IMPLEMENTING A TRINOMIAL CONVERTIBLE BOND PRICING MODEL

BRG
Berkeley Research Group

WHITE PAPER | SEPTEMBER 2015

Prepared By:

Stuart McCrary
smccrary@thinkbrg.com
312.429.7902

Berkeley Research Group | 1800 M Street NW, Second Floor | Washington, DC 20036

This manuscript is program documentation for a model to create a trinomial stock price tree with mean reversion to price convertible bonds and convertible preferred stock. Although information contained in this manuscript is believed to be accurate, the documentation is offered without warranty, and users agree to assume all responsibilities and consequences from using this documentation.

Introduction

This manuscript documents a model to value convertible bonds and preferred stock[1] subject to default. The model is based significantly on a published model that relies on a trinomial stock tree. The model relaxes several constraints:

- This model does not assume that stock dividends arrive continuously. Instead, the user specifies the past known dividend paid, the ex-dividend date, and a forecast for the growth of future dividends.
- This model can rely on many more steps per year than the published model, which allowed annual steps. The user can use semiannual steps to match the tenor of most convertible bonds, or create two or more steps between standard coupon periods.
- The steps occur on periodic intervals with actual days consistent with industry-standard treatment for the timing of nodes, mindful of weekends and market holidays.
- This model assumes that the price of the stock is subject to an adjustment that occurs on periodic ex-dividend dates, mindful of weekends and holidays.
- This model allows for a short first period, permitting pricing partway through a coupon period.
- This model accepts a curve for interest-rate inputs used for valuation, a curve of volatilities used for valuation, and a curve of default intensities.
- This model reconciles the timing of steps required to price a convertible bond to the timing of dividend payments, which can affect early exercise.

This document describes particulars used to implement a tree-based model to value convertible bonds, based in part on the model described by John Hull in *Futures, Options, and Other Derivatives*.[2] The model presented below will implement a similar model, except the user has much more flexibility to input key assumptions for interest rates, default probabilities, and volatilities over time. The model presented below will rely on a tree that captures the cash-flow patterns typical of common stocks, which may affect the value of early conversion.

Overview

The model presented by Hull begins by creating a trinomial tree representing stock prices and using probabilities and stock price to create a second tree representing bond prices. The three end states include: 1) an up move in the price of the stock; 2) a down move in the price of the stock; and 3) default. The model controls the size of the moves and the probability of each of the three end states to 1) represent the chance of default, 2) match the assumed volatility, and 3) satisfy no-arbitrage conditions given a continuously compounded dividend and a continuously compounded risk-free rate. The model expects that the time interval between steps is equal for all steps and that the level of volatility remains a constant over the span of the tree. The model as presented is a useful educational tool. However, it is possible to match market inputs more accurately by allowing the user to specify the inputs in considerably more detail.

1 In general, this manuscript will describe the application of the model to convertible bonds but can generally be used to value convertible preferred stock as well.

2 Hull, John, *Futures, Options and Other Derivatives*, 7th edition, 599–602.

The model considers three values for the convertible bond. The value described by Hull as Q_1 is the value of a bond evaluated by backward induction on a trinomial tree. At a particular node, the value described as Q_2 is the value if the bond is or could be called. The value described as Q_3 is the value of the bond at that node if converted to common stock. Generally, the bond is callable at a fixed call price or at a premium that declines over time and may not be callable at all for a period. The tree used to value the mode must account for all three prices. In particular, the bond is worth the greater of Q_1 and Q_3 subject to a maximum price of Q_2. Equation 1 describes the valuation while the bond is *not* callable:

$$\text{Bond Price} = \text{Max}(Q_1, Q_3) \tag{1}$$

Equation 2 describes the valuation while the bond is callable:

$$\text{Bond Price} = \text{Max}(\text{Min}(Q_1, Q_2), Q_3) \tag{2}$$

This formulation presumes that the bondholder can be forced to convert when the value of the equity is high enough; that the bondholder can avoid the call by converting to equity; and that the bondholder cannot capture a higher value shown on the tree because that requires time that the forced conversion does not allow.

Building the Stock Tree

The tree of stock prices must be built before the tree of bond prices is constructed. The stock tree is built by forward induction, starting at the left of the tree and moving to the right one step at a time.

