
6. The LIBOR and Swap Market Models (LFM and LSM)

*There is so much they don't tell. Is that why I became a cop?
To learn what they don't tell?*

John Jones in "Martian Manhunter: American Secrets" 1, DC Comics,
1992

6.1 Introduction

In this chapter we consider one of the most popular and promising families of interest-rate models: The market models.

Why are such models so popular? The main reason lies in the agreement between such models and well-established market formulas for two basic derivative products. Indeed, the lognormal forward-LIBOR model (LFM) prices caps with Black's cap formula, which is the standard formula employed in the cap market. Moreover, the lognormal forward-swap model (LSM) prices swaptions with Black's swaption formula, which again is the standard formula employed in the swaption market. Since the caps and swaptions markets are the two main markets in the interest-rate-options world, it is important for a model to be compatible with such market formulas.

Before market models were introduced, there was no interest-rate dynamics compatible with either Black's formula for caps or Black's formula for swaptions. These formulas were actually based on mimicking the Black and Scholes model for stock options under some simplifying and inexact assumptions on the interest-rates distributions. The introduction of market models provided a new derivation of Black's formulas based on rigorous interest-rate dynamics.

However, even with full rigor given separately to the caps and swaptions classic formulas, we point out the now classic problem of this setup: The LFM and the LSM are not compatible. Roughly speaking, if forward LIBOR rates are lognormal each under its measure, as assumed by the LFM, forward swap rates cannot be lognormal at the same time under their measure, as assumed by the LSM. Although forward swap rates obtained from lognormal forward LIBOR rates are not far from being lognormal themselves under the relevant measure (as shown in some empirical works and also in Chapter 8

of the present book), the problem still stands and reduces one's enthusiasm for the theoretical setup of market models, if not for the practical one.

In this chapter, we will derive the LFM dynamics under different measures by resorting to the change-of-numeraire technique. We will show how caps are priced in agreement with Black's cap formula, and explain how swaptions can be priced through a Monte Carlo method in general. Analytical approximations leading to swaption-pricing formulas are also presented, as well as closed-form formulas for terminal correlations based on similar approximations.

We will suggest parametric forms for the instantaneous covariance structure (volatilities and correlations) in the LFM. Part of the parameters in this structure can be obtained directly from market-quoted cap volatilities, whereas other parameters can be obtained by calibrating the model to swaption prices. The calibration to swaption prices can be made computationally efficient through the analytical approximations mentioned above.

We will derive results and approximations connecting semi-annual caplet volatilities to volatilities of swaptions whose underlying swap is one-year long. We will also show how one can obtain forward rates over non-standard periods from standard forward rates either through drift interpolation or via a bridging technique.

We will then introduce the LSM and show how swaptions are priced in agreement with Black's swaptions formula, although Black's formula for caps does not hold under this model.

We will finally consider the "smile problem" for the cap market, and introduce some possible extensions of the basic LFM that are analytically tractable and allow for a volatility smile.

6.2 Market Models: a Guided Tour

Before market models were introduced, short-rate models used to be the main choice for pricing and hedging interest-rate derivatives. Short-rate models are still chosen for many applications and are based on modeling the instantaneous spot interest rate ("short rate") via a (possibly multi-dimensional) diffusion process. This diffusion process characterizes the evolution of the complete yield curve in time. We have seen examples of such models in Chapters 3 and 4.

To fix ideas, let us consider the time-0 price of a T_2 -maturity caplet resetting at time T_1 ($0 < T_1 < T_2$) with strike X and a notional amount of 1. Caplets and caps have been defined in Chapter 1 and will be described more generally in Section 6.4. Let τ denote the year fraction between T_1 and T_2 . Such a contract pays out the amount

$$\tau(L(T_1, T_2) - X)^+$$

at time T_2 , where in general $L(u, s)$ is the LIBOR rate at time u for maturity s .

