

Bond and Option Pricing when Short Rates are Lognormal

This article describes a one-factor model for bond and option pricing that is based on the short-term interest rate and that allows the target rate, mean reversion and local volatility to vary deterministically through time. For any horizon, the distribution of possible short rates is lognormal, so the rate neither falls below zero nor reflects off a barrier at zero. A model like this allows one to match the yield curve, the volatility curve and the cap curve. Surprisingly, adding to future local volatility lowers the volatility curve.

A conventional binary tree with probabilities of 0.5 but variable time spacing is used to value bonds and options. When the inputs are constant, the slope of the yield curve starts out positive and ends up negative, while its curvature shifts from negative to positive. Even when mean reversion is zero, the volatility curve has a negative slope. The differential cap curve rises steeply (through the effects of volatility) and then falls steeply (through the effects of discounting and a falling forward rate used as the strike price).

THIS ARTICLE PRESENTS a one-factor model of bond prices, bond yields and related options. The single factor that is the source of all uncertainty is the short-term interest rate. We assume no taxes or transaction costs, no default risk and no extra costs for borrowing bonds. We also assume that all security prices are perfectly correlated in continuous time.

We can choose from among a number of models of the local process for the short-term interest rate—a normal process, a lognormal process, a “square-root” process or others.¹ The nominal interest rate cannot fall below zero as long as people can hold cash; it can become stuck at zero for long periods, however, as when prices fall persistently and substantially. None of the models we have to choose from allows for both these features. Lognormal models keep the rate away from zero entirely, while some square-root models make zero into a “reflecting barrier.”

The lognormal model we use is more general than others, because we allow the local process to change over time. So long as the process for

$\log r$ is linear in $\log r$ at each time, we will have a lognormal distribution for the possible values of the short rate at a given future time. In contrast, the square-root process does not give a square-root distribution at a given future time.

A Lognormal Model

A lognormal distribution has a mean and a variance. Assuming a different lognormal short-rate distribution for each future time allows both mean and variance to depend on time.

As Hull and White point out, however, a normal (or lognormal) model with mean reversion can depend on time in three ways, not just two ways.² In their notation, the continuous-time limit of the Black-Derman-Toy one-factor model is:³

$$d(\log r) = [\theta(t) - \phi(t) \log r]dt + \sigma(t)dz, \quad (1)$$

where r is the local interest rate and $\sigma(t)$ depends on $\phi(t)$.

To create a lognormal model that depends on time in three ways, we can simply drop the tie between $\sigma(t)$ and $\theta(t)$. Hull and White do this for a general model.⁴ Our model is a special case of theirs.

1. Footnotes appear at end of article.

We make one change in the way Hull and White write their model. We write $\mu(t)$ for the "target interest rate." When $\log r$ is above $\log \mu(t)$, it tends to fall, and when it is below $\log \mu(t)$, it tends to rise. Thus we rewrite Equation (1) as:

$$d(\log r) = \phi(t)[\log \mu(t) - \log r]dt + \sigma(t)dz. \quad (2)$$

We take $\mu(t)$ as the target rate, $\phi(t)$ as mean reversion and $\sigma(t)$ as local volatility in the expression for the local change in $\log r$. We choose these three functions ("inputs") to match three features of the world ("outputs").

For their outputs, Hull and White chose

- the yield curve;
- the volatility curve; and
- the future local volatilities $\sigma(t)$.

The yield curve gives for each maturity the current yield on a zero-coupon bond. The volatility curve gives for each maturity the current yield volatility on a zero-coupon bond. The

Table I Time Spacing

n/N	t_n
0	0.0
1/5	4.1
2/5	6.3
3/5	7.9
4/5	9.0
1	10.0

future local volatilities output is the same as the corresponding input.

For our outputs, we choose:

- the yield curve;
- the volatility curve; and
- the cap curve.