Consider first the shape of the trinomial tree. In its most general form, the tree described above contains nodes that branch to three separate nodes, as illustrated in Figure 1. The gray nodes are defaults, which terminate and do not branch further. The up node (higher stock price) and down node (lower stock price) each branch to additional nodes. Within each step, each vertical cross-section, the unconditional probabilities[3] must sum to 1 or 100%. At a given node, the conditional probabilities[4] also sum to 1. Finally, the sum of all of the individual probabilities to default on a particular step must match the default rate predicted by the assumed default intensity.

FIGURE 1. COMPLETE NON-RECOMBINING TREE

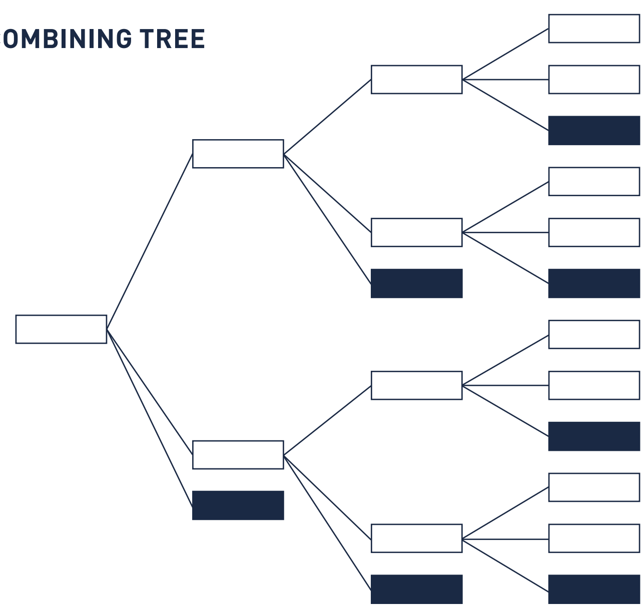

3 The unconditional probability of getting to a node is the probability of getting to that node from the starting point at the leftmost node of the tree and equals the product of all the marginal probabilities of going down a path necessary to get to that node. Note that, as drawn in Figure 1, a single path leads to each node.

4 $1=p_u+p_d+p_x$

Hull illustrates the tree with two boxes at each node, reflecting the value of the bond and the value of the underlying common stock. The extra boxes are not illustrated here or in the figures below. Instead, the model carries two trees as illustrated: one for the stock and one for the convertible bond or preferred stock. Conditions that apply to the stock tree also apply to the bond tree. The assumed volatility and market interest rates define the stock prices and probabilities at each node, which are then applied to the bond tree so that the value of the bond is derived from the information about the stock.

A first step toward simplifying the tree in Figure 1 is to realize that default is a single state at each time step, as illustrated in Figure 2. Default, however, must be considered a possibility from any node.

FIGURE 2. REDUCED NON-RECOMBINING TREE

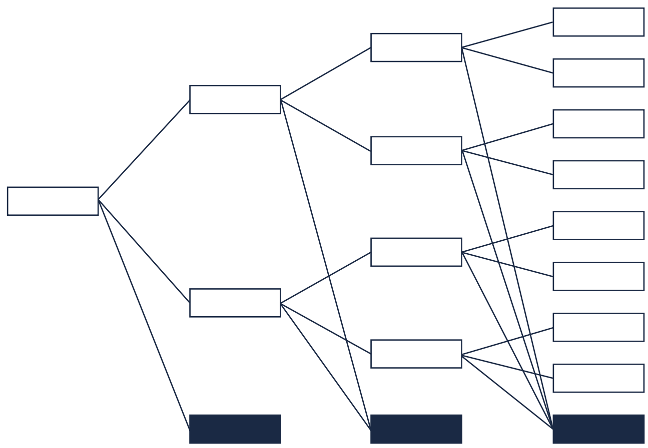

This non-recombining tree does not reduce the complexity of the problem much so it does not significantly reduce the total number of nodes to be evaluated. The number of nodes grows rapidly as the number of steps increases, so it is practically impossible to evaluate a tree as shown in Figure 2 that has more than about 24 steps (that is, more than about 24 columns).

Figure 3 represents the same process as the trees in Figures 1 and 2, but the tree has been forced to recombine. This occurs naturally when the volatility is constant because the node coming down from above exactly equals the node coming up from below. In particular, in Figure 2, the middle two unshaded nodes in the column with four unshaded nodes are actually the same price, and the nodes emanating to the right of one of these nodes also match up with the nodes emanating from the other duplicate node.

