Convertible Bond Valuation: 20 Out Of 30 Day Soft-call

Robert L. Navin
Quantitative Analysis and Quantitative Trading Strategy

Highbridge Capital Management, LLC
767 Fifth Avenue, 23rd floor
New York, NY 10153

Abstract

"Soft-call" in convertible bonds (CBs) usually means that the bond can be recalled by the issuer only if the stock price has previously closed above a specified trigger price for any 20 out of any 30 consecutive trading days. It is not an easy optionality to value and to my knowledge no method has been implemented besides Monte Carlo. The problem is not very well suited to Monte Carlo due to a large number of possible permutations of stock price closes above or below the trigger over a year period (i.e., 2^{260}) with the result that a Monte Carlo valuation requires a trade off between being slow and not smooth. The soft-call feature is typically modeled in the CB industry by presuming that the bond is called as soon as stock touches the trigger price. After discussion of the exact solution of this problem (requiring valuation of component derivatives on, of order, 2^{260} grids), a simple algorithm is presented to approximately value this feature for the general n out of m case of soft-call. The algorithm requires merely a subtle change to the call feature of the one-touch model and only one running of a grid or tree and hence it is very fast. The method boils down to making the bond "1-touch" callable on some days and not on others, the precise sequence being a function of the 29 day stock price close history. It gives smooth functional output (besides theta of course, the theoretical price jumps from day to day) and very compelling qualitative results. The results are accurate to a dime on the dollar of benefits due to provisional call, and this is determined by comparison to the exact solution for easily calculated cases.

Introduction

A convertible bond is a coupon paying bond or (possibly putable) zero-coupon bond with the embedded option to turn it into a fixed number of shares. It is clearly a hybrid interest rate and equity derivative, and thus more complicated than either. The convertible bond market is increasing its significance as a method for medium to better quality companies to raise capital. It is easy to argue then, that convertible bonds are one of the most actively traded and complex derivatives in the market place.

Convertible bond issues typically include a clause allowing the issuer to call the bond back from the holder, by paying a cash sum. This ensures the issuer can refinance if it is in their interest. The notice period that precedes this cash payment (often thirty calendar days) allows the holder the option to convert into a fixed number of shares instead of receiving cash. Thus a convertible bond may trade at a price above the cash call amount, even while currently callable. Clearly as soon as it is called, it will be worth a thirty day option to get cash or convert to stock.

However since the 1960's, after some issues were called even before the buyer had received one coupon, a call protection feature was introduced to ensure the bond would be outstanding for at least a year or two. This feature took two forms, and these are often combined. The first is hard-call protection or *hard-call*. The bond is not callable for the first one, two, or maybe three years (in some cases more) after issue. The second case, is soft-call protection or *soft-call*. The bond is callable only if the stock trades above a trigger price for 20 out of any consecutive 30 trading days. As soon as this condition is observed, the thirty day notice of cash call may be given leaving the holder with a thirty calendar day American option to convert to stock or take the offered cash (specified in the prospectus) at the end of the notice period. Other variations on this theme are practiced, such as 20 out of 20 provisional call, but the 20 out of 30 is the most typical.

A first attempt to value almost all of the features in convertible bonds, using a (say) one-factor model assuming stochastic stock prices (while interest rates, bond yields and stock borrow costs and dividend yields are assumed fixed forever), is generally a trivial extension of the Black-Scholes 'method' that results in the famous formula. A numerical method, such as a grid or tree implementation will have to be used. The only part of this recipe that is lacking is valuation of the provisional call feature.

I outline below a method to value soft-call numerically approximately, but qualitatively compellingly, on a single running of a grid, with a prescription to obtain arbitrary accuracy with multiple runs of the grid with different boundary conditions. It may be viewed as a perturbation expansion of the answer with the initial 'zeroth' order value requiring only one grid run.

Defining the Problem

A one-factor convertible-bond model, without provisional call is easy to construct and I shall not describe it here. See, for example, the article in the July 1997 issue of Bloomberg magazine (*"A Mathematical Edge for Convertible Bond Traders"* Bloomberg magazine, July 1997).

Let us simply note the following features of such a model of convertible bonds. At any time slice in the grid: the holder has the choice to either hold the bond, put the bond at a specified price, or convert into a specified number (the conversion-ratio) of shares; the issuer has the option to issue 30 days notice of cash call. If called, the holder retains an option to either hold for thirty days and receive cash or convert at any time. Thus, a rational exercise policy makes the valuation unique for some continuous Markovian stock price distribution, using the risk neutral pricing scenario.

Below, I will use such a model with slightly changed 'specs.' I describe two instruments and solve them perturbatively. One is the case of an option which knocks-out (immediately expires worthless) if the stock price is above the trigger on 2 out of any 5 trading days and pays $1 if it survives to expiration and the other is a convertible bond which typically has one year of hard-call, one year of soft-call (20 out of 30 day trigger say) and then for the remaining life it is callable.

