

Continuous Time Finance Notes, Spring 2004 – Section 8, March 24, 2004
Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use in connection with the NYU course Continuous Time Finance.

This section begins by wrapping up our discussion of HJM – discussing its strengths and weaknesses. Then we turn to the “Libor Market Model” – which though relatively new is rapidly becoming the method of choice for many purposes. Our discussion of the Libor Market Model follows Hull’s treatment (Section 24.3 of the 5th edition).

Some comments on the HJM approach to interest rate modelling:

HJM is a “framework” not a “model”. When modeling equity-based derivatives, our usual framework is to assume the underlying solves a diffusion process $ds = \mu(s, t) dt + \sigma(s, t) dw$. The market is complete for any choice of the functions $\mu(s, t)$ and $\sigma(s, t)$ (with some minor restrictions, for example $\sigma(s, t) > 0$). But what to choose for μ and σ ? The choice of μ is of course irrelevant for option pricing, because under the risk-neutral measure the underlying satisfies $ds = r dt + \sigma(s, t) dw$. But the choice of σ is crucial. We commonly choose $\sigma(s, t) = \sigma_0 s$ with σ_0 constant, i.e. we commonly assume the underlying has *lognormal* dynamics. (This is not the only possibility; in a week or two we’ll discuss “local vol” models, i.e. the idea of using the skew/smile of implied volatilities to infer an appropriate choice of $\sigma(s, t)$.)

The situation with HJM is analogous. For one-factor HJM, the framework is $d_t f(t, T) = \alpha(t, T) dt + \sigma(t, T) dw$. Working under the risk-neutral measure (i.e. assuming w is a Brownian motion in the risk-neutral measure) the drift α is completely specified by σ . The choice of σ is again crucial. It is no longer so obvious how to choose it; we saw how to choose $\sigma(t, T)$ to get either Ho-Lee or Hull-White. But many other choices are possible. Calibration to data is not easy, because when σ is not deterministic (e.g. if it depends explicitly on $f(t, T)$) we have no exact pricing formulas.

It is natural to suppose σ is a function of f as well as t and T . Indeed, if σ is a deterministic function of t and T then $f(t, T)$ is Gaussian, so there is a positive probability that it is negative. We tolerated this in Hull-White, and one often tolerates it in HJM for the same reason – namely analytical tractability of the model. However a model that permits negative interest rates cannot be the last word. Amongst short-rate models the simplest way to keep the short rate positive is to use a state-dependent volatility: for example the Cox-Ingersoll-Ross model assumes $dr = (\theta(t) - ar) dt + \sigma_0 \sqrt{r} dw$ (compare this to Hull-White: $dr = (\theta(t) - ar) dt + \sigma_0 dw$). Our discussion of HJM (in particular: our explanation how the volatility determines the drift) did *not* assume σ was deterministic. It applies with no change if, for example, σ has the form $\sigma_0(t, T)f^\alpha$ for some α . [Note however that problems 5 and 6 on HW4 implicitly assume $\sigma(t, T)$ is deterministic.]

It is tempting to suppose f is lognormal, but this leads to problems. The standard method for pricing a cap using Black’s formula assumes the forward term rate is lognormal

(more on this below). HJM considers the infinitesimal forward rate not the forward term rate, but one might hope it could also be lognormal. However there's a problem. If we assume $\sigma(t, T) = \sigma_0 f(t, T)$, the HJM theory tells us the risk-neutral process has a non-constant drift term

$$\alpha(t, T) = \sigma(t, T) \int_t^T \sigma(t, u) du = \sigma_0^2 f(t, T) \int_t^T f(t, u) du.$$

Since the drift is random (it depends on f) the solution will not be lognormal. But worse: the drift is *quadratic* in f . As a result, one can show that the associated HJM stochastic differential equation experiences finite-time blowup with positive probability. (Heuristically: it blows up for the same reason that the solution of $dy = y^2 dt$ becomes infinite in finite time; for an honest analysis of this issue for HJM see Section 13.6 of Avellaneda and Laurence.) By the way, other simple hypotheses on σ can also lead to problems. For example if we set $\sigma(t, T) = \sigma_0 f^\alpha(t, T)$ with $0 < \alpha < 1$ we must worry about the possibility that f reaches 0 in finite time.