FIGURE 3. REDUCED RECOMBINING TREE

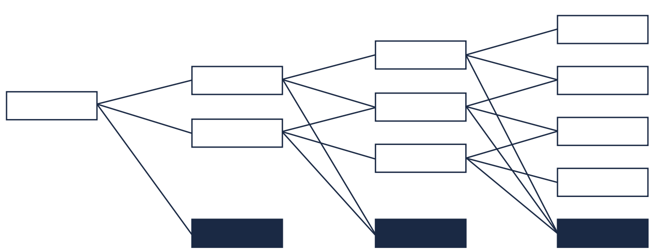

Notice that, as with Figure 2, only one default node is illustrated for each step (each column on the tree), although the probability of default during a particular step must reflect the continuing probability of surviving long enough to be able to default. The probability of ending at a particular defaulting node must be consistent with the assumed default intensity and provides the basis for allocating probabilities to the other nodes in the step.

If volatilities are unequal or if the time step size is unequal[5], the down node coming from above will generally not equal the up node coming from below. Figure 4[6] illustrates the problem. The increasing volatility during the second step causes the blue down node to be below the orange up node.

FIGURE 4. CHANGING DISPERSION

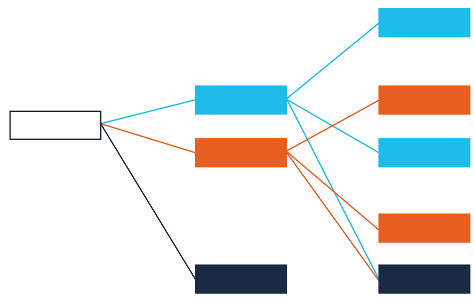

It might appear that the four non-defaulting outcomes in Figure 4 have a higher volatility than the three non-defaulting outcomes in Figure 3. In fact, the volatility from the pricing date to each step matches the inputted or interpolated volatilities. The dispersion of the outcomes at each step should be consistent with the level of volatility on a tree built from flat volatilities.

$$\sigma_{T_1-T_2} = \sqrt{\frac{dT_1 * \sigma_1 + dT_2 * \sigma_2}{dT_1 + dT_2}} \tag{3}$$

More generally:

$$\sigma_{T_1 \text{ to } T_N} = \sqrt{\frac{\sum_{1=1}^{N} dTi * \sigma_i}{T_n - T_1}} \tag{4}$$

In other words, it is easy to build a tree that has three nodes instead of four in this column by creating a tree using the weighted average volatility over the two steps. The resulting recombining distribution can have the same expected value and volatility as the four node non-recombining column in Figure 4.

Figure 5 illustrates this strategy to keep the recombining shape intact. The recombining tree is calibrated to a blended volatility, so that the spacing of the three non-defaulting nodes is consistent with the volatility observed in Figure 4. Upon inspection, individual nodes are not arbitrage-free. The leftmost green node expects to observe either a large gain or a small loss, while the leftmost red node expects to observe either a small gain or a large loss. It would be possible to find probabilities that make the local nodes arbitrage-free. However, levels and probabilities in this model are set so that the tree is arbitrage-free from the pricing date, the leftmost node. In particular, the stock prices in the nodes of a column are chosen so that their average matches the forward price of the stock and that their spacing is consistent with the dispersion expected given the time elapsed.

5 More generally, the problem arises if the de-annualized volatility over a step does not match the de-annualized volatility of the previous step.

6 Figure 4 is not drawn to scale. The way the final nodes are dispersed can vary.

FIGURE 5. CHANGING DISPERSION

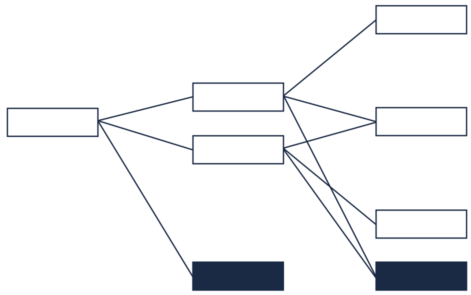

Figure 6 shows a numerical example of a non-recombining tree representing a half-year short first step followed by a full-year step. For simplicity, the probability of default is assumed to be 0% and is omitted from the figure. The de-annualized volatility (dσ) for the first step is 7.07%, based upon an annualized volatility of 10% for a half year ($d\sigma = \sigma * dT^{.5}$). The volatility driving the second step is 20%, the volatility prevailing at that point and described as a marginal volatility. The price given an up movement is given by Equation 5, and the price given a down movement is given by Equation 6. Because both the time interval and the volatilities are different, the nodes do not recombine, as can be seen on Figure 6.

$$\text{Price}_{\text{Up},t+1} = \text{Price}_t * e^{\text{Drift}+d\sigma} \tag{5}$$

$$\text{Price}_{\text{Down},t+1} = \text{Price}_t * e^{\text{Drift}-d\sigma} \tag{6}$$

The drift is found by searching for a return that sets the weighted average of prices at this first step equal to the forward price.