Triggered-Knockout Option

1. Setting up the problem

Consider an option that instantly expires worthless ("knocks out" or is "called away" worthless) when the following three conditions are satisfied: the stock closes above a trigger price of $80 say; during the previous 4 business days it closed above the trigger on at least (any) 2 days; and the last close above the trigger occurred within a 7 day period from expiration. Otherwise the option pays $1. This is the 3 out of 5 provisional trigger problem, with a 7 day window for exercise. I shall discuss valuation two days before the callable period begins, and so, if valuation is on day 1, it is callable from day 3 to day 9 inclusively and it expires on day 9 paying $1 if it has not been called.

Convertible bond practitioners will recognize the salient features of a 20 out of 30 day soft-call with a 1 year (260 business days etc.), 2 year or 3 year window for exercise, the worthless knock-out value corresponds to the 30 day cash or conversion option and the $1 payout resulting from 'failure' to trigger exercise corresponds to ending up holding a, typically callable, convertible bond.

The full solution to this problem is straight forward. Assume we know the relevant risk-neutral distribution Green's function and therefore evolution operator (non-mathematicians should simply think in terms of an "evolution operator:" it is the repeated application of the one-day "grid-algorithm" operator, which takes *any* payout function as an input and generates the price of this payout as a stock slide on the previous day). For this problem there are eight relevant possible 'payouts:' One payout of $1 at expiration and seven payouts of zero on each day from the third day after valuation. Each payout occurs under conditional probability measures, e.g. the path of stock resulting in a history of closes that satisfies the trigger will pay out $0 on the day it is satisfied, all other paths, never satisfying the trigger will pay $1 at expiration. We must value all of the payouts under these conditional probability measures. There are $2^7 = 128$ possible permutations of stock closes above or below on each day of the options life.

Two payouts are possible at expiration, a payout of $1 and a payout of zero for all stock prices, as in fig. 1. We then backward-evolve one day and chop up the price at the trigger resulting in two pieces, as shown in fig. 1 for the not called payout. This gives us the price slide on day 8 of getting $1 at maturity dependent on stock closing above the trigger and the price slide for stock closing below the trigger. We repeat the procedure until we get to the valuation date and have 128 different price slides. On all the days from day 8 to day 3 inclusive, a payout of $0 could be made if the knock-out is triggered. These payouts require the process to be applied to them resulting in a further valuation of $2^7 + 2^6 \ldots + 2^1 = 254$ grids. Obviously valuing the zero payouts is trivial but bear in mind these, in general, may be non-zero and will be non-zero in the case of convertible bonds and so we continue to count them all.

In this example we shall also assume that stock never closed above the trigger prior to the valuation date.

To recap, Fig. 1 shows how the not called contribution of a $1 payout is propagated backward by repeated application of one-day backward-diffusion alternating with splitting under the conditional probability of stock being above or below the trigger. By splitting, convolution of the price with step functions struck at the trigger is intended. Note that the totality of results (just for the not called case) will all add up to the present value of $1 if summed.

A collection of valuations of each of the seven possible payouts is obtained, each payout is resolved into the constituent prices under conditional probability measures of being above or below the trigger at each close. A d day period of soft-call, with an n out of m trigger, results in a total number of price slides of

$$2^{d+m-1} + \sum_{j=1}^{d+m-1} 2^j$$

The problem is now one of elimination. Many of these are mutually exclusive. For example if stock closes above on day 1, 2 and 3, then the option knocks out on day 3. Thus, all prices that have these three days with stock above the close and are called on day 4 to 9 or are not called should be thrown away.

Note that for a 2 year period we have a number of price slides of order 2^{520} to be valued, and we immediately see that Monte-Carlo methods will need to be seriously modified to succeed in valuing such things.

2. Combinatorics

To solve the problem completely, the idea is simply to select the prices under the various conditional probabilities that contribute to the option and discard all the others. I shall label the various contributions as in the following example:-

$$P_e(6;110010;S)$$

means the option price for stock price, S (or option prices over a stock price slide) that represents exercise on the 6[th] day under the conditional probability that the history of stock price closes was: on the day of call stock closed above (1), the day before call, stock closed above (1), the day before that, below (0), and so on: below (0), above (1), below (0). There are also price slides for non-exercise, or hold, P_h, for this problem these are the only ones that are non-trivial but we count them all for later use. I drop the S argument for brevity and in the tables below I drop the $P_e(\ )$ formalism and the number of days, while the column headings show whether it is a called or not-called final payout that is being priced. The complete list of all possible price slides under the conditional probability measure is:-

$$\begin{aligned}
&P_e(3;000)\\
&P_e(3;001)\\
&\vdots\\
&P_e(3;111)\\
&P_e(4;0000)\\
&P_e(4;0001)\\
&\vdots\\
&P_e(4;1111)\\
&P_e(5;00000)\\
&P_e(5;00001)\\
&\vdots\\
&P_e(9;000000000)\\
&\vdots\\
&P_e(9;111111111)\\
&P_h(9;000000000)\\
&\vdots\\
&P_h(9;111111111)
\end{aligned}$$