Again to fix ideas, let us choose a specific short-rate model and assume we are using the shifted two-factor Vasicek model G2++ given in (4.4). The parameters of this two-factor Gaussian additive short-rate model are here denoted by $\theta = (a, \sigma, b, \eta, \rho)$. Then the short rate r_t is obtained as the sum of two linear diffusion processes x_t and y_t , plus a deterministic shift φ that is used for fitting the initially observed yield curve at time 0:

$$r_t = x_t + y_t + \varphi(t; \theta).$$

Such model allows for an analytical formula for forward LIBOR rates F ,

$$\begin{aligned} F(t; T_1, T_2) &= F(t; T_1, T_2; x_t, y_t, \theta), \\ L(T_1, T_2) &= F(T_1; T_1, T_2; x_{T_1}, y_{T_1}, \theta). \end{aligned}$$

At this point one can try and price a caplet. To this end, one can compute the risk-neutral expectation of the payoff discounted with respect to the bank account numeraire $\exp\left(\int_0^{T_2} r_s ds\right)$ so that one has

$$E \left[\exp \left(- \int_0^{T_2} r_s ds \right) \tau(F(T_1; T_1, T_2, x_{T_1}, y_{T_1}, \theta) - X)^+ \right].$$

This too turns out to be feasible, and leads to a function

$$U_C(0, T_1, T_2, X, \theta).$$

On the other hand, the market has been pricing caplets (actually caps) with Black's formula for years. One possible derivation of Black's formula for caplets is based on the following approximation. When pricing the discounted payoff

$$E \left[\exp \left(- \int_0^{T_2} r_s ds \right) \tau(L(T_1, T_2) - X)^+ \right] = \dots$$

one first assumes the discount factor $\exp\left(-\int_0^{T_2} r_s ds\right)$ to be deterministic and identifies it with the corresponding bond price $P(0, T_2)$. Then one factors out the discount factor to obtain:

$$\dots \approx P(0, T_2) \tau E \left[(L(T_1, T_2) - X)^+ \right] = P(0, T_2) \tau E \left[(F(T_1; T_1, T_2) - X)^+ \right].$$

Now, inconsistently with the previous approximation, one goes back to assuming rates to be stochastic, and models the forward LIBOR rate $F(t; T_1, T_2)$ as in the classical Black and Scholes option pricing setup, i.e as a (driftless) geometric Brownian motion:

$$dF(t; T_1, T_2) = vF(t; T_1, T_2)dW_t, \quad (6.1)$$

where v is the instantaneous volatility, assumed here to be constant for simplicity, and W is a standard Brownian motion under the risk-neutral measure Q .

Then the expectation

$$E[(F(T_1; T_1, T_2) - X)^+]$$

can be viewed simply as a T_1 -maturity call-option price with strike X and whose underlying asset has volatility v , in a market with zero risk-free rate.

We therefore obtain:

$$\begin{aligned} \mathbf{Cpl}(0, T_1, T_2, X) &:= P(0, T_2)\tau E(F(T_1; T_1, T_2) - X)^+ \\ &= P(0, T_2)\tau[F(0; T_1, T_2)\Phi(d_1(X, F(0; T_1, T_2), v\sqrt{T_1})) \\ &\quad - X\Phi(d_2(X, F(0; T_1, T_2), v\sqrt{T_1}))], \end{aligned}$$

$$\begin{aligned} d_1(X, F, u) &= \frac{\ln(F/X) + u^2/2}{u}, \\ d_2(X, F, u) &= \frac{\ln(F/X) - u^2/2}{u}, \end{aligned}$$

where Φ is the standard Gaussian cumulative distribution function.

