Bus Ad 239B–Spring 2003
Lecture Notes, Part IV

33 Term Structure Models: Overview

We now turn to Chapter 7 of Nielsen, which concerns term structure models. We will tread fairly lightly; in particular, we will skip all discussion of forward rates.

If the main point of our work so far has been the pricing of options and other derivatives, the point of this material is to set up machinery to look for small arbitrages in the term structure of interest rates. Our models begin with an exogenously-specified process describing risk-free interest rates r. We will see that the Vasicek Model and Merton model are essentially indistinguishable in the short-run dynamics of the risk-free interest rate. Each of these two models completely determines the whole term structure of interest rates as a function of the current value of r and certain parameters. In these two models, it is assumed the parameter values are constant; under this assumption, the parameters can be econometrically estimated. Given the parameter values, the current value of r completely determines the yield curve. Empirically, the shape of the yield curve changes over time, and in particular is *not* determined solely by r, but also by expectations over future interest rates and inflation. Thus, these two models do not provide a good *explanation* of the term structure of interest rates, nor do they successfully predict future interest rates.

In the Extended Vasicek Model, the parameters are allowed to be deterministic functions of time. This additional freedom allows one to choose the parameters to fit essentially any yield curve; one calibrates the parameters of the model at each future date to give the model's predicted yield curve exactly the curvature in the observed yield curve. Thus, the Extended Vasicek Model cannot be said to explain the term structure either; since it is able to replicate any yield curve, it cannot hope to predict what the yield curve should look like. However, the Extended Vasicek Model is very useful for finding small arbitrages in the term structure, for example when the yield on one bond deviates from the predicted relationship to the yields on bonds of similar maturity.

Thus, our main goal is to develop the Extended Vasicek Model. We do the Vasicek and Merton models first, to develop intuition in a simpler setting. We will see that the analysis of the Extended Vasicek Model follows closely, with small changes, the analysis of the Vasicek Model.

34 Term Structure Notation

Section 7.1 of Nielsen is primarily devoted to notation. We will tread lightly.

If $0 \leq s \leq t$, $P(s;t)$ denotes the price (in other words, the martingale value) at time s of a zero-coupon bond paying off 1 at its maturity date t; by definition, zero-coupon bonds pay nothing at any time other than the maturity date. The continuously compounded yield is

$$\begin{aligned} R(s;t) &= \frac{1}{t-s} \ln\left(\frac{1}{P(s;t)}\right) \\ &= -\frac{\ln P(s;t)}{t-s} \\ R(s;t) > 0 &\Leftrightarrow P(s;t) < 1 \\ P(s;t) &= e^{-(t-s)R(s;t)} \end{aligned}$$

The zero-coupon yield curve or term structure of interest rates is the function

$$\tau \rightarrow R(s; s+\tau)$$

which maps $[0,\infty)$ to $\mathbf{R}$.

35 The Vasicek Model

Let W be a 1-dimensional Wiener Process (so $K = 1$), and let r_0, $a > 0$, b, $\sigma > 0$ be constants. In the Vasicek Model, the risk-free interest rate r follows an Ornstein-Uhlenbeck Process

$$dr = a(b-r)\,dt + \sigma\,dW, \;\; r(0) = r_0$$

Recall this says that the interest-rate is mean-reverting to b, but is being kicked away from b by the changes dW in the Wiener process.

Remark 35.1 Interest rates do exhibit mean reversion, so it would not be appropriate to model r or the value M of the money-market account in the way we have modelled stock prices. However, there are several problems in taking the Vasicek Model seriously as a model of interest rates:

1. In the Vasicek Model, r will certainly be negative at some times. However, since individuals always have the option of taking currency and putting it in a safe deposit box, which yields a zero interest rate, negative interest rates cannot be sustained in practice.

2. In the Vasicek Model, the mean-reversion target b stays constant; we know in practice that interest rates stay high or low over long periods, in particular in response to high or low rates of inflation. It probably makes more sense to think of the Vasicek Model as a model of real interest rates.

3. The risk-free interest rate is really a controlled rate. In the United States, the Federal Reserve's Open Market Committee effectively sets r, and so it may be predictable in a way that the future movements of stock prices are not, if markets are efficient.

4. In the Vasicek Model, the market price of risk λ is assumed to be a positive constant. We use r and λ to form a state price process, which we take as a primitive to price the zero coupon bonds. From an equilibrium perspective, this is unrealistic. Changes in securities prices and economic conditions affect individuals' wealths, and wealth effects alter individuals' willingness to bear risk.

As noted in the last remark, we assume that λ is a positive constant, and let

$$\Pi = \Pi(0)\eta[-r, -\lambda]$$

We set the time interval to be $\mathcal{T} = [0, T]$. Since λ is a constant, there is a probability measure Q with density $\eta[0, -\lambda](T)$ with respect to P. We will interpret Q as a risk-adjusted measure to price the zero-coupon bonds. Recall that

$$W^{\lambda}(t) = W(t) + \lambda t$$

is a Wiener Process under Q. If we let

$$\bar{r} = b - \frac{\sigma\lambda}{a}$$

so

$$\sigma\lambda = a(b - \bar{r})$$

we have

$$\begin{aligned}
dr &= a(b-r)\,dt + \sigma\,dW \\
&= a(b-r)\,dt + \sigma(dW^\lambda - \lambda dt) \\
&= [a(b-r) - a(b-\bar{r})]\;dt + \sigma dW^\lambda \\
&= a(\bar{r} - r)\,dt + \sigma\,dW^\lambda \\
r(t) &= \bar{r} + e^{-a(t-s)}(r(s) - \bar{r}) + \sigma\int_s^t e^{-a(t-u)}\,dW^\lambda
\end{aligned}$$

Thus, r is an Ornstein-Uhlenbeck Process under Q, with the same dispersion σ and speed of adjustment a as under P, *but with a different mean-reversion target* $\bar{r} < b$. This seems weird; our intuition is that changing the probabilities might alter the rate of mean-reversion, but shouldn't alter the target. This intuition, which seems to be based on the idea that r has to get to the mean-reversion target eventually, is incorrect. The Ornstein-Uhlenbeck Process doesn't converge to b under P or to $\bar{r}$ under Q. Since Q moves probability from the higher paths of W to the lower paths of W, the mean-reversion target $\bar{r}$ under Q is necessarily lower than the target b under P. Another way to see this is to note that W^λ has positive drift under P, so you need to set the mean-reversion level with respect to W^λ lower than the level with respect to W to compensate for the drift.

In addition to zero-coupon bonds, there is a money-market account

$$M(0) = 1, \;\; M(t) = e^{\int_0^t r(u)\,du}$$

If $s \le t$, set

$$I(s;t) = \int_s^t r(u)\,du$$

so

$$\begin{aligned}
M(t) &= M(s)e^{I(s;t)} \\
\frac{M(s)}{M(t)} &= e^{-I(s;t)}
\end{aligned}$$

Proposition 35.2 *Suppose that Y is a contingent claim, Y is $\mathcal{F}_t$-measurable, and $\frac{Y}{M(t)} \in L^1(Q)$. The Martingale Value Process of Q is*

$$\begin{aligned} V(s) &= E_Q\left(\left.\frac{YM(s)}{M(t)}\right|\mathcal{F}_s\right) \\ &= E_Q\left(\left.Ye^{-I(s;t)}\right|\mathcal{F}_s\right) \end{aligned}$$

Proof:

$$\begin{aligned} \frac{V(Y;\Pi)(s)}{M(s)} &= V\left(\frac{Y}{M(t)};Q\right) \\ V(Y;\Pi)(s) &= M(s)V\left(\frac{Y}{M(t)};Q\right) \\ &= M(s)E_Q\left(\left.\frac{Y}{M(t)}\right|\mathcal{F}_s\right) \\ &= E_Q\left(\left.\frac{YM(s)}{M(t)}\right|\mathcal{F}_s\right) \end{aligned}$$

since $M(s)$ is $\mathcal{F}_s$-measurable. ■

Corollary 35.3 *If $\frac{1}{M(t)} \in L^1(Q)$, the Martingale Value of a zero-coupon bond maturing at time t is*

$$\begin{aligned} P(s;t) &= E_Q\left(\left.\frac{M(s)}{M(t)}\right|\mathcal{F}_s\right) \\ &= E_Q\left(\left.e^{-I(s;t)}\right|\mathcal{F}_s\right) \end{aligned}$$

Proposition 35.4 (Proposition 7.3 in Nielsen) *For $0 \le s \le t$,*

$$I(s;t) = (t-s)\bar{r} + \frac{1}{a}\left(1-e^{-a(t-s)}\right)(r(s)-\bar{r}) + \frac{\sigma}{a}\int_s^t \left(1-e^{-a(t-v)}\right)dW^\lambda(v)$$

Proof:

$$\begin{aligned} I(s;t) &= \int_s^t r(u)\,du \\ &= \int_s^t \left(\bar{r} + e^{-a(u-s)}(r(s)-\bar{r}) + \sigma\int_s^u e^{-a(u-v)}\,dW^\lambda(v)\right)du \end{aligned}$$

$$
\begin{aligned}
&= (t-s)\bar{r} + \frac{e^{-a(u-s)}}{-a}(r(s)-\bar{r})\Big|_s^t + \sigma \int_s^t \int_s^u e^{-a(u-v)}\, dW^\lambda(v)\, du \\
&= (t-s)\bar{r} + \left(-\frac{e^{-a(t-s)}}{a} + \frac{1}{a}\right)(r(s)-\bar{r}) + \sigma \int_s^t \int_v^t e^{-a(u-v)}\, du\, dW^\lambda(v) \\
&= (t-s)\bar{r} + \left(\frac{1-e^{-a(t-s)}}{a}\right)(r(s)-\bar{r}) - \frac{\sigma}{a}\int_s^t \left(e^{-a(u-v)}\Big|_v^t\right) dW^\lambda(v) \\
&= (t-s)\bar{r} + \left(\frac{1-e^{-a(t-s)}}{a}\right)(r(s)-\bar{r}) + \frac{\sigma}{a}\int_s^t \left(1-e^{-a(t-v)}\right) dW^\lambda(v)
\end{aligned}
$$