Fig. 1

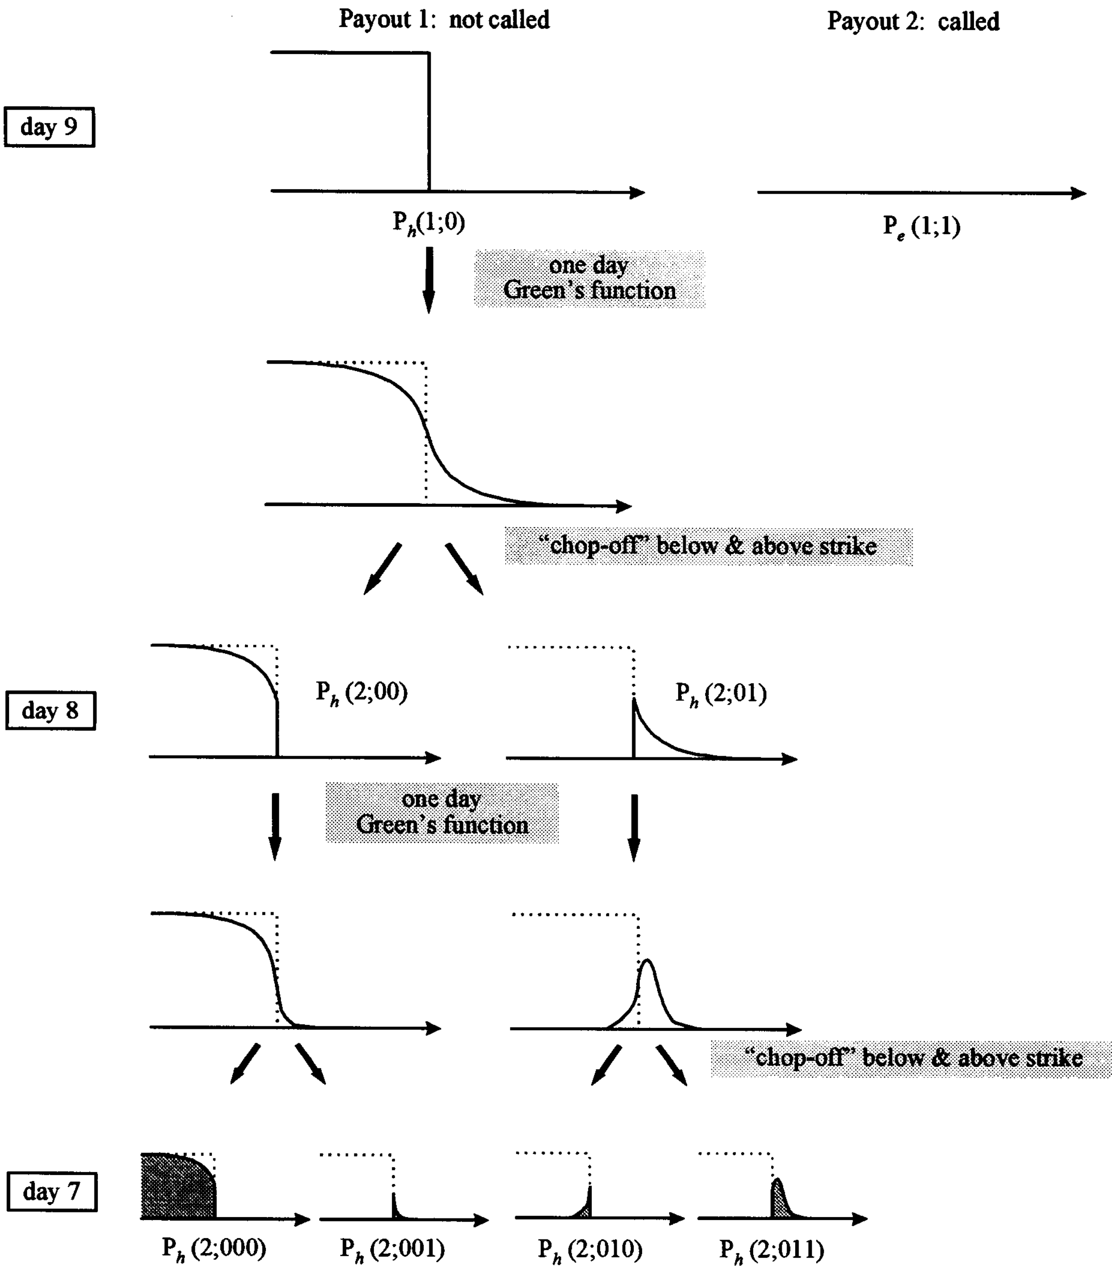

Repeated application of evolution operator and splitting under conditional probability measure of being above and below strike on close. For the $1 payout at expiration.