From the way we just introduced it, this formula seems to be partly based on inconsistencies. However, within the change-of-numeraire setup, the formula can be given full mathematical rigor as follows. Denote by Q^2 the measure associated with the T_2 -bond-price numeraire $P(\cdot, T_2)$ (T_2 -forward measure) and by E^2 the corresponding expectation. Then, by the change-of-numeraire approach, we can switch from the bank-account numeraire $\exp\left(\int_0^t r_s ds\right)$ associated with risk-neutral measure Q to the bond-price numeraire $P(t, T_2)$ and obtain

$$E\left[\exp\left(-\int_0^{T_2} r_s ds\right)\tau(L(T_1, T_2) - X)^+\right] = P(0, T_2)E^2[\tau(L(T_1, T_2) - X)^+].$$

What has just been done, rather than assuming deterministic discount factors, is a change of measure. We have “factored out” the stochastic discount factor and replaced it with the related bond price, but in order to do so we had to change the probability measure under which the expectation is taken. Now the last expectation is no longer taken under the risk-neutral measure but rather under the T_2 -forward measure. Since by definition $F(t; T_1, T_2)$ can be written as the price of a tradable asset divided by $P(t, T_2)$, it needs follow a martingale under the measure associated with the numeraire $P(t, T_2)$, i.e. under Q^2 . As we have hinted at in Appendix A, martingale means “driftless” when dealing with diffusion process. Therefore, the dynamics of $F(t; T_1, T_2)$ under Q^2 is driftless, so that the dynamics

$$dF(t; T_1, T_2) = vF(t; T_1, T_2)dW_t \quad (6.2)$$

is correct under the measure Q^2 , where W is a standard Brownian motion under Q^2 . Notice that the driftless (lognormal) dynamics above is precisely the dynamics we need in order to recover *exactly* Black's formula, without approximation. *We can say that the choice of the numeraire $P(\cdot, T_2)$ is based on this fact: It makes the dynamics (6.2) of F driftless under the related Q^2 measure, thus replacing rigorously the earlier arbitrary assumption on the F dynamics (6.1) under the risk-neutral measure Q .* Following this rigorous approach we indeed obtain Black's formula, since the process F has the same distribution as in the approximated case above, and hence the expected value has the same value as before.

The example just introduced is a simple case of what is known as "log-normal forward-LIBOR model". It is known also as Brace-Gatarek-Musiela (1997) model, from the name of the authors of one of the first papers where it was introduced rigorously. This model was also introduced by Miltersen, Sandmann and Sondermann (1997). Jamshidian (1997) also contributed significantly to its development. At times in the literature and in conversations, especially in Europe, the LFM is referred to as "BGM" model, from the initials of the three above authors. In other cases, colleagues in the U.S. called it simply an "HJM model", related perhaps to the fact that the BGM derivation was based on the HJM framework rather than on the change-of-numeraire technique. However, a common terminology is now emerging and the model is generally known as "LIBOR Market Model". We will stick to "Lognormal Forward-LIBOR Model" (LFM), since this is more informative on the properties of the model: Modeling forward LIBOR rates through a lognormal distribution (under the relevant measures).

Let us now go back to our short-rate model formula U_C and ask ourselves whether this formula can be compatible with the above reported Black's market formula. It is well known that the two formulas are not compatible. Indeed, by the two-dimensional version of Ito's formula we may derive the Q^2 -dynamics of the forward LIBOR rate between T_1 and T_2 under the short-rate model,

$$dF(t; T_1, T_2; x_t, y_t, \theta) = \frac{\partial F}{\partial(t, x, y)} d[t \ x_t \ y_t]' + \frac{1}{2} d[x_t \ y_t] \frac{\partial^2 F}{\partial^2(x, y)} d[x_t \ y_t]', \quad , Q^2, \quad (6.3)$$

where the Jacobian vector and the Hessian matrix have been denoted by their partial derivative notation. This dynamics clearly depends on the linear-Gaussian dynamics of x and y under the T_2 -forward measure. The thus obtained dynamics is easily seen to be incompatible with the lognormal dynamics leading to Black's formula. More specifically, for no choice of the parameters θ does the distribution of the forward rate F in (6.3) produced by the short-rate model coincide with the distribution of the "Black"-like forward rate F following (6.2). In general, no known short-rate model can lead to Black's formula for caplets (and more generally for caps).