Our first two outputs are the same as Hull and White's. The cap curve gives, for each maturity, the price of an at-the-money differential cap. A differential cap pays at a rate equal to the difference (if positive) between the short rate

Figure A Zero-Coupon Yield (initial rate = 10%; rate volatility = 12%; mean reversion = 0.02)





Figure B Zero-Coupon Yield Volatility (initial and target rates = 10%; mean reversion = 0.02)



and the strike price. For any maturity, an at-the-money cap has a strike equal to the forward rate for that maturity. A full cap is the integral of differential caps over all future horizons up to the full cap's maturity.

An advantage of these outputs over Hull and White's is that all of them are (in principle) observable. They all correspond to market prices.

We do not claim that our inputs imply a reasonable process for the short rate, except as a rough approximation. We choose inputs that give reasonable outputs, though they may be somewhat unreasonable themselves.⁵ We are following in the footsteps of Cox, Ross and Rubinstein, who value options by generating a "risk-neutral" distribution of future stock prices.⁶ This is not a true distribution, but it nonetheless gives correct option prices.

Input and Output Volatilities

We were surprised at the relation we found between input and output volatilities. The input


volatilities are the local volatilities for the short rate at all horizons. The output volatilities are the current yield volatilities for zero-coupon bonds at all maturities.

Imagine that we hold the local short-rate volatilities fixed out to a certain time, but raise the local volatilities after that time. We hold all other inputs fixed. For our lognormal model, raising the future input volatilities will, if anything, lower the output yield volatilities.

One way to raise the output volatilities, then, is to lower the future input volatilities. Another way is to reduce the amount of mean reversion. To arrive at very high output volatilities, we can even turn to negative mean reversion.

When all the inputs are constant in a lognormal model, the volatility curve will decline. The higher the local short-rate volatility, and the greater the mean reversion, the faster it will decline. A declining volatility curve is a persistent feature of the world. We can model this feature in a lognormal model without using the time-dependence of our inputs.

Figure C Zero-Coupon Yield Volatility (initial and target rates = 10%; rate volatility = 12%)



Building a Tree

Black, Derman and Toy show how to build a binomial tree for some lognormal models.⁷ They use the location and spacing of the nodes for each future time to vary the inputs. They are able to match two of the outputs (yield curve and volatility curve), but not the third (cap curve). In fact, with their models, choosing a yield curve and volatility curve implies choosing a cap curve. They cannot vary the target rate, local volatility and mean reversion separately.

Hull and White solve this problem by using a trinomial tree, rather than a binomial tree.⁸ From each node of a trinomial tree, you can move to one of three adjacent nodes one period later. What those three nodes are depends on the problem at hand. So do the probabilities associated with the three nodes.

We solve the problem by varying the spacing in the tree. This gives us another degree of freedom, so we can vary all three inputs within a simple binary tree. We can continue to assume that the probabilities of up and down moves are identical and both equal to 0.5. This helps make

our use of the tree efficient. We preserve the topology of the simple binary-tree method.

When mean reversion is positive, this method has a problem of its own—the spacing declines over time. For a reasonable number of nodes, the time separation of the early branches can be so large that we don't have the detail we want for applications such as valuing short-term options on long-term bonds.


We may thus need to prune the tree as we build it. We build it out to a certain point and then chop off half the nodes. When we are working back in the tree, we interpolate and extrapolate when we come to a place where the number of nodes doubles.

Possible Outputs

We can't use our process to match any output curves we might write down. For example, we can't match a yield curve that implies negative forward rates, or a discontinuous volatility curve.

We thus need to use smooth output curves. If we see bumpy curves in the world, we may be

Figure D Differential at-the-money Cap (initial and target rates = 10%; mean reversion = 0.02)



seeing swap or arbitrage opportunities, or we may be seeing data errors. In either case, we need to smooth if we hope to match successfully.

A "normal" model that allows negative interest rates may be easier to match. But we prefer our lognormal models, because curves we can match only with a "normal" model present profit opportunities.

With Known Inputs

Suppose we know the functions $\mu(t)$, $\phi(t)$ and $\sigma(t)$ (the target rate, mean reversion and local volatility). How can we build a tree to fit these functions?