■

Corollary 35.5 *Conditional on $\mathcal{F}_s$, $I(s;t)$ is normally distributed under Q, with conditional mean*

$$
\begin{aligned}
E_Q\left(I(s;t)|\mathcal{F}_s\right) &= (t-s)\bar{r} + \frac{1}{a}\left(1-e^{-a(t-s)}\right)(r(s)-\bar{r}) \\
&= m(r(s), t-s)
\end{aligned}
$$

where

$$
m(r,\tau) = \tau\bar{r} + \frac{1}{a}\left(1-e^{-a\tau}\right)(r-\bar{r})
$$

and conditional variance

$$
\mathrm{Var}\left(I(s;t)|\mathcal{F}_s\right) = v(t-s)
$$

where

$$
v(\tau) = \frac{\sigma^2}{2a^3}\left(4e^{-a\tau} - e^{-2a\tau} + 2a\tau - 3\right)
$$

In particular, $v(\tau)$ does not depend on λ or on $r(s)$.

Proof: The statement about the conditional mean is immediate from Proposition 35.4. Now, we turn to the conditional variance:

$$
\begin{aligned}
\mathrm{Var}_Q\left(I(s;t)|\mathcal{F}_s\right) &= \frac{\sigma^2}{a^2}\int_s^t \left(1-e^{-a(t-v)}\right)^2 dv \\
&= \frac{\sigma^2}{a^2}\int_s^t \left(1-2e^{-a(t-v)} + e^{-2a(t-v)}\right) dv
\end{aligned}
$$

$$
\begin{aligned}
&= \frac{\sigma^2}{a^2}\left(v - \frac{2e^{-a(t-v)}}{a} + \frac{e^{-2a(t-v)}}{2a}\Bigg|_s^t\right) \\
&= \frac{\sigma^2}{a^2}\left((t-s) - \frac{2}{a}\left(1 - e^{-a(t-s)}\right) + \frac{1}{2a}\left(1 - e^{-2a(t-s)}\right)\right) \\
&= \frac{\sigma^2}{2a^3}\left(4e^{-a(t-s)} - e^{-2a(t-s)} + 2a(t-s) - 3\right) \\
&= v(t-s)
\end{aligned}
$$

■

Proposition 35.6 $P(s;t) = P(r(s), t-s)$, *where*

$$P(r,\tau) = e^{-m(r,\tau)+v(\tau)/2}$$

Proof: This is an exercise in the expectation of a lognormal random variable, and is left to the reader.■

Remark 35.7 In Nielsen's notation, symbols following semi-colons denote maturity dates, as in $P(s;t)$, the price at time s of a zero-coupon bond maturing at time t. There is no semi-colon in $P(r,\tau)$ because τ is not the maturity *date*; it is the *length* of time (from the present) until the maturity date.

Remark 35.8 As promised in Section 33, Proposition 35.6 provides a complete specification of the term structure, assuming that the market price of risk is a constant, and given constant parameters σ, a and b (or $\bar{r}$) of the interest rate process r. The fact that the term structure empirically is *not* solely determined by $r(s)$ indicates that the model is incorrect, most likely because the parameters are in fact *not* constant.

We now turn to the stochastic dynamics of the bond price $P(s;t)$.

Definition 35.9 Let

$$
\begin{aligned}
B(\tau) &= \frac{1}{a}\left(1 - e^{-a\tau}\right) \\
A(\tau) &= (\tau - B(\tau))\bar{r} - \frac{v(\tau)}{2}
\end{aligned}
$$

so

$$
\begin{aligned}
P(r,\tau) &= e^{A(\tau)-B(\tau)r} \\
P(r(s), t-s) &= e^{-A(t-s)-B(t-s)r(s)}
\end{aligned}
$$

Proposition 35.10 *With t fixed and s as the time variable, $P(s;t)$ satisfies the stochastic differential equation*

$$\begin{aligned}\frac{dP(s;t)}{P(s;t)} &= \frac{dP(r(s),t-s)}{P(r(s),t-s)} \\ &= r(s)\,ds - \left(\frac{1-e^{-a(t-s)}}{a}\right)\sigma dW^{\lambda} \\ &= r(s)\,ds - B(t-s)\sigma\,dW^{\lambda}\end{aligned}$$

Proof:

$$\begin{aligned}P(s;t) &= P(r(s),t-s) \\ &= e^{-m(r(s),t-s)+v(t-s)/2} \\ &= e^{-\tau\bar{r}-\frac{1}{a}\left(1-e^{-a\tau}\right)(r-\bar{r})+v(t-s)/2} \\ \frac{\partial P}{\partial r} &= \left(e^{-\tau\bar{r}-\frac{1}{a}\left(1-e^{-a\tau}\right)(r-\bar{r})+v(t-s)/2}\right)\left(-\frac{1-e^{-a\tau}}{a}\right) \\ &= -P(s;t)\frac{1-e^{-a\tau}}{a} \\ \frac{dP(s;t)}{P(s;t)} &= -\frac{1-e^{-a\tau}}{a}\,dr\end{aligned}$$

Thus, the relative dispersion of $P(s;t)$ is

$$\left|-\frac{1-e^{-a\tau}}{a}\sigma\right| = \frac{1-e^{-a\tau}}{a}\sigma$$

Under the risk-adjusted probability measure Q, the relative drift of the bond price P must equal the risk-free interest rate r, so

$$\begin{aligned}\frac{dP(s;t)}{P(s;t)} &= \frac{dP(r(s),t-s)}{P(r(s),t-s)} \\ &= r(s)\,ds - \left(\frac{1-e^{-a(t-s)}}{a}\right)\sigma\,dW^{\lambda} \\ &= r(s)\,ds - B(t-s)\sigma dW^{\lambda}\end{aligned}$$

■

One can estimate the parameters of the model econometrically; Nielsen reports that Chan, Karolyi, Longstaff and Sanders [3] found $a = 0.18$, $\sigma = 0.02$, and $\bar{r} = 0.07$. Using these parameters, Nielsen sketches (Figure 7.2) two price curves corresponding to $r(0) = 0.03$ (3% initial interest rate) and $r(0) = 0.11$ (11% initial interest rate). The graph of $P(r, \tau)$ is a smooth function of τ because r denotes a fixed initial interest rate; changing τ changes the interval to maturity but not the initial interest rate. The price $P(s; t)$ will follow a jagged path because of the changes in $r(s)$; at any given s, it will lie on the curve $P(r(s), t - s)$.

Note that the pricing formula exhibits pull-to-par: $P(r, t - s) \to 1$ as $s \to t$ for fixed r, and $P(s; t) \to 1$ with probability 1 as $s \to t$.

Now, we consider the price $P(r(s), \tau)$ of a zero-coupon bond with s varying, but a fixed interval τ to maturity.

$$
\begin{aligned}
P(r(s), \tau) &= e^{-A(\tau) - B(\tau) r(s)} \\
d \ln P &= -B(\tau)\, dr(s) \\
&= -B(\tau) a(\bar{r} - r(s))\, ds - B(\tau)\sigma\, dW^{\lambda}
\end{aligned}
$$

The relative dispersion is

$$
-B(\tau)\sigma = -\frac{\sigma}{a}\left(1 - e^{-a\tau}\right)
$$

which is the same as the dispersion of $P(s; t)$, where the maturity date (rather than the time to maturity) is fixed. However, because $P(r(s), \tau)$ is *not* a fixed security (the maturity *date* is constantly receding into the future as s advances), there is no reason that the relative drift of $P(r(s), \tau)$ should equal the risk-free interest rate.

Proposition 35.11 *Under Q,*

$$
\frac{dP(r(s), \tau)}{P(r(s), \tau)} = \left(-B(\tau) a(\bar{r} - r(s)) + \frac{B(\tau)^2 \sigma^2}{2}\right) ds - B(\tau)\sigma\, dW^{\lambda}
$$

while under P,

$$
\frac{dP(r(s), \tau)}{P(r(s), \tau)} = \left(-B(\tau) a(b - r(s)) + \frac{B(\tau)^2 \sigma^2}{2}\right) ds - B(\tau)\sigma\, dW
$$

In particular, the stochastic differential equation with respect to P, W and b is the same as the stochastic differential equation with respect to Q, W^{λ} and $\bar{r}$.