What is then done with short-rate models is the following. After setting the deterministic shift φ so as to obtain a perfect fit of the initial term structure, one looks for the parameters θ that produce caplet (actually cap) prices U_C that are closest to a number of observed market cap prices. The model is thus calibrated to (part of) the cap market and should reproduce well the observed prices. Still, the prices U_C are complicated nonlinear functions of the parameters θ , and this renders the parameters themselves difficult to interpret. On the contrary, the parameter v in the above “market model” for F has a clear interpretation as a lognormal percentage (instantaneous) volatility, and traders feel confident in handling such kind of parameters.

When dealing with several caplets involving different forward rates, different structures of instantaneous volatilities can be employed. One can select a different v for each forward rate by assuming each forward rate to have a constant instantaneous volatility. Alternatively, one can select piecewise-constant instantaneous volatilities for each forward rate. Moreover, different forward rates can be modeled as each having different random sources W that are instantaneously correlated. Modeling correlation is necessary for pricing payoffs depending on more than a single rate at a given time, such as swaptions. Possible volatility and correlation structures are discussed in Sections 6.3.1 and 6.9. The implications of such structures as far as caplets and caps are concerned are discussed in Section 6.4, and their consequences on the term structure of volatilities as a whole are discussed in Section 6.5.

As hinted at above, the model we briefly introduced is the market model for “half” of the interest-rate-derivatives world, i.e. the cap market. But what happens when dealing with basic products from the other “half” of this world, such as swaptions? Swaptions are options on interest-rate swaps. Interest rate swaps and swaptions have been defined in Chapter 1 and will be again described in Section 6.7. Swaptions are priced by a Black-like market formula that is, in many respects, similar to the cap formula. This market formula can be given full rigor as in the case of the caps formula. However, doing so involves choosing a numeraire under which the relevant forward *swap* rate (rather than a particular forward LIBOR rate) is driftless and lognormal. This numeraire is indeed different from any of the bond-price numeraires used in the derivation of Black’s formula for caps according to the LFM. The obtained model is known as “lognormal (forward) swap model” (LSM) and is also referred to as the Jamshidian (1997) market model or “swap market model”.

One may wonder whether the two models are distributionally compatible or not, similarly to our previous comparison of the LFM with the shifted two-factor Vasicek model G2++. As before, the two models are incompatible. If we adopt the LFM for caps we cannot recover the market formula given by the LSM for swaptions. The two models (LFM and LSM) collide. We will point out this incompatibility in Section 6.8.

There are some recent works investigating the “size” of the discrepancy between these two models, and we will address this issue in Chapter 8 by comparing swap-rates distributions under the two models. Results seem to suggest that the difference is not large in most cases. However, the problem remains of choosing either of the two models for the whole market.

When the choice is made, the half market consistent with the model is calibrated almost automatically, see for example the cap calibration with the LFM in Section 6.4. But one has still the problem of calibrating the chosen model to the remaining half, e.g. the swaption market in case the LFM is adopted for both markets.

Brace, Dun and Barton (1998) indeed suggest to adopt the LFM as central model for the two markets, mainly for its mathematical tractability. We will stick to their suggestion, also due to the fact that forward rates are somehow more natural and more representative coordinates of the yield-curve than swap rates. Indeed, it is more natural to express forward swap rates in terms of a suitable preselected family of LIBOR forward rates, rather than doing the converse.

We are now left with the problem of finding a way to compute swaption prices with the LFM. In order to understand the difficulties of this task, let us consider a very simple swaption. Assume we are at time 0. The underlying interest-rate swap (IRS) starts at T_1 and pays at T_2 and T_3 . All times are equally spaced by a year fraction denoted by τ , and we take a unit notional amount. The (payer) swaption payoff can be written as