The percentage to the left of each node is the unconditional probability of getting to that node. The returns to the right of the node are the continuously compounded return from the drifted starting price. The data in Figure 6 confirm that the volatility of the returns corresponds with the expected or weighted average volatility.

FIGURE 6. NON-RECOMBINING EXAMPLE

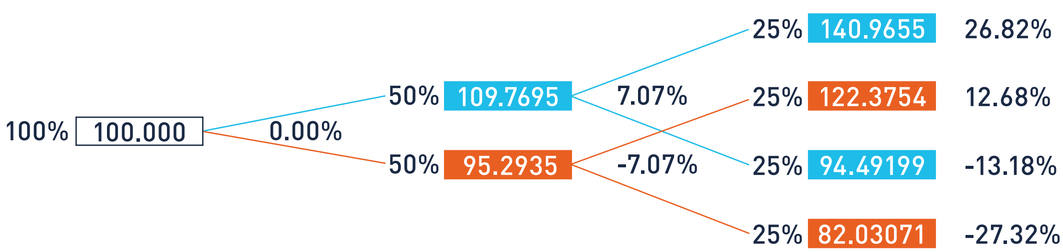

dT		0.50		1.00	
Forward Rate		5.00%		7.00%	
Marginal Volatility		10.00%		20.00%	
dVolatility		7.07%		20.00%	
Drift		2.25%		5.01%	
Forward Price	100.000	102.532		109.966	
Average		102.532	0.00%	109.966	-0.25%
Error		0.000		0.000	
Tree Volatility		7.07%		21.21%	
Expected Volatility		7.07%		21.21%	17.32%
Drifted Price	100.000	102.276		107.8027	

Figure 7 constructs an equally spaced tree with constant volatility equal to the volatility measured on the tree in Figure 6. This tree is constructed to discover the column of prices in this recombining tree that match the volatility of the non-recombining tree on the second step and also produce a weighted average price that matches both the non-recombining tree and the forward price. It is possible to derive the rightmost step without constructing the entire tree in Figure 7, which is what this model does, continuing from left to right to complete the tree.

After just two steps, the outcomes possible on the non-recombining tree and the recombining tree are quite different. However, the distributions of prices on both trees converge to the lognormal distribution with the same volatility, just as the original Cox-Ross-Rubenstein trees converged on the Black-Scholes assumptions. It is perhaps true that the recombining tree converges faster than a non-recombining tree, since a histogram of returns on this tree already looks more like a normal distribution than the returns on the non-recombining tree. In any case, substituting a binomial approximation for a continuous distribution is only possible if the number of steps is sufficiently large, in which case the substitute tree provides a robust description of possible stock prices sufficient to price a convertible bond.

FIGURE 7. SIMILAR RECOMBINING STEP

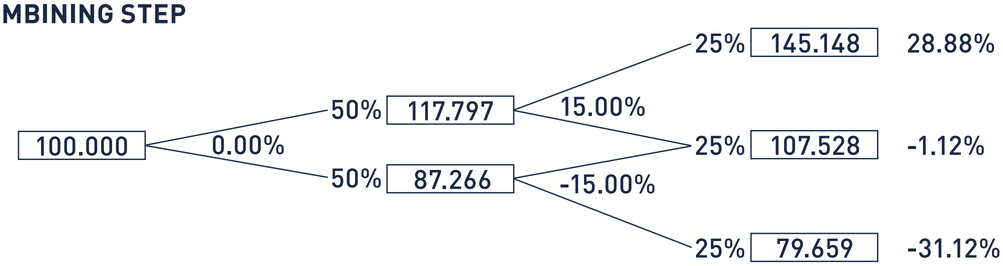

dT		0.75		0.75	
Forward Rate		5.00%		7.00%	
Marginal Volatility		17.32%		17.32%	
dVolatility		15.00%		15.00%	
Drift		1.38%		5.88%	
Forward Price	100.000	102.532		109.966	
Average		102.532	0.00%	109.966	-1.12%
Error		0.000		0.000	
Tree Volatility		15.00%		21.21%	
Expected Volatility		15.00%		21.21%	17.32%
Drifted Price	100.000	101.389		108.740	
K				0.741	