Proof: The proof is left as an exercise.■

36 Risk-Adjusted Probability as Primitive

In the rest of his treatment of term structure models, Nielsen takes the risk-adjusted measure Q and a Wiener Process (with respect to Q) $\hat{W}$ as primitives. $\hat{W}$ corresponds to W^λ, but it is not *constructed* from λ because we don't have λ or for that matter the true probability measure P in the model. For the Vasicek Model, we have

$$\begin{aligned} r(t) &= \bar{r} + e^{-at}(r_0 - \bar{r}) + \sigma e^{-at} \int_0^t e^{au}\, d\hat{W}(u) \\ dr &= a(\bar{r} - r)\, dt + \sigma\, d\hat{W} \\ \frac{dP(r(s), t-s)}{P(r(s), t-s)} &= r(s)\, ds - \sigma B(t-s) d\hat{W}(s) \\ \frac{dP(r(s), \tau)}{P(r(s), \tau)} &= \left(-B(\tau) a(\bar{r} - r(s)) + \frac{1}{2} B(\tau)^2 \sigma^2 \right) ds - B(\tau)\sigma\, d\hat{W}(s) \end{aligned}$$

If we know λ, we can recover the dynamics under the true probabilities. For example, if λ is a constant, we define P to have Radon-Nikodym derivative

$$\frac{1}{\eta[0, -\lambda](T)} = \eta[\lambda^2, \lambda](T)$$

with respect to Q and

$$W(t) = \hat{W}^{-\lambda}(t) = \hat{W}(t) - \lambda t$$

which is a Wiener Process under P. Let

$$b = \bar{r} + \frac{\sigma \lambda}{a}$$

Then

$$\begin{aligned} \frac{dP(r(s), t-s)}{P(r(s), t-s)} &= (r(s) - a(b - \bar{r})B(t-s))\; ds - \sigma B(t-s)\, dW(s) \\ \frac{dP(r(s), \tau)}{P(r(s), \tau)} &= \left(-B(\tau) a(b - r(s)) + \frac{B(\tau)^2 \sigma^2}{2} \right) ds - B(\tau)\sigma\, dW(s) \end{aligned}$$

37 Yields in the Vasicek Model

The continuously compounded rate of return on the zero-coupon bond is

$$\begin{aligned}
R(s;t) &= \frac{1}{t-s}\ln\left(\frac{1}{P(r(s),t-s)}\right)\\
&= \frac{1}{t-s}\left(A(t-s)+B(t-s)r(s)\right)\\
&= R(r(s),t-s)
\end{aligned}$$

where

$$\begin{aligned}
R(r,\tau) &= \frac{1}{\tau}\ln\left(\frac{1}{P(r,\tau)}\right)\\
&= \frac{1}{\tau}\left(m(r,\tau)-\frac{v(\tau)}{2}\right)\\
&= \frac{1}{\tau}\left(A(\tau)+B(\tau)r\right)
\end{aligned}$$

Notice that $R(s;t)$ is an affine function of $r(s)$; since the (unconditional) distribution of $r(s)$ is normal, so is the distribution of $R(s;t)$.

Nielsen's Figure 7.10 shows two yield curves, corresponding to initial interest rates $r_0 = 0.03$ and $r_0 = 0.11$, with the same parameters $\bar{r} = 0.07$, $a = 0.18$, and $\sigma = 0.02$ used in Figure 7.2. If $r_0 = 0.11$, the initial interest rate is above the mean-reversion target 0.07, so we expect interest rates to decline; hence, the price of the bond is above the price discounted at 11% interest, so the yield is below 11%; conversely, the yield for $r_0 = 0.03$ is above 3%. For long maturities, the current interest rate has very little effect on the yield, *which is empirically false*; it may be more reasonable for inflation-protected securities.

The average interest rate over the period $[s,t]$ is

$$\begin{aligned}
r_a &= \frac{1}{t-s}\int_s^t r(u)\,du\\
&= \frac{1}{t-s}I(s;t)\\
&= \frac{1}{t-s}\ln\left(\frac{M(t)}{M(s)}\right)
\end{aligned}$$

so r_a is the continuously compounded rate of return on the Money-Market account over the time interval $[s,t]$.

Proposition 37.1 *Conditional on $\mathcal{F}_s$, r_a is normal with conditional mean*

$$\begin{aligned} E_Q\left(r_a|\mathcal{F}_s\right) &= \frac{1}{t-s}m(r(s),t-s) \\ &= R(r(s),t-s)+\frac{v(t-s)}{2(t-s)} \end{aligned}$$

and conditional variance

$$\mathrm{Var}_Q\left(r_a|\mathcal{F}_s\right)=\frac{v(t-s)}{t-s}$$

Proof: Immediate. ■

Remark 37.2 Proposition 37.1 is remarkable, and very disturbing. It says that the expected yield on the Money-Market Account over the period $[s,t]$ exceeds the expected yield on the bond by the yield risk premium $\frac{v(t-s)}{2(t-s)}$. One of the best-established facts in Finance is that, on average, the yield on long bonds exceeds the long-run yield on money-market accounts. In the Vasicek Model, the zero-coupon long bond is riskless over $[s,t]$, whereas the return on the money-market account, which is instantaneously riskless, is risky over the period $[s,t]$. Thus, the Money-Market Account must pay a yield risk premium; this would also be true in an equilibrium pricing model, taking the dynamics of r as a primitive, if the utility payoff of the long bond were risk-free. The yield risk premium will be even higher if one uses the true probability P instead of Q, since Q shifts probability toward the lower branches of W, and hence toward lower short-term interest rates. In practice, the long bond is not risk-free, because of inflation risk. In an equilibrium model, pricing of securities is in units of marginal utility, and in these units, the payoff of the long bond is *not* risk-free because of inflation. Thus, it appears again that the Vasicek Model may be a better model for real interest rates than for nominal interest rates.

Proposition 37.3 (Proposition 7.4 in Nielsen) *The yield risk premium is*

$$\frac{v(t-s)}{2(t-s)}=\frac{\sigma^2}{4a^3(t-s)}\left(4e^{-a(t-s)}-e^{-2a(t-s)}+2a(t-s)-3\right)$$

$\frac{v(\tau)}{2\tau}$ *is strictly increasing in* τ*, and*

$$\begin{aligned} \lim_{t \to 0} \frac{v(\tau)}{2\tau} &= 0 \\ \lim_{t \to \infty} \frac{v(\tau)}{2\tau} &= \frac{\sigma^2}{2a^2} \end{aligned}$$

Proof: This is left as an exercise.■

Read the rest of Section 7.5 on your own.

38 The Merton Model

The Merton Model has no mean-reversion. It specifies

$$r(t) = r_0 + \alpha t + \sigma \hat{W}(t), \;\; r_0, \alpha, \sigma \text{ constant}, \sigma > 0$$

In the short-run dynamics of r, the Merton Model is indistinguishable from the Vasicek Model; simply take $\alpha = a(\bar{r} - r_0)$, and note that r is continuous in the Vasicek Model. However, the long-run dynamics and the pricing of long bonds is very different in the two models. As Nielsen shows in a straightforward calculation, in the Merton Model, we have

$$P(s;t) = e^{-r(s)(t-s) - \alpha(t-s)^2/2 + \sigma^2(t-s)^3/6}$$

For $t - s$ sufficiently large, the cubic term $\sigma^2(t-s)^3/6$ term will dominate, so $P(s;t) \to \infty$ as $t \to \infty$; a bond that pays off \$1 in 5000 years is worth much more than \$1 today! I still find this puzzling. In class, I suggested that r can go arbitrarily negative in the future, and the long bond is very valuable as insurance against that possibility, but I don't think this is a complete explanation. Note that by the Long Run Law of the Iterated Logarithm, if $\alpha > 0$, then for

$$Q\left(\left\{\omega : \exists_{t_0(\omega)} \;\; t > t_0(\omega) \Rightarrow r(t,\omega) > 0\right\}\right) = 1 \tag{1}$$

Also, conditional on $\mathcal{F}_s$, $r(s+\tau)$ is normally distributed, with mean $r(s) + \alpha\tau$ and variance $\sigma^2\tau$, so

$$E_Q\left(e^{-r(t)}|\mathcal{F}_s\right) = e^{-r(s) - \alpha\tau + \sigma^2\tau/2} \to 0 \tag{2}$$

provided $\alpha > \sigma^2/2$. However,

$$P(s;t) = E_Q\left(e^{-\int_s^t r(u)\,du}\middle|\mathcal{F}_s\right)$$

and it appears that Jensen's inequality plays a key role in the long-run behavior of $P(s;t)$, despite Equations (1) and (2). If I get any more insight on this, I will let you know.

Both the Vasicek and Merton Models are best thought of as short-run models; their ability to explain the yield curve (and in particular, changes in the shape of the yield curve over time) is very limited. As noted above, in the short run, we can't distinguish between them.

39 The Extended Vasicek Model

The Extended Vasicek Model is a generalization of both the Vasicek and Merton Models. However, the key difference is not the unification of the two models, but rather the fact that the parameters are allowed to be deterministic functions of time. The advantage of this is that it allows the parameters to be determined by calibration to whatever the observed yield curve happens to be at any time. The disadvantage is that, since the Extended Vasicek Model can reproduce essentially any yield curve, it cannot possibly *explain* why the yield curve has a given shape at a particular time.

The assumption that the parameters are deterministic functions of time, rather than random processes, is certainly incorrect, but in a sense that is beside the point. Deterministic functions of time allow enough freedom to reproduce essentially any yield curve, so more freedom is not needed in order to calibrate the model to the observed yield curve. Calibration of the model allows one to detect small arbitrages, as the relationship among yields of securities of similar maturities varies.

At the same time, the parameter values $\alpha(t)$, $a(t)$ and $\sigma(t)$ that one obtains at time s through calibration should not really be thought of as predictions of the values those parameters will actually hold when time t rolls around. For starters, the actual values of those parameters are surely random, viewed from the perspective of time s, so the calibrated values could at most be point estimates. But if we try to view them as point estimates, there's no reason to think the expectation of the actual value will equal the calibration-predicted value. The yield curve emerges from the parameters in a nonlinear way, and the calibrated parameter values are presumably expectations of a nonlinear function of the random future values. The calibrated value is what the parameter value would have to be if in fact it were a deterministic function of time, which it isn't.