Multifactor HJM is no more difficult than one-factor HJM. For the reasons summarized just above, it is common to assume $\sigma(t, T)$ is deterministic. But it is also common to permit more than one factor, e.g. to assume

$$d_t f(t, T) = \alpha(t, T) dt + \sum_{i=1}^n \sigma_i(t, T) dw_i \tag{1}$$

where w_1, \dots, w_n are independent Brownian motions under the risk-neutral measure. The advantage of a multifactor model over a one-factor model is simple: when there is only one factor, the prices of bonds with different maturities, say $P(t, T)$ and $P(t, S)$, are perfectly correlated. This is obviously not the case in the real world (a change in the short-term interest environment will, in practice, not have an entirely predictable effect on the yields of 30-year treasuries.) Permitting more one factor captures this effect. In practice one usually uses just a few factors (two or three); the coefficients $\sigma_i(t, T)$ may be chosen for analytical tractability (Baxter-Rennie gives an example in Section 5.7), or they may be chosen using a principal-component analysis of historical data (Section 13.4 of Avellaneda-Laurence discusses how this works and what it produces). Whatever the form of σ_i , they determine the drift in (1) by essentially the same argument we used in the one-factor case:

$$\alpha(t, T) = \sum_{i=1}^n \sigma_i(t, T) \int_t^T \sigma_i(t, u) du.$$

Pricing using HJM must usually be done computationally, via Monte Carlo simulation. Short-rate trees are not practical (except for a few special cases, like Hull-White) because the short-rate process is usually non-Markovian (so the associated tree cannot recombine). One might consider a tree approximation for the evolution of $f(t, T)$ with T fixed, if e.g. $\sigma(t, T)$ is deterministic; but remember the drift term couples all maturities, so this requires considering as many trees as you have maturities – not a happy thought. By the way, correlations between different maturities are more important when

considering swaptions than when considering caps, because a cap is a sum of its caplets (each of which can be valued independently) but a swaption is an option on a basket of payments, occurring at different times – so the correlations between the interest rates at each payment date are relevant. (Our trick for pricing swaptions using Hull-White avoided this issue, but the argument – due to Jamshidian – only works for one-factor short-rate models.)

The real problem with HJM however is that it does not make sufficient contact with (a) *market practice*, and (b) *market data*. Indeed: the simplest and most widely-used method of valuing caps is Black’s formula. There’s plenty of data concerning the implied vols associated with this formula. An ideal theory would be one that’s consistent with Black’s formula and easy to calibrate using the associated implied vols. (The other widely-quoted instruments are swaptions. One could just as well request a theory that’s consistent with Black’s formula for swaptions, and easy to calibrate using the associated implied vols. Unfortunately one must choose: there is a version of the Libor Market Model that’s consistent with Black’s formula for caps, and one that’s consistent with Black’s formula for swaptions, but none that has both features simultaneously. We’ll focus here on caps as the basic instrument.)

The Libor Market Model is the “ideal theory” just requested. It can be viewed as a special case of HJM, but it is much easier to derive it from scratch. As usual, I focus on the one-factor version because it captures all the main ideas in the most transparent form. The following discussion follows Section 24.3 of Hull’s 5th edition in both notation and content.

Let’s start by reviewing how Black’s formula is used to price caplets. A caplet is an option on the term interest rate available in the marketplace at time T_1 for a loan with maturity $T_2 = T_1 + \delta$. The payment is made at the end; therefore if the cap rate is K the holder receives

$$L\delta(R - K)_+$$

at time T_2 , where L is the notional principal and R is the term rate, defined by

$$\frac{1}{1 + \delta R} = P(T_1, T_2).$$

Since we want to value this instrument at a time $t < T_1$, it is natural to consider the associated *forward* term rate $F(t)$ for lending at T_1 with maturity T_2 . It is defined, for any $t \leq T_1$, by

$$\frac{1}{1 + \delta F(t)} = \frac{P(t, T_2)}{P(t, T_1)},$$

notice that $F(T_1) = R$. It is immediate from the definition that

$$F(t) = \frac{1}{\delta} \frac{(P(t, T_1) - P(t, T_2))}{P(t, T_2)}$$

so $F(t)$ is the price of a tradeable divided by $P(t, T_2)$. Therefore $F(t)$ is a martingale under the forward measure associated to time T_2 . (Indeed: by definition, this measure has the property that the value of any tradeable at time t normalized by $P(t, T_2)$ is a martingale under it.) Thus – always using the forward measure – the mean of $F(t)$ is independent of t . In particular, the mean of $F(T_1)$ is known in advance: it equals $F(t)$ at time t . Since we are working in the forward measure and the option is a tradeable, the option value satisfies

$$\frac{\text{option value at } t}{P(t, T_2)} = \text{expected payoff at time } T_2$$

i.e.