Finally, notice that the spacing on the recombining tree follows the pattern observed by Jamshidian.[7] That is, the ratio, K, of the middle price divided by the top price (107.528 / 145.148 = 0.741) equals the ratio of the bottom price divided by the middle price (79.659 / 107.528 = 0.741) and is constant for all prices on a step. Therefore, the strategy for completing a step is to find K, which is determined by the volatility input. This method requires the unconditional probability for each node, mindful of the probability of default. After determining for forward price of the stock, it is necessary to iterate on stock prices (Jamshidian uses the top node and defines all other prices in this step with K relative to that top price) until the weighted average price matches the forward price.

$$K = e^{-2*\bar{\sigma}*\sqrt{dT}} \tag{7}$$

The volatility to use in Equation 7 is the quoted volatility for a stock option expiring at the same time as the step being analyzed.

7 Jamshidian, Farshid, "Forward Induction and Construction of Yield Curve Diffusion Models," *Journal of Fixed Income* (June 1991), 62–74.

Risk-Free Rates and Forward Rates

This model accepts a list of interest rates associated with different times to maturity. In particular, the user passes the spot rates or zero-coupon bond rates associated with particular maturities. These spot rates can be directly observed from U.S. Treasury Strips or derived or "bootstrapped" from a series of Treasury bonds or interest rate swaps.[8] For example, Table 1 shows the spot rates derived from the USD Libor Swap Curve as of June 15, 2015:

Table 1: Bootstrapped Spot Rates as of 06/15/2015	
1 Month	0.19%
3 Months	0.30%
6 Months	0.44%
1 Year	0.65%
2 Years	1.06%
3 Years	1.42%
4 Years	1.68%
5 Years	1.92%
10 Years	2.53%
30 Years	3.02%

The rates and associated years to maturity provide the basis for interpolation so that, it is not necessary to input yields for all time horizons addressed by the model. The model can fill in missing inputs using linear interpolation or a cubic spline. In either case, spot rates for two adjacent maturities provide the information needed to calculate a forward rate between the two maturities. For example, the interpolated rate[9] for 4.25 years is 1.736%, and the interpolated rate for 4.5 years is 1.797%. The present value for 4.25 years is 0.929, and the present value for 4.50 years is 0.922. The forward price or present value for the 0.25-year interval spanning the two maturities is 0.993, which implies a forward rate of 2.707%. The calculations appear below:

$$PV_{4\ 25} = e^{-1\ 736\%*4\ 25} = 0\ 929 \tag{8}$$

$$PV_{4.50} = e^{-1.797\%*4.50} = 0.922 \tag{9}$$

$$PV_{4.25-4.50} = \frac{0.922}{0.929} = 0.993 \tag{10}$$

$$Rate_{4.25-4.50} = \frac{LN(0.993)}{4.50 - 4.25} \tag{11}$$

8 Because the bootstrapping occurs outside the model described here, it is outside the scope of this document. See, for example, Vasicek, Oldrich A., and H. Gifford Fong, "Term Structure Modeling Using Exponential Splines," *Journal of Finance* 37 (1982), 339–56.

9 The interpolated rates were very similar using either linear interpolation or cubic splines.

The model uses these forward rates to calculate the arbitrage-free level of the stock at different points in time. For June 15, 2015, forward rates are above the spot, as they generally are with an upward-sloping yield curve. These forward rates produce a higher forward stock price, all things equal, than if the spot rates were used (Figure 8).