The key to the calibration will be the curvature of the yield curve at each future time t.

In the model,

$$dr = (\alpha - ar)\,dt + \sigma\,d\hat{W}, \quad r(0) = r_0$$

where r_0 is constant, and α, a and σ are determintic functions of time, $\sigma \in \mathcal{L}^2$, $\alpha, a \in \mathcal{L}^1$. The model reduces to the Vasicek Model if we impose the constraints that α, a and σ are constant and $a > 0$; we get a mean-reversion target of $\bar{r} = \alpha/a$. The model reduces to the Merton Model if we impose the constraint that $a = 0$ and α and σ are constants. In the short-run dynamics of r, we can't distinguish the Extended Vasicek Model from the Vasicek and Merton Models, provided α, a and σ are continuous.

Proposition 39.1 (Proposition 3.7 in Nielsen) *Let* $K(t) = \int_0^t a\,ds$. *Then*

$$\begin{aligned} r(t) &= e^{-K(t)}\left(r_0 + \int_0^t e^{K(s)}\alpha\,ds + \int_0^t e^{K(s)}\sigma\,d\hat{W}(s)\right) \\ &= e^{-K(t)}\left(e^{K(s)}r(s) + \int_s^t e^{K(u)}\alpha\,du + \int_s^t e^{K(u)}\sigma\,d\hat{W}(u)\right) \end{aligned}$$

As in the Vasicek model, we define

$$\begin{aligned} I(s;t) &= \int_s^t r(u)\,du \\ B(s;t) &= \int_s^t e^{-K(u)+K(s)}\,du \\ &= e^{K(s)}\int_s^t e^{-K(u)}du \\ A(s;t) &= \int_s^t e^{-K(u)}\int_s^u e^{K(x)}\alpha\,dx\,du - \frac{v(s;t)}{2} \end{aligned}$$

Proposition 39.2 (Proposition 7.5 in Nielsen) *For* $0 \le s \le t$,

$$I(s;t) = B(s;t)r(s) + \int_s^t e^{-K(u)}\int_s^u e^{K(x)}\alpha\,dx\,du + \int_s^t \sigma(x)B(x;t)d\hat{W}(x)$$

Define

$$\begin{aligned} m(r,s;t) &= B(s;t)r + \int_s^t e^{-K(u)} \int_s^u e^{K(x)} \alpha \, dx \, du \\ v(s;t) &= \int_s^t \sigma(x)^2 B(x;t)^2 \, dx \end{aligned}$$

so that, as in the Vasicek Model, $I(s;t)$, conditional on $\mathcal{F}_s$, is Normal with mean $m(r(s), s;t)$ and variance $v(s;t)$.

Proposition 39.3 *The value at time s of a zero-coupon bond maturing at time t is*

$$\begin{aligned} P(s;t) &= E_Q\left(e^{-I(s;t)} \middle| \mathcal{F}_s\right) \\ &= e^{-m(r(s),s;t)+v(s;t)/2} \\ &= P(r(s), s;t) \end{aligned}$$

where

$$\begin{aligned} P(r,s;t) &= e^{-m(r,s;t)+v(s;t)/2} \\ &= e^{-A(s;t)-B(s;t)r} \end{aligned}$$

Proof: Exercise.■

Notice that $P(r,s;t)$ is not a function of r and $t-s$ alone; it is a function of r, s, and $t-s$ (equivalently, r, s and t) because the parameters are time-varying, and hence $P(r,s;t)$ is not stationary.

Proposition 39.4 (Proposition 7.6 in Nielsen)

$$\begin{aligned} m(r(s),s;t) &= \int_s^t E_Q\left(r(u)|\mathcal{F}_S\right) \, du \\ v(s;t) &= \int_s^t \int_s^t \mathrm{Cov}_Q\left(r(u), r(y)|\mathcal{F}_s\right) \, du \, dy \end{aligned}$$

Consequently,

$$\begin{aligned} v(s;t) &\to 0 \text{ as } s \to t \\ A(s;t) &\to 0 \text{ as } s \to t \\ A(s;t) &\to 0 \text{ as } s \to t \\ P(r,s;t) &\to 1 \text{ as } s \to t \ \ (r \text{ fixed}) \\ P(r(s),s;t) &\to 1 \text{ as } s \to t \quad \text{with probability } 1 \end{aligned}$$

Proposition 39.5

$$\frac{dP(s;t)}{P(s;t)} = r(s)\,ds - B(s;t)\sigma(s)\,d\hat{W}(s)$$

In particular, the stochastic differential with respect to Q is not affected by α

Proof:

$$\begin{aligned} d\ln P(s;t) &= d\left(-B(s;t)r(s) - A(s;t)\right) \\ &= -B(s;t)\,dr(s) - r(s)\,dB(s;t) - dA(s;t) \end{aligned}$$

Therefore, the relative dispersion with respect to $\hat{W}$ is $-B(s;t)\sigma(s)$. The relative drift coefficient of the bond price with respect to the risk-adjusted probability Q equals the interest rate $r(s)$. Therefore,

$$\frac{dP(s;t)}{P(s;t)} = r(s)\,ds - B(s;t)\sigma(s)\,d\hat{W}(s)$$

■

Remark 39.6 In the Vasicek Model, the stochastic differential of $P(s;t)$ with respect to Q doesn't depend on $\bar{r}$; in the Extended Vasicek Model, the stochastic differential doesn't depend on α. α affects $P(r(s),s;t)$, but conditional on this value, α doesn't affect the relative drift or relative dispersion of the bond price. α affects the distribution of interest rates after t, but these rates are not taken into account in $P(s;t)$ as s increases.

Proposition 39.7 *v, B and A are continuously differentiable with respect to t, in particular*

$$\begin{aligned} v_t(s;t) &= 2\int_s^t \sigma(x)^2 B(x;t)\frac{\partial B}{\partial t}(x;t)\,dx \\ B_t(s;t) &= e^{-K(t)+K(s)} \\ A_t(s;t) &= \int_s^t B_t(x;t)e^{K(x)}\alpha(x)\,dx - \frac{v_t(s;t)}{2} \end{aligned}$$

Proof: See Nielsen. ■

We now turn to yields in the Extended Vasicek Model.

Proposition 39.8 *The continuously compounded yield on the zero-coupon bond is*

$$\begin{aligned} R(s;t) &= \frac{1}{t-s}\left(m(r(s),s;t) - \frac{v(s;t)}{2}\right) \\ &= R(r(s),s;t) \end{aligned}$$

where

$$R(r,s;t) = \frac{1}{t-s}\left(A(s;t) + B(s;t)r\right)$$

In particular, $R(s;t)$ is an affine function of $r(s)$, so it is normally distributed; it is not stationary.

$$E_Q\left(r_a|\mathcal{F}_s\right) = R(r,s;t) + \frac{v(s;t)}{2(t-s)}$$

Proof: See Nielsen. ■

Remark 39.9 As in the Vasicek Model, the expectation (with respect to Q) of the interest rate over the period $[s,t]$ exceeds the rate of return on the bond by the yield risk premium $\frac{v(s;t)}{2(t-s)} > 0$. Measured with respect to the true probabilities, the yield risk premium will be even greater. Empirically, the yield risk premium for nominal bonds is negative: the return from long bonds exceeds the expected return from the money-market. In our discussion of the Vasicek Model, we suggested the reason for the divergence between theory and empirics was that the Vasicek Model makes more sense as a model for real interest rates than as a model for nominal interest rates. With the Extended Vasicek Model, it is harder to make that argument convincingly. As we shall see, with the Extended Vasicek Model, we can calibrate the parameters to reproduce essentially any yield curve, and in particular the empirical yield curve for nominal bonds. The calibrated values of the parameters then lead to a prediction of the future interest rate process, but this prediction is systematically wrong; empirically, the realized rates will be systematically lower than those predicted from the calibrated parameter values.

40 Calibrating the Extended Vasicek Model

In order to calibrate the Extended Vasicek Model to fit the observed yield curve at time s, we need to know

- $r(s)$, the current Money-Market interest rate
- $\sigma(t)$, the standard deviation of interest rates at all future times $t \geq s$
- $R(r(s), s; t)$, the continuously compounded yield of the zero-coupon bond maturing at all future times $t \geq s$
- $\frac{B(s;t)\sigma(s)}{t-s}$, the instantaneous standard deviation of yields of zero-coupon bonds maturing at time t, for all future times $t \geq s$.

Except for $\sigma(t)$, all of these are observable at time s for a finite (but relatively closely spaced) number of future values of t up to approximately thirty years from s. Since the data are in fact finite, it is necessary to interpolate C^2 functions that pass through the observed data points. The observability of $\sigma(t)$ *is* problematic, and Nielsen doesn't address this issue. The best justification I can give is that, of the parameters in the model, $\sigma(t)$ is the one which is most likely to remain approximately constant over time, and whatever variations there might be in $\sigma(t)$ in the future seem virtually impossible to predict, so it may make sense to do the calibration as if $\sigma(t)$ were constant and equal to $\sigma(s)$, which can be measured.

The calibration proceeds as follows:

1. Compute $B(s;t)$ from $\frac{B(s;t)\sigma(s)}{t-s}$ and $\sigma(s)$.

2. Calculate the derivatives $B_t(s;t)$ and $B_{tt}(s;t)$ and compute
$$a(t) = -\frac{B_{tt}(s;t)}{B_t(s;t)}$$
(Proposition 7.8 in Nielsen). Thus, a is determined by the curvature of $B(s;t)$.

