated than we had in the deterministic case: The process need not be terminated at a fixed time T for all possible histories; instead, the planner could be stopped at different times for different histories of the world. However, stating all this mathematically is somewhat tedious, so just bear in mind the simple case where we have a "bequest function" $v_T(x_T)$ for some fixed terminal point T in time.

Then, the (non-stationary) HJB is:

$$-\frac{\partial v(t,x)}{\partial t} = \max_{\bar{h}} \left\{ e^{-\beta t} u(x,\bar{h}) + \frac{\partial v(x,t)}{\partial x} a(x,\bar{h}) + \frac{\partial^2 v(x,t)}{\partial x^2} \frac{s(x,\bar{h})^2}{2} \right\} \quad (12)$$

Of course we have $v(T,x)$ as a boundary condition. This, again, is a *partial differential equation*, as was the case for the deterministic case.