$$\text{option value at } t = P(t, T_2) \cdot L\delta E[(F(T_1) - K)_+].$$

When we use Black’s formula, we evaluate the expectation by making the *assumption* that $F(T_1)$ has lognormal statistics (under the forward measure associated with maturity T_2). The expectation is then given by a Black-Scholes-like formula, whose only inputs are the the mean $E[F(T_1)] = F(t)$ and the standard deviation of $\log F(T_1)$. In summary: to be consistent with Black’s formula, a theory must predict that each term rate has lognormal statistics under the forward measure associated to its maturity time; nothing else is needed.

OK, now the Libor Market Model. We want to consider many caplets with different maturities; for simplicity let’s suppose the maturities of interest are all multiples of a single parameter δ . So the present time is $t = 0$, and we’re only interested in term rates for lending at time $t_k = k\delta$ with maturity $t_{k+1} = (k + 1)\delta$. We now have many forward rates; let’s distinguish them notationally, writing

$$F_k(t) = \frac{1}{\delta} \frac{(P(t, t_k) - P(t, t_{k+1}))}{P(t, t_{k+1})} \quad (2)$$

for the term rate associated with (t_k, t_{k+1}) . The basic hypothesis we need for Black’s formula to hold is that these forward rates evolve with lognormal statistics; if there is only one factor this means

$$dF_k(t) = \zeta_k(t)F_k(t) dw_k \quad (3)$$

for each k , where w_k is Brownian motion under the martingale measure associated with maturity t_{k+1} and $\zeta_k(t)$ is a *deterministic* function of time. (We should of course choose $\zeta_k(t)$ to match the implied volatilities of forward term rates, obtained from market prices of caplets via Black’s formula.)

But do we have the right to do this? Our experience with HJM suggests that the answer is yes – in that setting the volatility $\sigma(t, T)$ was ours to choose. But it’s hard to be sure when we use a different measure for each forward rate. In HJM we checked consistency by writing everything in the risk-neutral measure; this corresponds of course to using the money-market fund as numeraire. In the present setting the money-market fund isn’t a natural object, because we’re working with term rates and a discrete set of maturities. Rather, the natural object is the *rolling CD*, which earns interest during the time interval (t_k, t_{k+1}) at term rate $F_k(t_k)$.

Every numeraire has an associated martingale measure. The one associated to the rolling CD has the property that the value of any option divided by the value of the rolling CD is a martingale. Hull calls this the *rolling forward risk-neutral measure*.

Our goal is now to express the evolution of all the forward rates F_k under the rolling-forward risk-neutral measure. To this end, we recall the change-of-numeraire calculation done on page 9 of Section 4: if Q is the risk-neutral measure and N is any tradeable, define σ_N by

$$dN = rN dt + \sigma_N N dw$$

where w is Brownian motion under the risk-neutral measure Q . Then the equivalent martingale measure Q_N associated with numeraire N is characterized by the property that

$$dw_N = -\sigma_N dt + dw$$

is a martingale under Q_N .

In the present setting we are interested in many numeraires, but none of them is the money market fund. No problem: just apply the preceding result to any pair of numeraires and subtract. We find that if M and N are two numeraires, then their martingale measures are related by the fact that

$$dw_N = (\sigma_M - \sigma_N) dt + dw_M$$

is a Brownian motion under Q_N if w_M is a Brownian motion under Q_M . Thus: when we change numeraire, we introduce a drift equal to the difference of the volatilities of the two numeraires.