3. Calculate $A(s;t)$ from the formula
$$A(s;t) = -\ln P(s;t) - B(s;t)r(s)$$
where $P(s;t)$ can be inferred from $R(r(s), s; t)$.

4. Calculate the derivatives $A_t(s;t)$ and $A_{tt}(s;t)$ and compute
$$\alpha(t) = a(t)A_t(s;t) + A_{tt}(s;t) + e^{-2K(t)} \int_s^t \sigma^2(x)e^{2K(x)}\, dx$$
(Proposition 7.8 in Nielsen). a is used to match the standard deviations of the yields, while α is used to match the yields themselves.

41 Introduction to Lévy Processes

Actual stock market prices show fatter tails than the log normal distribution. The normal distribution is completed specified by its mean μ and standard deviation σ; σ can be determined from the probability that the variable lies in any *one* interval of the form $[\mu - \alpha, \mu + \alpha]$; actual log prices put less weight near μ and more far from μ than is consistent with normality. Indeed, since the price of a stock can move by a very large percentage in a matter of minutes following an important announcement, it seems at least as reasonable to model stock prices as having discontinuities as it is to model them as continuous.

Whether or not stock prices exhibit actual discontinuities, the fat tails of stock prices manifest themselves in the price of options. At any given time t, the stock price $S(t)$ is known, as are the prices of options trading at various exercise prices $X_j : j = 1, \ldots J$. If we know the price of the option with a given exercise price X_j, and we assume that the Black-Scholes Model is correct, we can infer σ_j, the volatility the stock must have to justify the observed price of the option. If the Black-Scholes Model were correct, all of the σ_j would be equal; if we plotted implied volatility against the exercise price, we would see a horizontal line. In fact, the implied volatility is not a horizontal line; it has the shape of a "smile," in which the implied volatility is lowest when X is close to $S(t)$, and rises as X moves away from $S(t)$ in either direction. This says that the probability that an option that is far out of the money ($S(t)$ is much lower than X) at time t will end up in the money at the exercise date T is higher than predicted by Black-Scholes, and hence the right to purchase the stock at the exercise price at the exercise date T is more valuable than predicted by Black-Scholes. Similarly, the probability than an option that is far in the money ($S(t)$ is much higher than X) at time t will end up out of the money at the exercise date T is higher than predicted by Black-Scholes, and hence the fact that the option does not *obligate* the option-holder to buy the stock is more valuable than predicted by Black-Scholes.

What properties should our stock price processes have? For mathematical tractability, we presumably want a one-parameter family of distributions such that the distribution of $\frac{S(t)}{S(s)}$ lies in this family for all s, t. Logged prices exhibit an additive structure: if $r < s < t$,

$$\begin{aligned}\ln S(t) - \ln S(r) &= \ln\left(\frac{S(t)}{S(r)}\right) \\ &= \ln\left(\frac{S(t)}{S(s)} \times \frac{S(s)}{S(r)}\right) \\ &= (\ln S(t) - \ln S(s)) + (\ln S(s) - \ln S(r))\end{aligned}$$

Therefore, our family of distributions (or rather their logs) must have the property than the sum of two independent random variables with distributions in the family must also have distribution in the family.

Definition 41.1 Suppose we are given a filtration $\{\mathcal{F}_t\}$. A Lévy Process is an adapted process X with $X_0 = 0$ which satisfies the following properties:

1. X has independent increments, i.e. $X_t - X_s$ is independent of $\mathcal{F}_s$ whenever $0 \leq s < t < \infty$;

2. X has stationary increments, i.e. the distribution of $X_t - X_s$ equals the distribution of X_{t-s} whenever $0 \leq s < t < \infty$; and

3. X_t is continuous in probability, i.e. for all t, $\lim_{s \to t} X(s) = X(t)$ almost surely.

Remark 41.2 Note that it follows immediately from the definition that a Wiener Process is a Lévy Process. But we shall see in a moment that the Poisson Process, which has discontinuous paths, is also a Lévy Process. The third condition in the definition requires that the process be continuous at time t, except for a null set *which may depend on* t; by moving the null set around, we can construct a process whose paths are discontinuous with probability one, but the set of times at which any given path is discontinuous will be countable. The third condition allows for sudden unanticipated discontinuities, but it does not allow for anticipated discontinuities; if it is known in advance that an announcement will be made at 2:00pm, and it will have a big effect on the stock price, the stock price is not continuous in probability, even if the direction of the effect remains unknown until 2:00pm. Hence, Lévy Processes are suitable for modelling unanticipated discontinuities in stock prices, but not anticipated discontinuities.

Remark 41.3 If a Lévy Process X has continuous paths and finite variances, then there exists σ sutch that σX is a (standard) Wiener Process. Thus, continuous Levy Processes other than Wiener Processes have infinite variances, which creates a variety of mathematical and modelling problems in using them to describe stock prices. We will not pursue these issues in these notes, and will focus on Poisson Processes. Theorem 41 on page 31 of Protter [8] says, essentially, the following: If X is a Lévy Process with bounded jumps, then $X(t) - E(X(t))$ is the sum of a continuous martingale Lévy Process and a martingale which is a mixture of compensated Poisson Processes.[1] Theorem 42 on page 32 provides another representation theorem.

42 The Poisson Process

Fix a parameter $\rho > 0$. We give two formulations of the Poisson Process with parameter ρ:

- **Discrete Formulation:** Given $n \in \mathbf{N}$, let the time axis be $\mathcal{T}_n = \left\{0, \frac{1}{n}, \frac{2}{n}, \ldots\right\}$, $\Omega = \{0,1\}^{\mathbf{N}}$. We put a measure P_n on Ω by requiring

$$P_n(\omega_j = 1) = \frac{\rho}{n} = \rho dt, \;\; P_n(\omega_j = 0) = 1 - \frac{\rho}{n}$$

and we extend to a measure $P - n$ on Ω by requiring that ω_j and ω_k be independent if $j \neq k$; equivalently, we take P_n to be the product measure induced on Ω by the above distributions on the individual ω_j. We define

$$X_n(\omega, t) = \sum_{j \leq nt} \omega_j$$

Notice that X_n has the following properties:

1. $X_n : \Omega \times [0, \infty) \to \mathbf{N} \cup \{0\}$;
2. X_n is nondecreasing;
3. X_n exhibits stationary, independent increments; the conditional distribution $(X_n(t) - X_n(s) | X_n(s))$ is the same as the distribution of $X_n(t-s)$;

[1] If X is a Poisson Process, then $X(t) - E(X(t))$ is called a compensated Poisson Process; as we shall see, for a Poisson Process with parameter ρ, $X(t) - E(X(t)) = X(t) - \rho t$.

4. $P_n(X_n(t) = 0) = \left(1 - \frac{\rho}{n}\right)^{nt}$
5. X_n is CADLAG[2], i.e. it is right-continuous, and has limits from the left everywhere; moreover,

$$P_n\left(\lim_{s \nearrow t} X_n(s) = X_n(t)\right) = 1 - \frac{\rho}{n} \text{ if } t \in \mathcal{T}_n$$

Given a CADLAG process X, define

$$X(\omega, t-) = \lim_{s \nearrow t} X(\omega, s) \text{ and } X_-(\omega, t) = X(\omega, t-)$$

- **Continuous Formulation:** There are two continuous formulations of the Poisson Process. First, in an analogous manner to the construction of Wiener measure as the limit of measures on the set of continuous functions induced by random walks, we can view P_n and X_n as generating a sequence of measures μ_n on the space of CADLAG integer-valued processes, then define the Poisson measure as the limit of these measures as $n \to \infty$, then use this to recover a probability space (Ω, P) and a Poisson process X.[3] The resulting process X has the following properties:

1. $X : \Omega \times [0, \infty) \to \mathbf{N} \cup \{0\}$;
2. X is nondecreasing;
3. X exhibits stationary, independent increments; the conditional distribution $(X(t) - X(s)|X(s))$ equals the distribution of $X(t-s)$ if $s < t$.
4.

$$\begin{aligned} P(X(t) = 0) &= \lim_{n\to\infty} \left(1 - \frac{\rho}{n}\right)^{nt} \\ \ln\left(P(x(t) = 0)\right) &= \lim_{n\to\infty} \ln\left(\left(1 - \frac{\rho}{n}\right)^{nt}\right) \end{aligned}$$

[2]CADLAG, sometimes written càdlàg, is an acronym for the French "continu à droite avec limites à gauche."

[3]Indeed, we could take Ω to be the set of CADLAG functions, P to be the limit of the μ_n, and X to be identity map from Ω to itself.

$$
\begin{aligned}
&= \lim_{n\to\infty} nt \ln\left(1-\frac{\rho}{n}\right) \\
&= \lim_{n\to\infty} nt \left(-\frac{\rho}{n} + \left(\frac{\rho}{n}\right)^2 + \cdots\right) \\
&= -\rho t \\
P(X(t)=0) &= e^{-\rho t} \to 0 \text{ as } t \to \infty
\end{aligned}
$$

5. X is CADLAG; moreover, for every t,

$$P\left(\lim_{s\nearrow t} X(s) = X(t)\right) = 1$$

Therefore, X is a Lévy Process.