This problem is now solved most simply by going forward in time, reflecting the non-Markovian nature of the solution, although it is only non-Markovian in so far as the instrument that is held on any day during the soft-call period has specs which depend on the lessor of *m-1* days of close history, and the number of days since the date of call start minus *m-1* days.

We now build up all the permutations by splitting each price into the part conditional upon the stock closing above or below (essentially counting in base two), and identifying exercise when it occurs and putting the relevant previously calculated value in. For the first day's close we have: close above, hold, and close below, hold, because the trigger did not cause exercise. We then split each hold again, until we see exercise and then set the value equal to the price of exercise on that day for that permutation of closes above or below. When exercise is observed the branch stops 'growing.'

Fig. 2 counts all of the contributions of splitting the forward propagating price distribution under the conditional probability measures, going forward three days. It shows that one of the permutations (perms) of close price history to the first callable date (day 3) results in exercise. We know the value of this and then the remaining probabilities carry on being divided, resulting in three more exercises on day four as shown in fig. 3.

Continuing to maturity we will be left with the perms that never result in call. These are all valued selecting from the collection of not called price slides.

This method simply ensures no double counting of mutually exclusive possibilities. The final step is merely to sum up the contributions to the price. The remaining perms of conditional probability measures under which the pay outs could have been valued are all irrelevant.

3. Perturbation Expansion.

Thinking of any given contribution, say $P_e(6;110010)$, we see that the more times it is cut up and then down the smaller the value will be. The values may be ranked:
largest:

$$P_e(6;111111)$$
$$P_e(6;000000)$$

this is zeroth order and then the next largest terms are first order

$$P_e(6;100000)$$
$$P_e(6;000001)$$
$$P_e(6;111110)$$
$$P_e(6;011111)$$

Fig. 2

Finding relevant price distributions: "How to cut-up call on various dates." Proceeding forward in time (right to left), we want the perms of price closes above or below trigger. When exercise is observed, i.e., the soft-call conditions are satisfied, the branch stops dividing and the payout is valued on valuation date under the relevant conditional probability measure.

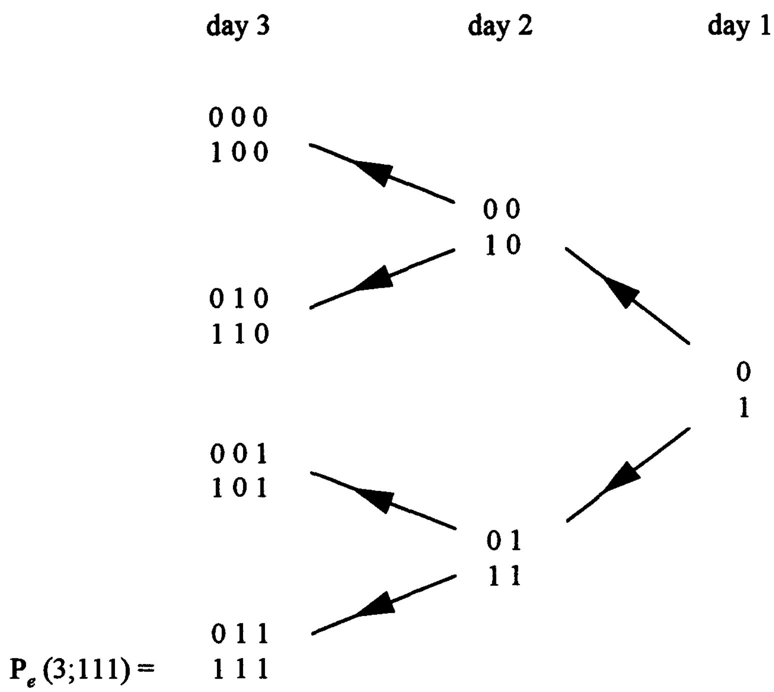

The 111 branch ends: in call on third day, assuming history prior to valuation date (day 1) was all closes below trigger.

Fig. 3

By day 4 we get 3 more calls.