We want to apply this with $M = P(t, t_k)$ and N equal to the rolling CD. Let's write $v_k(t)$ for the volatility of $P(t, t_k)$. What is the volatility of the rolling CD? Well, for any time t there is a unique "next maturity date" $t_{m(t)}$, and the rolling CD provides a known payment on this maturity date $t_{m(t)}$. So the volatility of the rolling CD is precisely the volatility of $P(t, t_{m(t)})$. In short: the volatility of the rolling CD at time t is

$$v_{m(t)}(t), \quad \text{where } m(t) \text{ is the first maturity date } \geq t.$$

OK, let's pull this together. Applying the change of measure result, the SDE (3) becomes

$$dF_k(t) = \zeta_k(t)[v_{m(t)}(t) - v_{k+1}(t)]F_k(t) dt + \zeta_k(t)F_k(t) dz$$

where z is Brownian motion in the rolling forward risk-neutral world.

Finally, we need to know the relation between $v_{m(t)}(t)$ and $v_{k+1}(t)$. We get it from the relation between bond prices and forward rates (2). Rearranging then taking the logarithm gives

$$\log P(t, t_k) - \log P(t, t_{k+1}) = \log(1 + \delta F_k(t)).$$

Applying Ito's formula, using the definition

$$d_t P(t, t_k) = rP(t, t_k)dt + v_k(t)P(t, t_k) dw = (\text{stuff}) dt + v_k(t)P(t, t_k) dz$$

and matching the dz terms we get

$$v_k(t) - v_{k+1}(t) = \frac{\delta F_k(t)\zeta_k(t)}{1 + \delta F_k(t)}.$$

Thus the laws (3) can be expressed simultaneously in a single measure – the rolling risk-neutral measure – as

$$dF_k = \zeta_k F_k \left(\sum_{i=m(t)}^k \frac{\delta F_i(t)\zeta_i(t)}{1 + \delta F_i(t)} \right) dt + \zeta_k F_k dz. \quad (4)$$

This SDE is *equivalent* to the separate equations (3). It expresses the fact that our hypotheses are consistent. Indeed, to satisfy our hypothesis that $dF_k = \zeta_k(t)dw_k$ with w_k a Brownian motion under the forward measure associated with maturity t_k for each k , we need merely solve the single SDE (4), under a single probability measure (associated with the rolling CD). Then the relations (3) are all recovered by reversing the changes of measure done above.

Notice that (4) is simply a discrete analogue of the HJM equation

$$d_t f(t, T) = \alpha(t, T) dt + \sigma(t, T) dw$$

with

$$\alpha(t, T) = \sigma(t, T) \int_t^T \sigma(t, u) du$$

and $\sigma(t, T) = \zeta_i(t)f_i(t)$ when $T \approx t_i$. Indeed: formally as $\delta \rightarrow 0$ we expect $F_k(t) \approx f(t, t_k)$ since f is the infinitesimal forward rate. And formally as $\delta \rightarrow 0$ we have $1 + \delta F_i(t) \approx 1$ while

$$\sum_{i=m(t)}^k \delta F_i(t)\zeta_i(t) \approx \int_t^T f(t, u)\zeta(t, u) du.$$

Thus the Libor Market Model is essentially the discrete analogue of HJM with lognormal forward rates.

Question: HJM with lognormal statistics was problematic due to blowup. Why don't we have the same trouble here? Answer: at finite δ the drift term

$$\zeta_k(t)F_k(t) \sum_{i=m(t)}^k \frac{\delta F_i(t)\zeta_i(t)}{1 + \delta F_i(t)}$$

is *not* quadratic in F ; rather it has linear growth for large values of F_k due to the denominator $1 + \delta F_i(t)$. So blowup is not expected (and indeed does not occur).

The fact that the model makes sense is nice of course. But the real point is that it is easy to make this model reproduce the predictions of Black's formula, by simply choosing the functions $\zeta_k(t)$ appropriately. That's how the model is used in practice: choose the $\zeta_k(t)$ (typically piecewise constant) to match market data on caps; then use the theory to price more exotic or complicated instruments for which no Black-like formula is possible.

I recommend reading Hull's section 24.3, which goes much farther than the preceding discussion. In particular he addresses (a) the multifactor version of the model; (b) analytic approximations for the prices of swaptions [useful for calibration, along with data on caps]; (c) some typical applications e.g. to mortgage-backed securities.