Let

$$
\begin{aligned}
Y_1(\omega) &= \min\{t : X(\omega,t) = 1\} \\
Y_2(\omega) &= \min\{t : X(\omega,t) = 2\} - Y_1(\omega) \\
&\vdots \\
Y_{k+1}(\omega) &= \min\{t : X(\omega,t) = k+1\} - Y_k(\omega) \\
&\vdots
\end{aligned}
$$

so that Y_{k+1} is the time between the k^{th} and the $k+1^{st}$ jumps. Let $f(s)$ be the probability density of the common distribution of the Y_k.

$$
\begin{aligned}
P(X(t)=0) &= e^{-\rho t} \\
&= \int_t^\infty f(s)\,ds \\
-\rho e^{-\rho t} &= -f(t) \qquad \text{(Fundamental Theorem of Calculus)} \\
f(t) &= \rho e^{-\rho t} \text{ for } t \geq 0 \\
X(\omega,t) &= \max\{k : Y_1(\omega) + \cdots + Y_k(\omega) \leq t\}
\end{aligned}
$$

The alternative formulation of the Poisson Process is to start with a family $\{Y_k : k \in \mathbf{N}\}$ of independent random variables with common density $f(t) = \rho e^{-\rho t}$ for $t \geq 0$, and *define*

$$X(\omega,t) = \max\{k : Y_1(\omega) + \cdots + Y_k(\omega) \leq t\}$$

Proposition 42.1 *The augmented filtration $\mathcal{F}^X$ generated by a Poisson Process X is right-continuous, i.e.*

$$\mathcal{F}_s^X = \bigcap_{t>s} \mathcal{F}_t^X$$

The idea of the proof is that as $t \searrow s$, the probability of a jump in the interval $(s,t]$ goes to zero; if $A \in \mathcal{F}_t$ for every $t > s$, there must be a set $\hat{A} \in \mathcal{F}_s$ such that $P((A \setminus \hat{A}) \cup (\hat{A} \setminus A)) = 0$, but since the filtration is augmented, $\hat{A} \in \mathcal{F}_s$.

43 Stochastic Integration with Respect to Poisson Processes

Because the Poisson Process X is almost surely of bounded variation on compact time intervals, the stochastic integral with respect to X is simply a Stieltjes Integral. Let $t_k(\omega)$ be the time of the k^{th} jump, so

$$t_k(\omega) = Y_1(\omega) + \cdots + Y_k(\omega)$$

Then

$$\begin{aligned}
\int_0^t \Delta\, dX &= \sum_{k=1}^{X(\omega,t)} \Delta(\omega, t_k(\omega))(X(\omega, t_k(\omega) - X(\omega, t_k(\omega)-))) \\
&= \sum_{k=1}^{X(\omega,t)} \Delta(\omega, t_k(\omega))
\end{aligned}$$

While the definition makes sense for all Δ, we will generally want to assume that Δ is predictable, a slight strengthening of adapted. If we only require that Δ be adapted, then $\Delta(\omega,t)$ can be conditioned on whether X has a jump at (ω,t); if we think of Δ as a trading strategy, then it as if a trader could observe the price of a stock has doubled, but still buy it at the old, pre-doubled price; small changes in Δ result in big changes in $\int \Delta\, dX$. In the discrete framework, we say that Δ is *predictable* if $\Delta\left(\omega, \frac{k}{n}\right)$ is $\mathcal{F}_{\frac{k-1}{n}}$-measurable.

Definition 43.1 The *predictable σ-algebra* is the σ-algebra on $\Omega \times [0,\infty)$ generated by the adapted left-continuous processes. A stochastic process Δ is *predictable* if it is measurable with respect to the predictable σ-algebra.

44 Silly First Model

In this section, we explore a Poisson-based stock price model. To help us get going, it is simple. It is silly because it only has upward jumps. Clearly, stock prices can exhibit jumps in either direction; indeed, casual empiricism suggests that downward jumps are more frequent and larger than upward jumps.

Let X be a Poisson Process with parameter ρ. Our securities are a stock $S = e^X$ and a Money-Market Account $M(t) = e^{rt}$, with $r > 0$ a constant.

We address a series of questions:

1. **Is S a stochastic integral with respect to X?** The answer is yes. S jumps exactly when X jumps. If X has a jump at t,

$$\begin{aligned} X(t) - X(t-) &= 1 \\ e^{X(t)} - e^{X(t-)} &= e^{X(t-)+1} - e^{X(t-)} \\ &= (e-1)e^{X(t-)} \end{aligned}$$

so we set

$$\Delta(t) = (e-1)e^{X(t-)} = (e-1)S(t-)$$

Notice that Δ is predictable, and

$$S(t) = 1 + \int_0^t \Delta \, dX(s)$$

2. **Is there arbitrage?** The answer is no, even if one allows doubling strategies which generate arbitrage with respect to geometric Wiener processes. Since $r > 0$, and for every t, $P(X(t) = 0) > 0$, there is a positive probability that borrowing to invest in the stock will end up losing money over the time interval $[0, t]$.

3. **Is there a state price process?** We first have to decide what we mean by a state price process. Since X is not a martingale, stochastic integrals $\int \Delta \, dX$ have drift under any reasonable definition of drift, so we can't define zero drift to mean the coefficient of dt in the stochastic differential is zero.[4] So we will say that Π is a state price process if ΠS

[4] We could probably replace the Poisson Process X with the compensated Poisson Process $\hat{X}(t) = X(t) - \rho t$, which *is* a martingale, and define zero drift as meaning the coefficient of dt is zero in the stochastic differential with respect to dt and $d\hat{X}$. Since this is a Silly Model, we will not pursue that.

and ΠM are martingales. Put aside M for the moment, and focus on finding Π such that ΠS is a martingale. By analogy with the Itô case, we expect that after we normalize $\Pi(0) = 1$, we will have a 1-parameter indeterminacy in Π. Here are two state price processes:

(a)

$$\Pi_j(t) = \frac{1}{S(t)} = e^{-X(t)}$$

Then $\Pi_j S$ is identically one, so it must be a martingale. Π_j is adapted (but not predictable). Π_j is a stochastic integral: if X jumps at t, then

$$\begin{aligned}
\Pi_j(t) - \Pi_j(t-) \quad &- \quad e^{-X(t)} - e^{-X(t-)} \\
&= \quad e^{-(X(t-)+1)} - e^{-X(t-)} \\
&= \quad e^{-X(t-)} \left(\frac{1}{e} - 1\right) \\
&= \quad \frac{1}{S(t-)} \left(\frac{1}{e} - 1\right) \\
&= \quad \Pi_j(t-) \left(\frac{1}{e} - 1\right)
\end{aligned}$$

so if we set

$$\Delta(t) = \Pi_j(t-) \left(\frac{1}{e} - 1\right)$$

then Δ is predictable and

$$d\Pi_j = \Delta\, dX, \quad \frac{d\Pi_j(t)}{\Pi_j(t-)} = \left(\frac{1}{e} - 1\right) dX$$

(b) There is also a continuous state price process Π_c. In the time interval $(t, t + \delta t]$, the probability of a jump is $\rho\, dt$.[5] If a jump occurs at time $s \in (t, t + \delta t]$,

$$\begin{aligned}
S(t + \delta t) - S(t) \quad &= \quad e^{X(t)+1} - e^{X(t)} \\
&= \quad e^{X(t)}(e - 1) \\
&= \quad S(t)(e - 1)
\end{aligned}$$

[5] The probability of two or more jumps, $(\rho\, dt)^2 + (\rho\, dt)^3 + \ldots$, is of order $O(dt^2)$. Exercise for the reader: fix the following calculation, which pretends that the probability of two or more jumps is zero.

Therefore, if $\Pi_c S$ is a martingale,

$$\begin{aligned}
\Pi_c(t)e^{X(t)} &= E\left(\Pi_c(t+dt)e^{X(t+dt)}\middle|\mathcal{F}_t\right) \\
&= \Pi_c(t+dt)\left(\rho\,dt\,e^{X(t)+1} + (1-\rho\,dt)\,e^{X(t)}\right) \\
&= \Pi_c(t+dt)e^{X(t)}\,(1+\rho(e-1)\,dt) \\
\Pi_c(t) &= \Pi_c(t+dt)\,(1+\rho(e-1)\,dt) \\
\Pi_c(t+dt)-\Pi_c(t) &= \frac{\Pi_c(t)}{1+\rho(e-1)\,dt} - \Pi_c(t) \\
&= \Pi_c(t)\left(\frac{1-(1+\rho(e-1)\,dt)}{1+\rho(e-1)\,dt}\right) \\
&= \Pi_c(t)\left(\frac{-\rho(e-1)\,dt)}{1+\rho(e-1)\,dt}\right) \\
&= -\Pi_c(t)\rho(e-1)\,dt\left(1-\rho(e-1)\,dt + (\rho(e-1)\,dt)^2 + \ldots\right) \\
&= -\Pi_c(t)\rho(e-1)\,dt + O\left(dt^2\right) \\
\frac{d\Pi_c}{\Pi_c} &= -\rho(e-1)\,dt \\
\Pi_c(t) &= e^{-\rho(e-1)t} \\
&= \eta[-\rho(e-1), 0]
\end{aligned}$$

Now, let's put the Money-Market Account M back in and see if we can find a state price process. Before doing this, we take a detour to consider Equivalent Martingale Measures.