Given the values of $\log r$, the tree will be rectilinear. For a given time, it will have equal spacing, though the time spacing will vary. The spacing for a given time will fit local volatility. The drift of the points from one time to the next will fit the target rate. And the time spacing will fit mean reversion.

Write τ_n for $t_{n+1} - t_n$, ϕ_n for $\phi(t_n)$ and σ_n for $\sigma(t_n)$. Then the formula for mean reversion is:

$$\phi_n = \frac{1}{\tau_n} \left(1 - \frac{\sigma_n \sqrt{\tau_n}}{\sigma_{n-1} \sqrt{\tau_{n-1}}} \right). \quad (3)$$

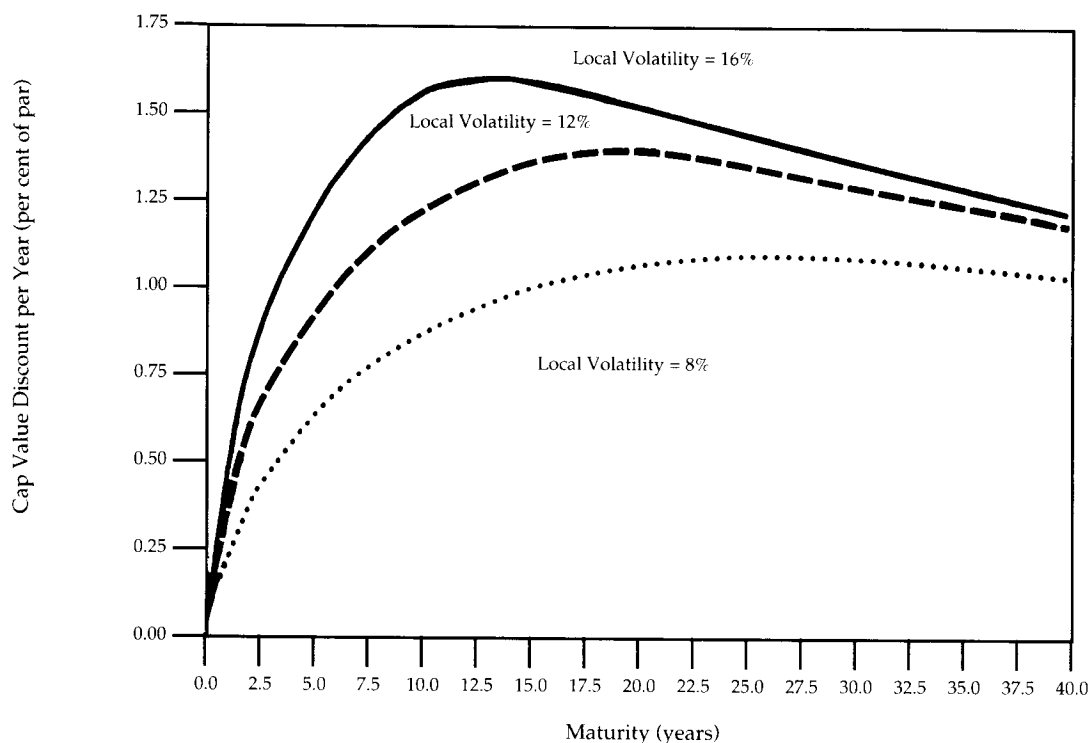
Solving this formula for τ_n , we have:

$$\tau_n = \tau_{n-1} \frac{4 \left(\frac{\sigma_{n-1}}{\sigma_n} \right)^2}{\left(1 + \sqrt{1 + 4\phi_n \left(\frac{\sigma_{n-1}}{\sigma_n} \right)^2 \tau_{n-1}} \right)^2}. \quad (4)$$

We can choose τ_0 as we wish. The smaller it is, the finer the tree and the more accurate the answers. We have σ_0 from $\sigma(0)$, σ_1 from $\sigma(t_0)$ and ϕ_1 from $\phi(t_0)$. We can use Equation (4) to find τ_1 and then repeat. We will gradually build up the full tree. At each time, we will use $\mu(t)$ to help locate the first point.