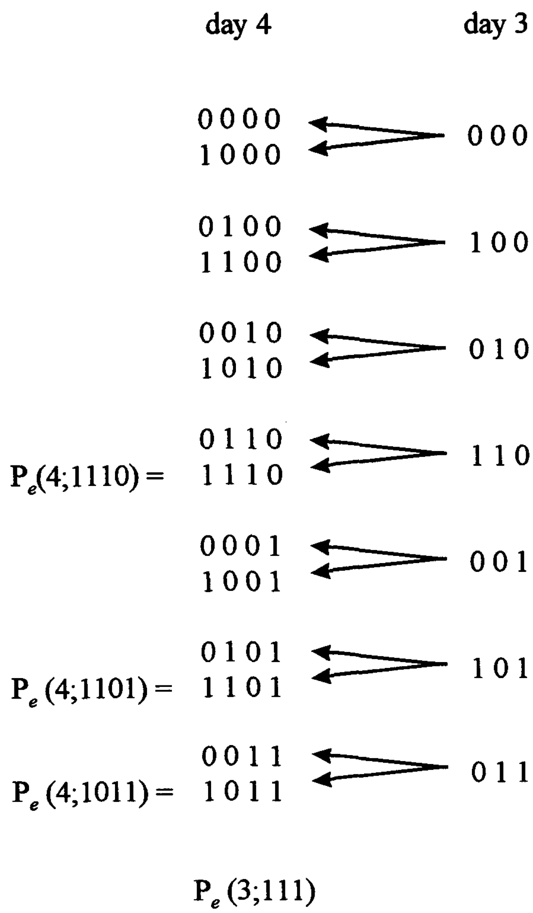

These three branches stop dividing also and the remaining 11 branches continue to be split.

and second order;

$$P_e(6;010000)$$
$$P_e(6;001000)$$
$$P_e(6;000100)$$
$$P_e(6;000010)$$
$$P_e(6;101111)$$
$$P_e(6;110111)$$
$$P_e(6;111011)$$
$$P_e(6;111101)$$

and so on. The ranking is clearly by number of times the distribution is 'chopped' rather than 'shaved,' i.e., chopped being, then, the conditional case of the stock price close *changing* from above to below or from below to above. This will be the form of the perturbative expansion.

We want to collect together the relevant terms in the price expansion grouped by exercise on any given day and then grouped into their perturbative expansion orders. The result is shown in fig. 4.

We may now 'compactify' the series by 'adding together' terms that differ only by a single digit and are in the same column in page 4, one of course being a 0 and the other a 1. By 'adding together' these terms I mean valuing on one grid run as not chopped up at all on that day (this will be signified by the digit 2). A moment's reflection should convince the reader of this, adding the price under the conditional probability of being below and the price for stock being above on any one day, all other conditions being equal, results in the price without a condition on stock being above or below on that day. This allows a faster calculation. The resulting perms are shown in fig. 5.

This then, is a formulation of the full solution to the problem. We may calculate the price slides of all of the grids listed in fig. 5. The solution may then be found to arbitrary accuracy by including all the grids in successively higher order rankings.

To recap, the payout of the called-on-day-4 contribution, for instance, is propagated backward using the Green's function on three different grids, the highest order being 1. This order one algorithm is: chop off payout-on-day-4 below strike; one day (backward) evolution; chop off below strike; one day evolution; chop off below strike; one day evolution; chop off above strike; one day evolution.

Valuation to zeroth order requires the grid to be run twice, to first order requires ten grid runs and valuation to all orders requires 71 grid runs.

Fig. 4

3 out of 5 provisional call: valuation on day 1 and callable on day 3 to day 9. The perms of stock closes above and below are arranged by day called, and order (number of contrary chops).

called on day 3	called on day 4	called on day 5	called on day 6	called on day 7	called on day 8	called on day 9	not called	order
111							000000000	0
	1110	11100	111000	1110000	11100000	111000000	000000001	1
							000000011	1
							100000000	1
							110000000	1
	1011	10011		1110001	11100001	111000011	000000010	2
	1101	11001			11100011	111000001	000000100	2
							000000110	2
							000001000	2
							000001100	2
							000010000	2
							000011000	2
							000100000	2
							000110000	2
							001000000	2
							001100000	2
							010000000	2
							011000000	2
							100000001	2
							100000011	2
							110000001	2
							110000011	2
		10110	110100	1001100	10011000	100110000	000000101	3
		11010	100110	1011000	10110000	101100000	000001001	3
			101100	1100100	11001000	110010000	000010001	3
			110010	1101000	11010000	111000010	000100001	3
					11100010	111000100	000100011	3
						111000110	000110001	3
						110100000	001000001	3
							001000011	3
							001100001	3
							001100011	3
							010000001	3
							010000011	3
							011000001	3
							011000011	3
							100000010	3
							100000100	3
							100000110	3
							100001000	3
							100001100	3
							100010000	3
							100011000	3
							100100000	3
							101000000	3
							110000010	3
							110000100	3
							110000110	3
							110001000	3
							110001100	3
		10101		1100101	10110001	100110001	000001010	4
				1101001	11001001	101100001	000010010	4
					11010001	101100011	000010100	4
						110010001	000100010	4
						111000101	000100100	4
						110100001	000101000	4
						110100011	001000010	4
							001000100	4
							001000110	4
							001001000	4
							001010000	4
							001100010	4
							010000010	4
							010000100	4
							010000110	4
							010001000	4
							010001100	4
							010010000	4
							010100000	4
							011000010	4
							011000100	4
							011000110	4
							100000101	4
							100001001	4
							100010001	4
							100100001	4
							100100011	4
							101000001	4
							101000011	4
							110000101	4
							110001001	4
			101010	1010100	10101000	101010000	000100101	5
					11001010	101100010	000101001	5
					11010010	110010010	001000101	5
						110010100	001001001	5
						110100010	001010001	5
						110100100	010000101	5
							010001001	5
							010010001	5
							010100001	5
							010100011	5
							011000101	5
							100001010	5
							100010010	5
							100010100	5
							100100010	5
							100100100	5
							100101000	5
							101000010	5
							101000100	5
							101000110	5
							101001000	5
							110001010	5
					10101001	101010001	001001010	6
						110100101	001010010	6
							010001010	6
							010010010	6
							010010100	6
							010100010	6
							010100100	6
							100100101	6
							100101001	6
							101000101	6
							101001001	6
						101010010	010100101	7
							101001010	7