(c) **Is there an Equivalent Martingale Measure?** There's no equivalent measure Q which makes S into a martingale. Since $S = e^X$ is increasing, for any equivalent measure Q and any $t > s$,

$$E_Q\left(e^{X(t)}\middle|\mathcal{F}_s\right) > e^{X(s)}$$

But that's not the point. We want to find an equivalent measure Q such that $e^{-rt}S$ is a martingale, and this we *can* do. Let's return to the discrete model.

$$E\left(e^{-r(k+1)/n}S\left(\frac{k+1}{n}\right) - e^{-rk/n}S\left(\frac{k}{n}\right)\middle|\,S\left(\frac{k}{n}\right)\right)$$

$$
\begin{aligned}
&= \left(e^{-r/n}\left(\left(\frac{\rho e}{n}\right) + \left(1 - \frac{\rho}{n}\right)\right) - 1\right) e^{-rk/n} S\left(\frac{k}{n}\right) \\
&= \left(e^{-r/n}\left(1 + \frac{\rho(e-1)}{n}\right) - 1\right) e^{-rk/n} S\left(\frac{k}{n}\right) \\
&= \left(\left(1 - \frac{r}{n} + O\left(n^{-2}\right)\right)\left(1 + \frac{\rho(e-1)}{n}\right) - 1\right) e^{-rk/n} S\left(\frac{k}{n}\right) \\
&= \left(-\frac{r}{n} + \frac{\rho(e-1)}{n} + O\left(n^{-2}\right)\right) e^{-rk/n} S\left(\frac{k}{n}\right)
\end{aligned}
$$

Thus, we could make $e^{-rt}S(t)$ into a martingale if we had

$$\rho = \frac{r}{e-1} + O\left(\frac{1}{n}\right)$$

This gives us the recipe for Q. In the continuous case, if we fix a time interval $[0, T]$, Q is the measure on the space of nondecreasing, integer-valued paths that generates a Poisson Process with parameter

$$\rho_Q = \frac{r}{e-1}$$

Q is equivalent to P on finite horizons.[6]

(d) **Is There a State Price Process (Revisited)** From the discussion of the equivalent martingale measure, we see that there is a state price process for S and M. Let $\frac{dQ}{dP}$ be the Radon-Nikodym derivative of the measure Q which makes X into a Poisson Process with parameter $\rho_Q = \frac{r}{e-1}$. Then

$$\Pi = e^{-rt}\frac{dQ}{dP}$$

is a state price process for S and M. I believe it is unique (up to the constant $\Pi(0)$), but I can't swear to this at this point. The advantage of this representation is tractability. One can readily compute expectations with respect to Q because these will be expectations with respect to a Poisson Process.

[6] This is exactly what is true in the Itô case. In Girsanov's Theorem, the measure Q is equivalent to P for each finite time horizon T, but not with an infinite time horizon.

45 Better Second Model

In our second model, we suppose that

$$X = X_1 - X_2, \;\; S = e^X$$

where X_1 and X_2 are independent Poisson Processes with parameter $\frac{\rho}{2}$. In the discrete model, we could take $\Omega = \{-1, 0, 1\}^{\mathbf{N}}$,

$$P_n(\omega_j = -1) = P_n(\omega_j = 1) = \frac{\rho}{2n}, \;\; P_n(w_j = 0) = 1 - \frac{\rho}{n}$$

Note that in the discrete model, the positive and negative Poisson processes are not quite independent, but they become independent as $n \to \infty$.

Notice that X is a martingale, so it is reasonable to say that a process with differential

$$0\,dt + a\,dX$$

has zero drift.

However, we immediately encounter a problem:

Proposition 45.1 *S is not a stochastic integral with respect to X.*

Proof: S jumps exactly when X jumps. However, whereas the upward and downward jumps of X are of the same magnitude (indeed, both jumps are of size 1), the upward and downward jumps of S are of unequal magnitude. If X jumps up at t, we have

$$\frac{e^{X(t)} - e^{X(t-)}}{X(t) - X(t-)} = (e-1)e^{X(t-)}$$

while if X jumps down at t, we have

$$\frac{e^{X(t)} - e^{X(t-)}}{X(t) - X(t-)} = \left(1 - \frac{1}{e}\right) e^{X(t-)} = \left(\frac{e-1}{e}\right) e^{X(t-)}$$

and these are not equal. Since it is not known at times $s < t$ the direction in which X will jump (or, for that matter, whether it will jump at all at time t), there is no predictable process Δ such that $dS = \Delta dX$.■

In particular, this tells us the statement of Itó's Lemma has to be modified to extend to this case.

46 Itô's Lemma for Discontinuous Processes

The natural class of stochastic processes that can be used as stochastic integrators is the class of semimartingales. The definition of a semimartingale is abstract and somewhat cumbersome. Roughly speaking, a semimartingale is the sum of a local martingale (which we don't define) and a process of bounded variation. In practice, every stochastic process that you would want to use as a stochastic integrator is a semimartingale: in particular, all Lévy Processes (including Wiener Processes and Poisson Processes) are semimartingales; the sum of two semimartingales is a semimartingale; every stochastic integral with respect to a semimartingale is a semimartingale; and every C^2 function of a semimartingale is a semimartingale. Semimartingales can be assumed to be CADLAG. See Protter [8] for details.

Definition 46.1 If Y is a semimartingale, the *quadratic variation* of Y is

$$[Y,Y] = Y^2 - 2\int Y_-\,dY$$

where we recall that $Y_-(\omega,t) = \lim_{s\nearrow t} Y(\omega,s)$.

Remark 46.2 If W is a Wiener Process, recall from Problem Set 2 that

$$[W,W](t) = W^2(t) - 2\int_0^t W_-\,dW = W^2 - 2\frac{W^2-t}{2} = t$$

which is the quadratic variation of the Wiener Process.

For X a Poisson process, or for $X = X_1 - X_2$, where X_1 and X_2 are Poisson Processes, the quadratic variation clearly ought to be the number of jumps. The process is constant except at the jump points, and the square of the jump is always equal to one. We'll see that the quadratic variation, as just defined, does indeed equal the number of jumps. If X is the standard Poisson Process (i.e., the one with only positive jumps), let the jumps in $[0,t]$ be at $t_1, t_2, \ldots, t_{X(t)}$

$$\begin{aligned}
[X,X](t) &= X(t)^2 - 2\int_0^t X_- dX \\
&= X(t)^2 - 2\sum_{j=1}^{X(t)} X(t_j-)(X(t_j) - X(t_j-))
\end{aligned}$$

$$
\begin{aligned}
&= (t)X^2 - 2\left(0\cdot 1 + 1\cdot 1 + 2\cdot 1 + \cdots + (X(t)-1)\cdot 1\right) \\
&= X(t)^2 - 2\left[\frac{(X(t)-1)X(t)}{2}\right] \\
&= X(t)^2 - X(t)^2 + X(t) \\
&= X(t)
\end{aligned}
$$

which is exactly the number of jumps. The calculation for the difference of Poisson Processes $X = X_1 - X_2$ is very hard to write down on paper, and much easier to do on the board, so I hope you got the idea from the lecture.

Definition 46.3 If Y is a semimartingale, then $[Y,Y]$ is a process of finite variation on compact intervals, hence each path has at most countably many jumps. Thus, we can split off all the jumps of $[Y,Y]$, leaving a continuous process which is denoted $[Y,Y]^c$.

Theorem 46.4 (Itô's Lemma (Theorem 32, page 71 in Protter)) *If Y is a semimartingale and $f : \mathbf{R} \to \mathbf{R}$ is C^2, then $f(Y)$ is a semimartingale and*

$$
\begin{aligned}
f(Y(t)) &= f(Y(0)) + \int_{0+}^{t} f'(Y(s-))\,dY(s) + \frac{1}{2}\int_{0+}^{t} f''(Y(s-))\,d[Y,Y]^c(s) \\
&\quad + \sum_{0<s\le t} \left(f(Y(s)) - f(Y(s-)) - f'(Y(s-))(Y(s)-Y(s-))\right)
\end{aligned}
$$

Remark 46.5 The last sum is over an uncountable index set. However, a semimartingale can have only countably many jumps, and the summand is zero except at points where Y jumps. We get a correction term at each jump because we need to add in the change in f over the jump $(f(Y(s)) - f(Y(s-)))$ and remove the term $f'(Y(s-))(Y(s)-Y(s-))$ that appears in the first integral. If $X = X_1 - X_2$ is the difference of Poisson Processes, Itô's Lemma doesn't tell us anything we didn't know: It says

$$
\begin{aligned}
f(X(t)) &= 1 + \sum_{0<s\le t} f'(X(s-))\,(X(s)-X(s-)) + 0 \\
&\quad + \sum_{0<s\le t} \left((f(X(s)) - f(X(s-))) - f'(X(s-))\,(X(s)-X(s-))\right) \\
&= \sum_{0<s\le t} \left(f(X(s)) - f(X(s-))\right)
\end{aligned}
$$

The good news is, that if Y is a continuous semimartingale, Itô's Lemma goes through provided we substitute $[Y, Y]$ for the term $(dW)^2 = dt$ in the version for Wiener Processes. The second bit of good news is that for CADLAG semimartingales, there is an Itô Calculus, provided one is a little careful about the jumps. The bad news is that the jump terms that appear in the statement of Itô's Lemma do not appear in the capital gains process we need for Finance, so while we may be able to price options on discontinuous securities, it appears we are going to have serious trouble replicating them.

47 More on the Second Model

Now that we have finished our detour through Itô's Lemma, we resume our discussion of the second model, in which $S = e^X$, where $X = X_1 - X_2$ is a difference of Poisson Processes. In addition to our stock X, we have a Money Market account $M(t) = e^{rt}$.