Suppose, for example, that mean reversion is constant at 0.1, that local volatility is a constant annual 0.20, and that we divide a 10-year period into $N = 160$ subperiods. Table I gives the resulting time spacing.

Figure E Differential at-the-money Cap Divided by Discount
(initial and target rates = 10%; mean reversion = 0.02)



Recall that t_n is years until time interval n . Note how the spacing declines over time with positive mean reversion. The level of mean reversion in the table is high (0.1) to exaggerate its effect on time spacing.

With Known Outputs

Suppose we want to choose the inputs [$\mu(t)$, $\phi(t)$ and $\sigma(t)$] to match known outputs (yield curve, volatility curve and cap curve). How can we do it?

We divide time into segments and each segment into many time intervals. We choose μ , ϕ and σ to match the outputs at the end of the first segment. Then we choose μ , ϕ and σ from the start to the end of the second segment to match the output at the end of the second segment, and so on.


We might call the resulting inputs "implied target rate," "implied mean reversion" and "implied local volatility." We do not say that the short rate follows a process with these features. We do say that securities behave *as if* we lived in a one-factor world and the short rate followed this process.

After a time, we can estimate the model again and get a new process. The implied process changes each time security prices and volatilities change.

We go through similar steps when we figure implied volatility: When the option price changes, the implied volatility changes. When we value the option, we are assuming that its volatility is known and constant. But a minute later, we start using a new volatility. Similarly, we can value fixed income securities by assuming we know the one-factor short-rate process. A minute later, we start using a new process that is not consistent with the old one.

Another approach is to search for an interest-rate process general enough that we can assume it is true and unchanging. It will have many variables and constants. We estimate the constants and hope that when we repeat the estimation, the variables change, but not the constants. While we may reach this goal, we don't know enough to use this approach today. For now, we must continue to reestimate simple models.

Figure F Differential at-the-money Cap Divided by Discount and Strike
(initial and target rates = 10%; mean reversion = 0.02)



Autocorrelation

Because the distribution of short rates at any horizon is lognormal, we need only a mean and a variance to describe it. Yet we need a target rate, mean reversion and local volatility to describe the short-rate process in our model. For a given distribution of short rates for every horizon, our model has a whole family of possible processes. How can this be?

It turns out that processes differ in the autocorrelations they imply for future short rates. When mean reversion is strong, a rise in the short rate above the target will probably be largely reversed before long. When it is weak, we won't see this.

If we have a narrow distribution of possible short rates in the future, we can infer that either mean reversion is strong or local volatility is low between now and then.

Examples

To see the relation between inputs and outputs, consider what happens when we keep all the inputs [target rate $\mu(t)$, mean reversion $\phi(t)$ and local volatility $\sigma(t)$] constant.

Figure A shows the yield curve for three cases. In each curve, mean reversion is constant at 0.02, local volatility is constant at 12% and the current short rate is 10%.

For the middle curve, the target rate is the same as the current rate. Note that the curve rises slightly up to about seven years and then falls off. This is a typical pattern when the target rate is constant and equal to the current rate. Note that the curvature eventually reverses from concave down to concave up. This is typical too.

The top curve shows what happens when the current rate is 10% and the target rate 15%. The bottom curve shows a current rate of 10% and a target rate of 5%. Even when the target rate is 15%, the yield curve never approaches 15%. In fact, for the case we show, it never goes above 11%. We don't know how to create a model that shows a consistently rising yield curve with reasonable assumptions—even by adding features that our model doesn't have.

Figure B shows the volatility curve for three cases. In each case, the current rate is 10%, the

target rate is constant at 10%, and mean reversion is constant at 0.02. The three curves show (from the top down) constant local volatilities of 16, 12 and 8%.

The volatility curve starts at the local volatility and declines. The higher the volatility, the faster it declines. This is also a typical pattern in real-world volatility curves: Yield volatilities for shorter-term bonds tend to be higher than yield volatilities for longer-term bonds.

Figure C shows that mean reversion affects the slope of the volatility curve. The higher the mean reversion, the more negative the slope. But the curve slopes down even when mean reversion is zero. Actually, the slope of the curve at zero maturity is proportional to the negative of the mean reversion.