FIGURE 8. SPOT AND FORWARD RATES - 06/15/2015

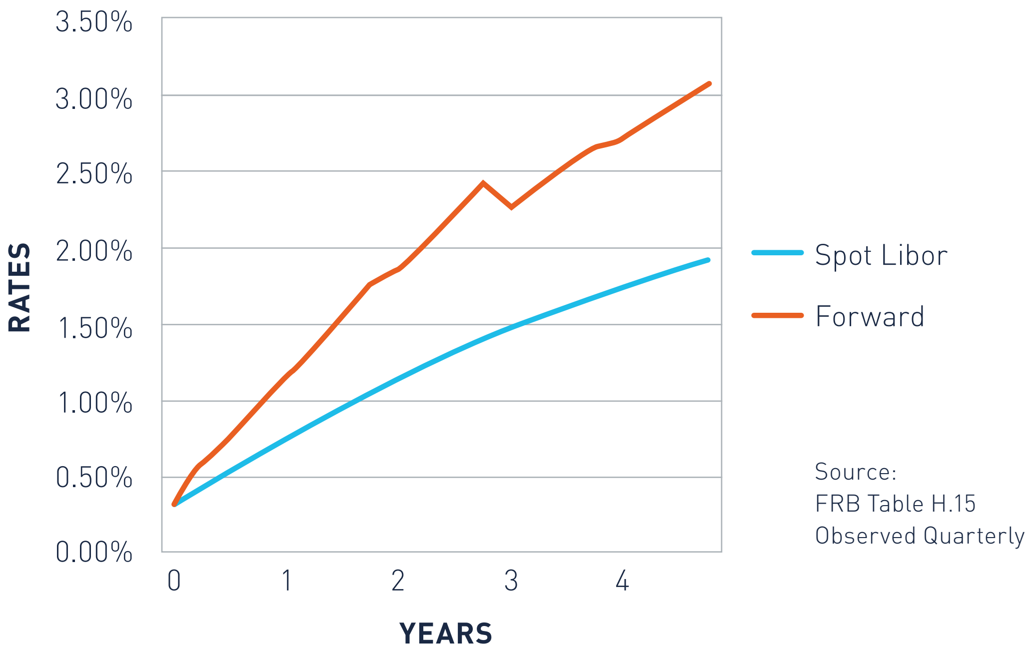

Stock Return Volatilities

The model also allows for volatilities to be different over time. Table 2 shows some hypothetical volatility levels for a stock over different forward-time horizons.

Table 2: Stock Volatilities as of 06/15/2015	
1 Year	14.0%
2 Years	15.0%
3 Years	15.5%
5 Years	16.0%
10 Years	17.0%
30 Years	18.0%

Note that the number of volatility inputs and the years observed do not have to match the tenor of the spot rates provided. The model interpolates volatilities for periods between inputs on the table. Potentially every node on the stock tree is subject to a unique volatility input. The model expects to receive constant volatilities for each time horizon. That is, with reference to the inputs on Table 2, the model assumes that a stock will experience 16% for the entire five years. The pattern shown would require not just an increase in volatility apparent from the table. It would also require higher still marginal volatilities so that the average volatility will match the table.

Table 2 represents the volatilities for expiration lengths. Volatilities for horizons not listed on Table 2 are approximated by linear interpolation or cubic spline interpolation. For nodes spanning a long time or for situations in which forward volatility changes rapidly over the time horizon, it would be more accurate to interpolate the volatility at more than one point in each period and blend the assumed volatilities, using Equation 3 or Equation 4. The model as implemented averages the beginning and ending volatility, equally weighted, for each step.

Default Probabilities

The intensity of default, which Hull calls lambda (λ), is input much like the spot rates and volatilities. Again, the user can create a table of any desired length and can include as many entries as necessary to describe the chance of default over time. Table 3 shows a representative table of inputs, reflecting the continuously compounded forward[10] default intensity.

Table 3: Default Intensities as of 06/15/2015	
1 Year	1.00%
2 Years	1.25%
3 Years	1.50%
5 Years	1.75%
10 Years	2.00%
30 Years	2.00%

Hull points out that the chance of default over the interval of a step, which I will call p_x, is:

$$p_x = 1 - e^{\lambda \Delta t} \tag{12}$$

Table 3 represents the default intensities at certain points in time. Default intensities for times between points listed on Table 3 are approximated by linear interpolation or cubic spline interpolation. For nodes spanning a long time or for situations in which default intensity changes rapidly over maturity, it would be more accurate to interpolate the default intensity at more than one point in each period and blend the assumed default intensities. The model as implemented averages the beginning and ending default intensity, equally weighted, for each step.

Probability of Up and Down Moves

Hull proffers formulas for the probability of an up-move p_u and the probability of a down-move p_d. Hull's probabilities reflect the conditions described above (matching the single, constant volatility, allowing for default at a single, constant intensity, consistent with no-arbitrage forward pricing against a single risk-free rate, and allowing for a fixed dividend paid continuously). The way this model handles the no-arbitrage conditions is necessarily different and will be explained below. For now, realize that this model attaches equal weight to the up and down legs. So p_u and p_d must sum to the probability that there is no default. Therefore, in the first (leftmost) node:

$$p_u = p_d = \frac{100\% - p_x}{2} \tag{13}$$

10 That is, lambda is the chance of default going forward assuming you arrive at the step without experiencing a default. In other words, lambda is the conditional probability of default, not the unconditional probability.