1. **Is there an Equivalent Martingale Measure?** The answer is yes. We work in the discrete model.

$$
\begin{aligned}
&E\left(e^{-r(k+1)n} S\left(\frac{k+1}{n}\right) - S\left(\frac{k}{n}\right) \middle| S\left(\frac{k}{n}\right)\right) \\
&= \left(e^{-r/n}\left(\frac{\rho}{2n}e + \frac{\rho}{2n}\frac{1}{e} + \left(1 - \frac{\rho}{n}\right)\right) - 1\right) e^{-rk/n} S\left(\frac{k}{n}\right) \\
&= \left(\left(1 - \frac{r}{n} + O\left(\frac{1}{n^2}\right)\right)\left(1 + \frac{\rho}{2n}\left(e + \frac{1}{e}\right) - \frac{\rho}{n}\right) - 1\right) e^{-rk/n} S\left(\frac{k}{n}\right) \\
&= \left(-\frac{r}{n} + \frac{\rho}{2n}\left(e + \frac{1}{e} - 2\right) + O\left(\frac{1}{n^2}\right)\right) e^{-rk/n} S\left(\frac{k}{n}\right)
\end{aligned}
$$

We can make the discrete model a martingale (in the limit as $n \to \infty$) provided

$$
\begin{aligned}
\frac{\rho}{2n}\left(e + \frac{1}{e} - 2\right) &= \frac{r}{n} \\
\rho &= \frac{2r}{e + \frac{1}{e} - 2}
\end{aligned}
$$

so if we let Q be the measure which turns X_1 and X_2 into Poisson Processes with parameter

$$\frac{\rho_Q}{2} = \frac{r}{e + \frac{1}{e} - 2}$$

on $[0, T]$, then Q is an Equivalent Martingale Measure and

$$\Pi(t) = e^{-rt}\frac{dQ}{dP}$$

is a state price process.

2. **Is the Equivalent Martingale Measure Unique?** The answer is no. The problem is that the set of probability measures on the three-point space $\{-1, 0, 1\}$ is two-dimensional, and we should not expect to get a single equation (the martingale condition for S) to produce a unique solution. Instead of adjusting the parameter ρ (which determines the rate at which jumps occur), we can keep ρ constant and adjust the relative probability of up and down jumps. In the discrete model, if we set the probabilty of an up jump to $(1-\alpha)\frac{\rho}{n}$, the probability of a down jump to $\alpha\frac{\rho}{n}$, and leave the probability of no jump at $1-\frac{\rho}{n}$, we will find that in the limit as $n \to \infty$, e^X satisfies the martingale equation provided that
$$\alpha = \frac{e - 1 - \frac{r}{\rho}}{e - \frac{1}{e}}$$
In order for this to generate a probability measure, we need to know that $\alpha \in [0, 1]$. Since $1 + \frac{r}{\rho} > 1 > \frac{1}{e}$, we automatically have $\alpha < 1$. We will have $\alpha \geq 0$ provided that $\frac{r}{\rho} \leq e - 1$. Thus, provided $\frac{r}{\rho} < e - 1$, we get a one-parameter family $\mathcal{P}$ of pairs (α, ρ). At each node, we can make a different choice $(\alpha(\omega, t), \rho(\omega, t))$. As long as we do this in a measurable way, the resulting measure will be an Equivalent Martingale Measure. Eberlein and Jacod [4] show that for "most" purely discontinuous processes, the range of values of the option, using all equivalent martingale measures, is the interval

$$\Big(\max\{S(t) - e^{-r(T-t)}X, 0\}, S(t)\Big)$$

which says the martingale method has essentially no predictive value.

48 Third Model

The third model has $X = X_1 - X_2$ a difference of Poisson Processes, and two stocks, so

$$\bar{S} = \begin{pmatrix} e^{rt} \\ e^{\mu_1 t + \sigma_1 X} \\ e^{\mu_2 t + \sigma_2 X} \end{pmatrix} = \begin{pmatrix} M \\ S_1 \\ S_2 \end{pmatrix}$$

Once again, we look in the discrete model for conditions to make $e^{-rt}\bar{S}$ into a martingale. The equation for S_1 is

$$\begin{aligned} 1 &= e^{(\mu_1 - r)n}\left((1-\alpha)\frac{\rho}{n}e^{\sigma_1} + \alpha\frac{\rho}{n}e^{-\sigma_1} + \left(1 - \frac{\rho}{n}\right)\right) \\ &= \left(1 + \frac{\mu_1 - r}{n} + O\left(\frac{1}{n^2}\right)\right)\left(\frac{\rho}{n}\left(\alpha e^{-\sigma_1} + (1-\alpha)e^{\sigma_1} - 1\right) + 1\right) \\ &= 1 + \frac{\mu_1 - r}{n} + \frac{\rho}{n}\left(\alpha e^{-\sigma_1} + (1-\alpha)e^{\sigma_1} - 1\right) + O\left(\frac{1}{n^2}\right) \end{aligned}$$

so to make both S_1 and S_2 martingales, we need

$$\begin{aligned} \rho_1 &= \frac{r - \mu_1}{(1-\alpha)e^{\sigma_1} + \alpha e^{-\sigma_1} - 1} \\ \rho_2 &= \frac{r - \mu_2}{(1-\alpha)e^{\sigma_2} + \alpha e^{-\sigma_2} - 1} \\ \rho_1 &= \rho_2 \end{aligned}$$

Proposition 48.1 *Generically in σ_1, σ_2, μ_1, and μ_2, there is a unique $\alpha(\sigma_1, \sigma_2, \mu_1, \mu_2)$ such that $\rho_1 = \rho_2$.*

Proof: Let

$$f(\alpha) = \frac{(1-\alpha)e^{\sigma_1} + \alpha e^{-\sigma_1} - 1}{r - \mu_1} - \frac{(1-\alpha)e^{\sigma_2} + \alpha e^{-\sigma_2} - 1}{r - \mu_2}$$

$f(\alpha)$ is linear in α, and generically it is not constant, hence generically there exists a unique α_0 such that $f(\alpha_0) = 0$. If $\sigma_1 \neq \sigma_2$, then at least one of $(1-\alpha)e^{\sigma_1} + \alpha e^{-\sigma_1} - 1$ and $(1-\alpha)e^{\sigma_1} + \alpha e^{-\sigma_1} - 1$ is nonzero for each α, in particular for α_0. Since $f(\alpha_0) = 0$, we must have

$$(1-\alpha_0)e^{\sigma_1} + \alpha_0 e^{-\sigma_1} - 1 \neq 0 \neq (1-\alpha_0)e^{\sigma_2} + \alpha_0 e^{-\sigma_2} - 1$$

so

$$\frac{r-\mu_1}{(1-\alpha)e^{\sigma_1}+\alpha e^{-\sigma_1}-1}=\frac{r-\mu_2}{(1-\alpha)e^{\sigma_2}+\alpha e^{-\sigma_2}-1}$$

■

Conjecture 48.2 (Probability 0.95) For a generic set of $(\sigma_1, \sigma_2, \mu_1, \mu_2)$, if $\bar{S}$ does not admit arbitrage, there is a unique state price process.

Remark 48.3 The proof should go along the following lines. In the discrete model, generically we have a unique ρ and α that satisfy the equation that asymptotically makes $\bar{S}$ into a martingale. There is a little problem with the $O(n^{-2})$ term; putting that aside, in the absence of arbitrage, the discrete version of the Fundamental Theorem of Finance asserts that there is an equivalent martingale measure, which says $\rho > 0$ and $\alpha \in [0, 1]$. Then take limits.

Thus, it appears very likely that we get unique martingale pricing of options in a generic model with two stocks generated by X, a difference of Poisson processes.

49 More Complicated Jump Processes

Consider the following model. Let X be a standard (nondecreasing) Poisson Process with parameter ρ. Fix $Z_1, Z_2, \cdots$ independent, identically distributed random variables, $E(Z_n) = 0$, where each Z_n is independent of the filtration generated by X. Let

$$Z(\omega, t) = \sum_{j=1}^{X(\omega,t)} Z_j(\omega)$$

Note that our process $X_1 - X_2$ is a special case in which the Z_n are IID, taking the value 1 with probabilty $\frac{1}{2}$ and -1 with probability $\frac{1}{2}$. If Z_n has a continuous distribution such as the normal, we will need infinitely many stocks in order to get a unique state price process. However, we conjecture (P=0.5) that some version of the following statement is true: given any $\varepsilon > 0$, there exists a finite number N of stocks so that, for any state price process on these stocks and the money market account, the martingale value of each option must lie in an interval of length less than ε.

50 Stochastic Volatility

In Lecture 30, we briefly discusses models with stochastic volatility and presented a conjecture. There will be an extended discussion of this in Lecture 31.

References

[1] Anderson, Robert M., "A Nonstandard Representation of Brownian Motion and Itô Integration," *Israel Journal of Mathematics* **25**(1976), 15-46.

[2] Billingsley, Patrick, *Convergence of Probability Measures*, John Wiley and Sons, 1968.

[3] Chan, K. C., G. A Karolyi, F. A. Longstaff, and A. Sanders, "An Empirical Comparison of Alternative Models of the Short-Term Interest Rate," *Journal of Finance* 47(1992), 1209-1227.

[4] Eberlein, Ernst and Jean Jacod, "On the Range of Options Prices," *Finance and Stochastics* 1(1997), 131-140.

[5] Harison, J. Michael and David M. Kreps, "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory* 20(1979), 381-408.

[6] Karatzas, Ioannis and Steven E. Shreve, *Methods of Mathematical Finance*, volume 39 in the series *Applications of Mathematics: Stochastic Modeling and Applied Probability*. New York: Springer-Verlag, 1998.

[7] Nielsen, Lars Tyge, *Pricing and Hedging of Derivative Securities*, Oxford University Press, 1999.

[8] Protter, Philip, *Stochastic Integration and Differential Equations: A New Approach*. Volume 21 in *Applications of Mathematics*. Berlin: Springer-Verlag, 1990.