Comparing Figures B and C, we see that we can match both the short end and the long end of the volatility curve even with constant inputs. We choose local volatility to match the short end and mean reversion to match the long end.

Figure D shows the differential cap curves for different levels of local volatility. For each curve, we hold all the inputs fixed and have the current rate equal to the target rate. Note that doubling the local volatility roughly doubles the cap values at every maturity.

For every level of local volatility, the differential cap curve rises sharply to a maximum and then falls sharply. The rise is due to a gain in the effect of volatility as we increase maturity. But two other factors overpower this—an increase in the discount factor and a fall in the forward rate, which is the cap's strike price.

Figures E and F analyze this. In Figure E, we divide the differential cap value by the discount factor (the value of a zero-coupon bond with corresponding maturity and unit face value). In Figure F, we divide by the strike price as well. The resulting curves rise smoothly with maturity, reflecting the impact of volatility.⁹ ■

Footnotes

1. For a normal process, see O. Vasicek, "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, November 1977, and F. Jamshidian, "An Exact Bond Option Formula," *Journal of Finance*, March 1989. For a lognormal process, see U. L. Dothan, "On the Term Structure of Interest Rates," *Journal of Financial Economics*, March 1978. For a square-root process, see J. C. Cox, J. E. Ingersoll, Jr. and S. A. Ross, "A Theory of the Term Structure of Interest Rates," *Econometrica*, March 1985; G. M. Constantinides, "Theory of the Term Structure of Interest Rates: The Squared Autoregressive Instruments Nominal
2. J. Hull and A. White, "Pricing Interest Rate Derivative Securities," *Review of Financial Studies* 3 (1990), pp. 573-592.
3. See Black, Derman and Toy, "A One-Factor Model," *op. cit.*
4. Hull and White, "New Ways with the Yield Curve," *op. cit.*
5. Cox, Ingersoll and Ross ("A Theory of the Term Structure," *op. cit.*) and Constantinides ("Theory of the Term Structure," *op. cit.*) imagine they are choosing a sensible process for the short rate and look at the resulting outputs, including option prices. Ho and Lee ("Term Structure Movements," *op. cit.*), on the other hand, believe their outputs are reasonable but do not claim that their inputs are reasonable. In fact, their inputs are quite unreasonable, as shown by P. H. Dybvig, "Bond and Bond Option Pricing Based on the Current Term Structure" (Working paper, February 1989).
6. J. C. Cox, S. A. Ross and M. Rubinstein, "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, September 1979.
7. Black, Derman and Toy, "A One-Factor Model," *op. cit.*
8. Hull and White, "New Ways with the Yield Curve," *op. cit.* and J. Hull and A. White, "One Factor Interest-Rate Models and the Valuation of Interest-Rate Derivative Securities" (Working paper, June 1990).
9. We thank Emanuel Derman and Francis Longstaff for their helpful comments.

Term Structure (SAINTS) Model" (Working paper, September 1990); F. Longstaff, "The Valuation of Options on Yields" (Working paper, February 1990); and F. Longstaff, "The Valuation of Options on Coupon Bonds" (Working paper, January 1990). Other processes are given in T. S. Y. Ho and S-B Lee, "Term Structure Movements and Pricing Interest Rate Contingent Claims," *Journal of Finance*, December 1986; D. Heath, R. Jarrow and A. Morton, "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation" (Working paper, October 1990); D. Heath, R. Jarrow and A. Morton, "Contingent Claim Valuation with a Random Evolution of Interest Rates," *Review of Futures Markets*, forthcoming; J. Hull and A. White, "New Ways with the Yield Curve," *Risk*, October 1990; and F. Black, E. Derman and W. Toy, "A One-Factor Model of Interest Rates and its Application to Treasury Bond Options," *Financial Analysts Journal*, January/February 1990. Some of these models of short-rate behavior are tested in K. C. Chan, G. A. Karolyi, F. A. Longstaff and A. B. Sanders, "Alternative Models of the Term Structure: An Empirical Comparison" (Working paper, October 1990).