Forward Pricing of Dividend-Paying Stock

The Hull trinomial convertible model and many option-pricing models assume that the stock pays a continuously compounded dividend. In fact, dividends do not accrue. Instead, the company declares a discrete dividend payable to holders of record as of the ex-dividend date payable on the dividend payment date. Holders of the stock outside of the ex-dividend date have no claim on part of the next dividend. Likewise, when a dividend is declared, the full amount of the dividend is paid to the investor holding the stock on the ex-dividend date, even if the holder just bought the stock.

Forward pricing is not a prediction of the future price of the stock. Rather, it is the price in the future that is economically equivalent to the present price in light of the time value of money and dividends received. In the absence of dividends, the forward price of the stock rises at the risk-free rate:

$$\text{Stock Price}_T = \text{Stock Price}_0 * e^{R*T} \tag{14}$$

where R refers to a single risk-free rate and T refers to the interval of time forward. In practice, the model considers a term structure of rates, so it predicts a forward price that may rise at different rates at different points in the future.

If the stock pays a dividend, the forward price of the stock immediately drops on the ex-dividend date by the present value of the dividend (i.e., the dividend discounted from the dividend payment date to the ex-dividend date), using the appropriate risk-free forward rate for that interval. The model accepts a dollar dividend amount plus an annual growth rate and assumes the dividend recurs quarterly.

Figure 9 is an illustration of the forward price of a stock. The stock begins at 100 and pays a $2 dividend annually, growing at 3% per year. In the early stages, the initial 2% dividend yield is below the risk-free rate. For this reason, the forward price of the stock generally declines. Later, the dividend yield is below the forward rate, so the stock price generally increases. All through, the stock price moves down discretely on the dividend payment date. This pattern of stock price can affect the value of a convertible bond and can affect the timing required to maximize value in conversion.

FIGURE 9. STOCK FORWARD PRICE

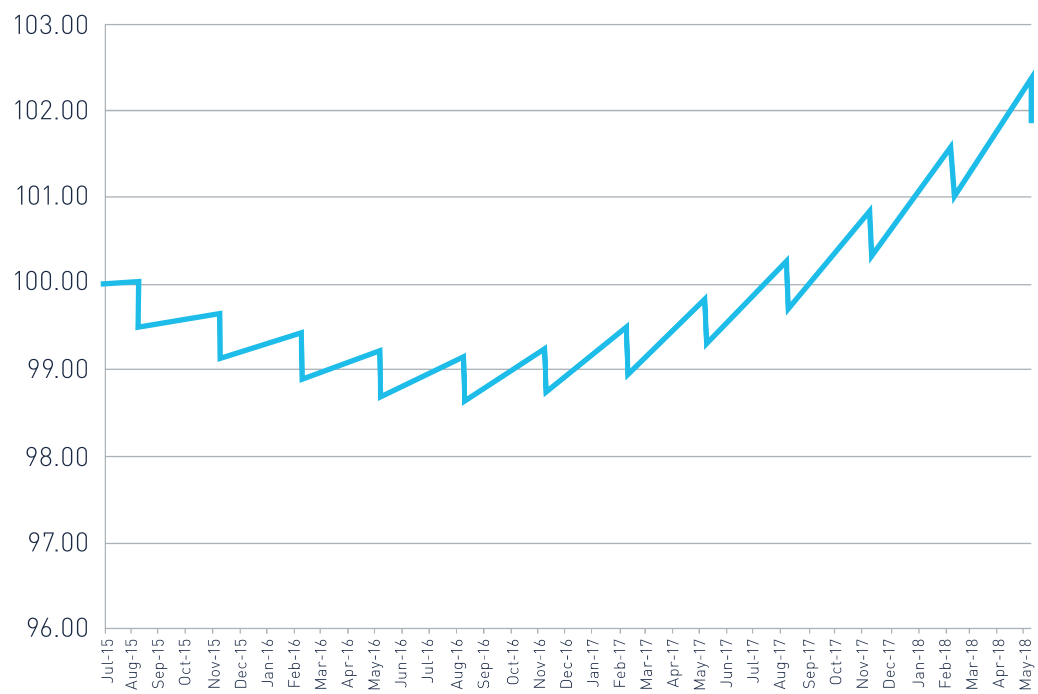

The Convertible Bond Tree

The final step pulls all of the above information together to value the bond. Work from right to left. On the last step (maturity), the bond is worth the greater of the face value of the bond (e.g., $100) or the value of the bond converted into common stock (Q3). In prior periods, if the bond is callable, the price of the bond is worth the lesser of the call price (Q2) and the discounted expected value of the bond tree prices (Q1). If the bond is not callable, it is worth the discounted expected value of the bond tree prices (Q1). Whether or not the bond is callable, it is additionally worth the greater of the value of the bond converted into common stock (Q3) and the capped expected value.