02/11/99

Fig. 5

'Compacted' 3 out of 5 provisional call. Fig. 4 perms that differ by one digit are reduced to not callable on that day. The process is repeated until minimum version is obtained.

called on day 3	called on day 4	called on day 5	called on day 6	called on day 7	called on day 8	called on day 9	not called	order
111							002200022	0
	1110	11100	111000	1110002	11100022	111000000	100200022	1
							110000022	1
	1011	10011				111000201	020200102	2
	1101	11001					022000110	2
							202001002	2
							220001100	2
							000210002	2
							200011000	2
							010200022	2
							011000022	2
		10110	100110	1001100	10011000	100110002	210001002	3
		11010	101100	1011000	10110002	101100022	100200102	3
			110010	1100102	11001020	110010200	102000110	3
			110100	1101002	11010020	110100202	100010002	3
						111000012	101000022	3
						111000120	110000102	3
							110000110	3
		10101			11001001	110010001	202001010	4
					11010001		002010010	4
							020010100	4
							200101002	4
							021000102	4
							001010002	4
							010010002	4
			101010	1010100	10101002	101010020	210001010	5
						110010010	100010010	5
						110100012	100010100	5
							101000102	5

The final step is due to a simple observation. We may avoid the multiple grid runs to value the zeroth order by valuing the option as callable only on certain days, and this reproduces the zeroth order term approximately to order one.

Consider only the not-called zeroth order term of fig. 5. This is an option, as described above that at expiration pays $1 for all stock prices. Then, chop off above strike and backward evolve one day, and repeat. Then backward evolve (no choppings) for two days, then chop off below and evolve three times, then two days of just evolution. Note that the zero order terms are a sum of, the above option and an option: call on day 3, chopped off above strike and one day evolve repeated three times.

Focus on the zeroth order term in the $1 payout, i.e., the not-called term. Now value a new option that is callable on the days we have a zero in this term, and not callable on the days we have a 2 in the term. This generates a new option that is expressible in the above notation as a sum of terms. The terms in this sum are shown in fig. 6. The difference between the sums of the fig. 6 series and the zero-order fig. 5 series is first order. This approximation to the zeroth order term is the central result of this paper, the proposed approximate and qualitatively compelling solution to the soft-call valuation problem.

I leave it to the reader to work out the algorithm that calculates on which trading days the option should be modeled as callable and on which it should be non-callable, given the call specs of the general soft-call case, n out of m, and start- and end-dates, together with the history of closes (the lessor of $m-1$ closes and the days since $m-1$ days before call start).

To recap, we have a term-by-term perturbative expansion of the solution, and we now also have a single-grid-run algorithm that approximates the zeroth order solution.

The approximate zeroth order solution and the sum of all exact terms up to fourth order are plotted (together with the difference) for a 4 out of 6 knock-out option struck at $80, expiring in 8 days to get one dollar, in fig. 6a and fig. 6b. Also the approximate zeroth order and exact sum to zeroth order for a 20 out of 30 knock-out dollar struck at $80, expiring in 60 days time, are plotted in fig. 7a and fig. 7b.

Fig. 6

called on day:

3	4	5	6	7	8	9	not called on 9
122	1022	10022			12200022	102200022	002200022

Expansion of an option which is callable on some days and non-call on other dates: 002200022. It is called on days with a zero and non-callable on days with a 2.