$$[P(T_1, T_2)\tau(F_2(T_1) - K) + P(T_1, T_3)\tau(F_3(T_2) - K)]^+ ,$$

where in general we set $F_k(t) = F(t; T_{k-1}, T_k)$. Recall that the discount factors $P(T_1, T_2)$ and $P(T_1, T_3)$ can be expressed in terms of $F_2(T_1)$ and $F_3(T_2)$, so that this payoff actually depends only on the rates $F_2(T_1)$ and $F_3(T_2)$. The key point, however, is the following. The payoff is not additively “separable” with respect to the different rates. As a consequence, when you take expectation of such a payoff, the *joint* distribution of the two rates F_2 and F_3 is involved in the calculation, so that the correlation between the two rates F_2 and F_3 has an impact on the value of the contract. This does not happen with caps. Indeed, let us go back to caps for a moment and consider a cap consisting of the T_2 - and T_3 -caplets. The cap payoff as seen from time 0 would be

$$\exp\left(-\int_0^{T_2} r_u du\right) \tau(F_2(T_1) - K)^+ + \exp\left(-\int_0^{T_3} r_u du\right) \tau(F_3(T_2) - K)^+ .$$

This time, the payoff is additively separated with respect to different rates. Indeed, we can compute

$$\begin{aligned}
& E \left[\exp \left(- \int_0^{T_2} r_u du \right) \tau(F_2(T_1) - K)^+ \right. \\
& \quad \left. + \exp \left(- \int_0^{T_3} r_u du \right) \tau(F_3(T_2) - K)^+ \right] \\
&= E \left[\exp \left(- \int_0^{T_2} r_u du \right) \tau(F_2(T_1) - K)^+ \right] \\
& \quad + E \left[\exp \left(- \int_0^{T_3} r_u du \right) \tau(F_3(T_2) - K)^+ \right] \\
&= P(0, T_2) \tau E^2 [(F_2(T_1) - K)^+] + P(0, T_3) \tau E^3 [(F_3(T_2) - K)^+] .
\end{aligned}$$

In this last expression we have two expectations, each one involving a *single* rate. The joint distribution of the two rates F_2 and F_3 is not involved, therefore, in the calculation of this last expression, since it is enough to know the marginal distributions of F_2 and F_3 separately. Accordingly, the terminal correlation between rates F_2 and F_3 does not affect this payoff.

As a consequence of this simple example, it is clear that adequately modeling correlation can be important in defining a model that can be effectively calibrated to swaption prices. When the number of swaptions to which the model has to be calibrated is large, correlation becomes definitely relevant. If a short-rate model has to be chosen, it is better to choose a multi-factor model. Multi-dimensional instantaneous sources of randomness guarantee correlation patterns among terminal rates F that are clearly more general than in the one-factor case. However, producing a realistic correlation pattern with, for instance, a two-factor short-rate model is not always possible.

As far as the LFM is concerned, as we hinted at above, the solution is usually to assign a different Brownian motion to each forward rate and to assume such Brownian motions to be instantaneously correlated. Manipulating instantaneous correlation leads to manipulation of correlation of simple rates (terminal correlation), although terminal correlation is also influenced by the way in which average volatility is distributed among instantaneous volatilities, see Section 6.6 on this point.

Choosing an instantaneous-correlation structure flexible enough to express a large number of swaption prices and, at the same time, parsimonious enough to be tractable is a delicate task. The integrated covariance matrix need not have full rank. Usually ranks two or three are sufficient, see Brace Dun and Barton (1998) on this point. We will suggest some parametric forms for the covariance structure of instantaneous forward rates in Section 6.9.

Brace (1996) proposed an approximated formula to evaluate analytically swaptions in the LFM. The formula is based on a rank-one approximation of the integrated covariance matrix plus a drift approximation and works well in specific contexts and in non-pathological situations. More generally,

a rank- r approximation is considered in Brace (1997). See again Brace, Dun and Barton (1998) for numerical experiments and results.

The rank-one approximated formula is reviewed in Section 6.11, while the rank- r approximation is reviewed in Section 6.12. There are also analytical swaption-pricing formulas that are simpler and still accurate enough for most practical purposes. Such formulas are based on expressing forward swap rates as linear combinations of forward LIBOR rates, to then take variance on both sides and integrate while freezing some coefficients. This “freezing the drift and collapsing all measures” approximation is reviewed in Section 6.13 and has also been tested by Brace, Dun and Barton (1998). Given its importance, we performed numerical tests of our own. These are reported in Section 8.2 of Chapter 8, and confirm that the formula works well.