The value of the convertible bond is the value at the left-most node of the bond price tree. It is also possible to estimate the sensitivity of the bond price to changes in the stock price. This sensitivity, similar to the delta of an option, can be estimated by looking at the first step to the right of the left-most mode. The numerator is the change in bond price observed on these two nodes on the bond tree. The denominator is the change in stock price observed on these two nodes of the stock tree.

Conclusion

The trinomial model presented is an extension of the binomial stock option pricing strategy, but with a third step added to provide for default. The resulting stock price tree can be reduced by collecting together the default nodes of each step. The model as presented allows volatilities to be different at every step. This creates a non-recombining stock price tree. An equivalent recombining tree is created to substitute for the non-recombining tree. This tree provides the information necessary to construct a convertible bond pricing tree.

This model relaxes a number of restrictions, allowing for more detailed market inputs. The tree adapts to the periodic timing of bond payments and accounts for the possibly mismatched timing of stock dividends. It allows for discrete dividends that grow over time. The model accepts term structures of interest rates, of volatilities, and of default intensities. Finally, the model can subdivide these calendar-driven periods into as many sub-periods as desired, where price discovery occurs but coupon payments are not necessarily made.

References

Bjerksund, Petter, and Gunner Stensland, "Implementation of the Black-Derman-Toy Model," *Journal of Fixed Income* 6:2 (September 1996), 67–75.

Hull, John, *Futures, Options and Other Derivatives*, 7th edition (2008), 599–602.

Jamshidian, Farshid, "Forward Induction and Construction of Yield Curve Diffusion Models," *Journal of Fixed Income* (1991), 62–74.

Nelson, Charles, and A. Siegal, "Parsimonious Modeling of Yield Curves," *Journal of Business* 6 (1987), 473–489.

Rendleman, Richard J., "Interpolating the Term Structure from Par Yield and Swap Curves," *Journal of Fixed Income* (2004), 80–89.

Vasicek, Oldrich A., and H. Gifford Fong, "Term Structure Modeling Using Exponential Splines," *Journal of Finance* 37 (1982), 339–56.

About the Author

Stuart McCrary

Stuart McCrary is a trader and portfolio manager who specializes in traditional and alternative investments, quantitative valuation, risk management, and financial software. Before joining BRG, he spent 13 years consulting on a wide range of capital markets issues including litigation consulting, valuation, modeling, and risk management. Previously, he was president of Frontier Asset Management, a market-neutral hedge fund. He held positions with Fenchurch Capital Management as senior options trader and CS First Boston as vice president and market maker, where he traded OTC options and mortgage-backed securities. Prior to that, he was a vice president with the Securities Groups and a portfolio manager with Comerica Bank.

Mr. McCrary is the author of several books, including *How to Create and Manage a Hedge Fund: A Professional's Guide* (John Wiley & Sons, 2002).

The author thanks Sanjay Nawalkha for reviewing this manuscript and offering comments. Dr. Nawalkha is a professor and chairman of the Department of Finance, Isenberg School of Management, at the University of Massachusetts.

About Berkeley Research Group

Berkeley Research Group, LLC (www.thinkbrg.com) is a leading global strategic advisory and expert consulting firm that provides independent advice, data analytics, authoritative studies, expert testimony, investigations, and regulatory and dispute consulting to Fortune 500 corporations, financial institutions, government agencies, major law firms, and regulatory bodies around the world. BRG experts and consultants combine intellectual rigor with practical, real-world experience and an in-depth understanding of industries and markets. Their expertise spans economics and finance, data analytics and statistics, and public policy in many of the major sectors of our economy, including healthcare, banking, information technology, energy, construction, and real estate. BRG is headquartered in Emeryville, California, with offices across the United States and in Australia, Canada, Hong Kong, Latin America, and the United Kingdom.