Fig. 6a

4 out of 6 knock-out dollar. Struck at $80.
The approximate order 0 and exact order 4 solutions

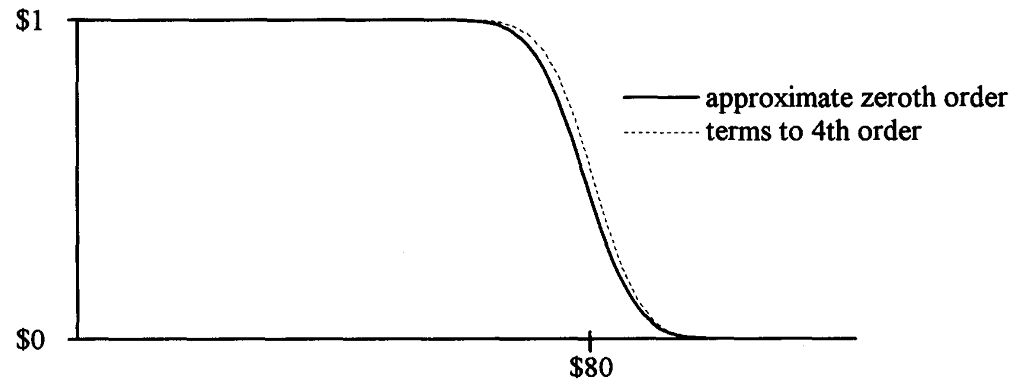

Fig. 6b

4 out of 6 knock-out dollar. Struck at $80.
Difference between approximate order 0 and exact order 4 solutions

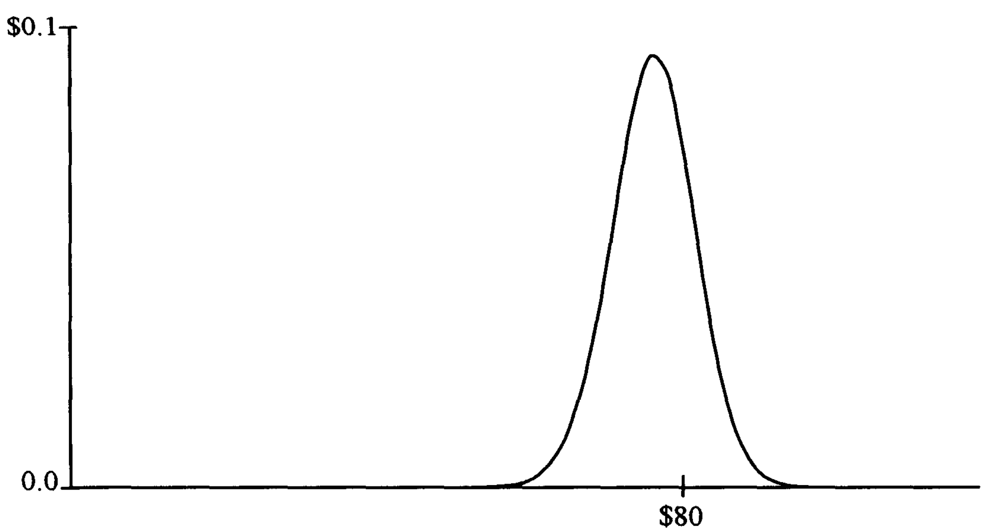

Fig. 7a

Knock-out dollar with $80 trigger: 60 days to maturity of a 20 out of 30 day provisional call. The zeroth order and exact first order approximations are shown.

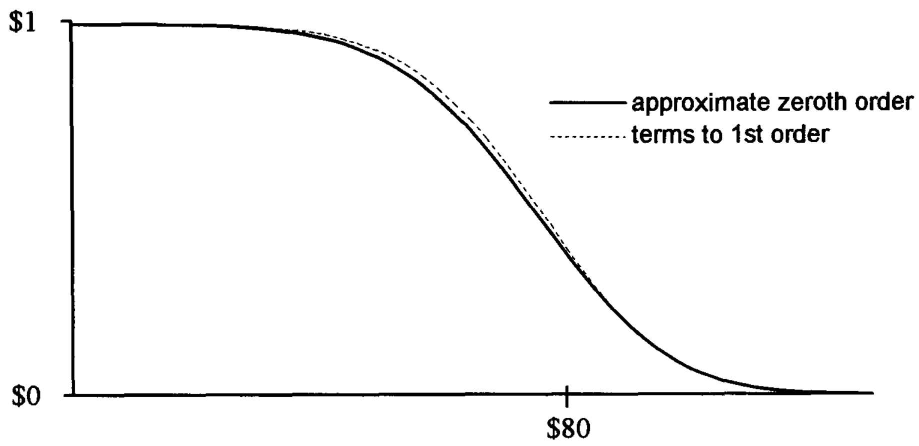

Fig. 7b

The difference between the approximate zeroth order and the exact first order term.