More generally, to evaluate swaptions and other payoffs with the LFM one has usually to resort to Monte Carlo simulation. Once the numeraire is chosen, one simulates all forward rates involved in the payoff by discretizing their joint dynamics with a numerical scheme for stochastic differential equations (SDEs).

Notice that each forward rate is driftless only under its associated measure. Indeed, for example, while F_2 is driftless under Q^2 (recall (6.2)), it is not driftless under Q^1 . The dynamics of F_2 under Q^1 is derived below and leads to a process that we need to discretize in order to obtain simulations, whereas in the driftless case (6.2) the transition distribution of F is known to be lognormal and no numerical scheme is needed. This point is discussed in detail again in Section 6.10. As an introductory illustration, consider again the above swaption payoff, now discounted at time 0, and take its risk-neutral expectation:

$$E \left\{ \exp \left(- \int_0^{T_1} r_u du \right) [P(T_1, T_2)\tau(F_2(T_1) - K) + P(T_1, T_3)\tau(F_3(T_2) - K)]^+ \right\} = \dots$$

Take $P(\cdot, T_1)$ as numeraire, which corresponds to choose the T_1 -forward-adjusted measure Q^1 . One then obtains:

$$\dots = P(0, T_1)E^1 [P(T_1, T_2)\tau(F_2(T_1) - K) + P(T_1, T_3)\tau(F_3(T_2) - K)]^+ .$$

It is known that F_1 is driftless under the chosen measure Q^1 , while F_2 and F_3 are not. For example, take for simplicity the one-factor case, so that

$$dF_2(t) = v_2 F_2(t) dW_t, \tag{6.4}$$

where W is now a standard Brownian motion under Q^2 . Then if Z is a standard Brownian motion under Q^1 , it can be shown that the change-of-numeraire technique leads to

$$dF_2(t) = \frac{v_2^2 F_2^2(t) \tau}{1 + \tau F_2(t)} dt + v_2 F_2(t) dZ_t. \quad (6.5)$$

Now it is clear that, as we stated above, the no-arbitrage dynamics (6.5) of F_2 under Q^1 is not driftless, and moreover its transition distribution (which is necessary to perform exact simulations) is not known, contrary to the Q^2 driftless case (6.4).

In order to price the swaption, it is the Q^1 dynamics of F_2 that matters, so that we need to discretize the related equation (6.5) in order to be able to simulate it. Numerical schemes for SDEs are available, like the Euler or Milstein scheme, see also Appendix A. Alternative schemes that guarantee the (weak) no-arbitrage condition to be maintained in discrete time and not just in the continuous-time limit have been proposed by Glasserman and Zhao (2000).

Whichever scheme is chosen, Monte Carlo pricing is to be performed by simulating the relevant forward LIBOR rates. While path-dependent derivatives can be priced by this approach in general, as far as early-exercise (e.g. American-style or Bermudan-style) products are concerned, the situation is delicate, since the joint dynamics of the LFM usually does not lead to a recombining lattice for the short rate, so that it is not immediately clear how to evaluate a Bermudan- or American-style product with a tree in the LFM.

Usually, ad-hoc techniques are needed, such as Carr and Yang's (1997) who provide a method for simulating Bermudan-style derivatives with the LFM via a Markov chain approximation, or Andersen's (1999) who approximates the early exercise boundary as a function of intrinsic value and "still-alive" nested European swaptions. However, a general method for combining backward induction with Monte Carlo simulation has been proposed recently by Longstaff and Schwartz (2000), and this method is rather promising, especially because of its generality. We briefly review all these methods in Chapter 10.

Several other problems remain, like for example the possibility of including a volatility smile in the model. We will address the smile problem for the cap market by proposing both the CEV dynamics suggested by Andersen and Andreasen (2000) with possibly shifted distributions, and the shifted "lognormal-mixture" dynamics model illustrated in Brigo and Mercurio (2000a, 2000b, 2001b).