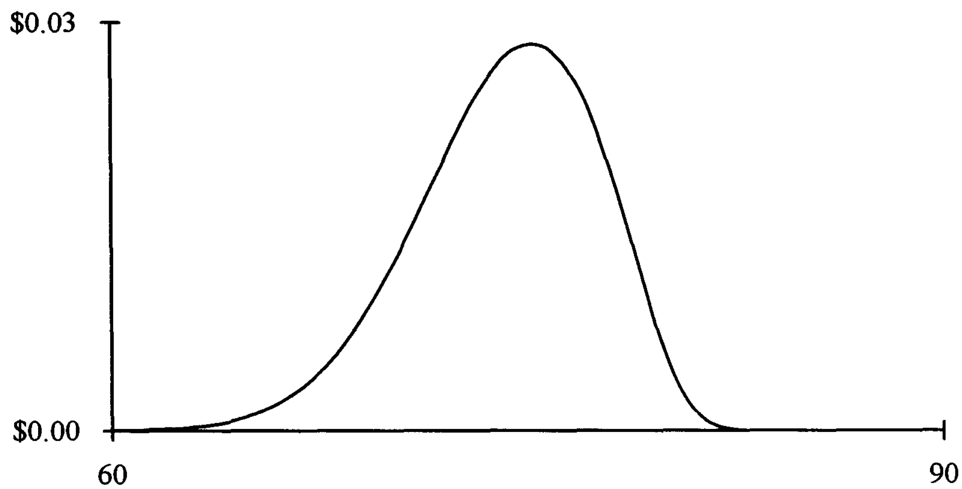

Full Convertible Bond

The generalization to CBs is now trivial. The approximate zeroth order solution is nothing more than making the bond callable on some days and not on others during the provisionally callable period.

The additional inputs beside those required for the one-touch model are the history of closes for the lessor of: *m-1* days and back *(m-1)* days before the soft-call period start. From this a grid is constructed which values an 'effective bond' callable only on certain days until the actual bond becomes callable, and the effective bond is then callable on every trading day. The grid must have at least a few time steps per day, but this is only during the provisionally callable period. The result is qualitatively as one would expect. The accuracy is, roughly, between a few pennies and a dime on the dollar of the benefits of provisional call protection (which is usually getting or losing the next coupon).

The result for usual vols and bond specs, only gives significant differences to the one-touch model when the bond becomes provisionally callable within the next few months (or is already provisionally callable) and for stock forward within 5-10 percent of the trigger.

A concrete example is illustrative: fig. 8 shows a plot of price, delta and gamma over a stock slide for the one-touch and approximate provisionally callable (as described above) one-factor lognormal model of the Alza 5% of 2006 convertible bond, as it might be seen on 6/1/1999 (assuming values for credit spreads and interest rate environment). This valuation date is a month after the bonds become provisionally callable, and I have assumed the stock price has never closed above the trigger. While the price difference is not more than a bond point, the gamma plot is vastly different, and underlines the very short dated nature of the provisional call feature. Variation of the input history shows each extra close above the trigger in the past to be worth, roughly, a loss of a nickel from the bond's premium. This is the right order of magnitude: it is about one bond point divided by 19, and if 19 closes above have been observed the bond will look very similar to its one-touch model value.

Fig. 8. - Price

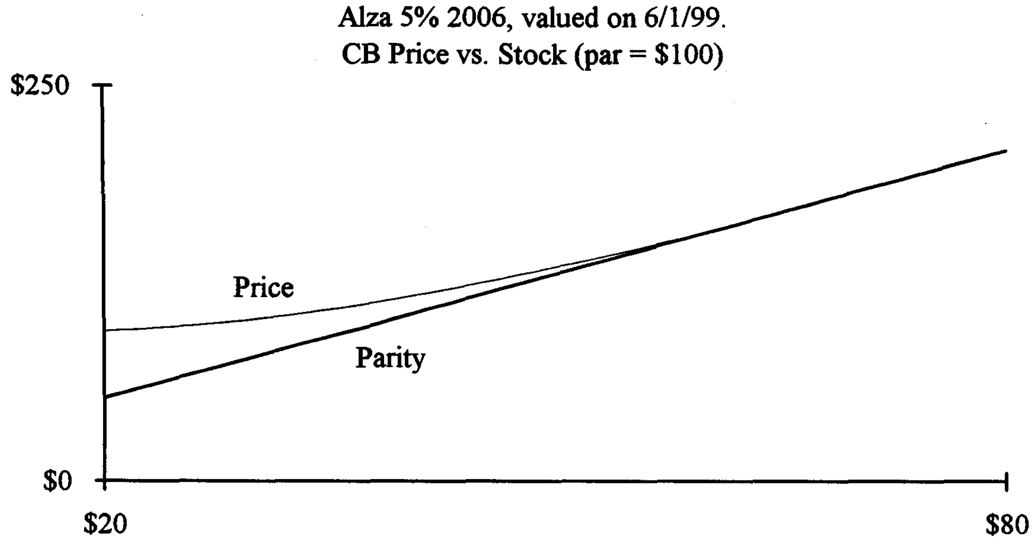

Alza 5% 2006, valued on 6/1/99.
CB Price vs. Stock (par = $100)
$250
Price
Parity
$0
$20
$80

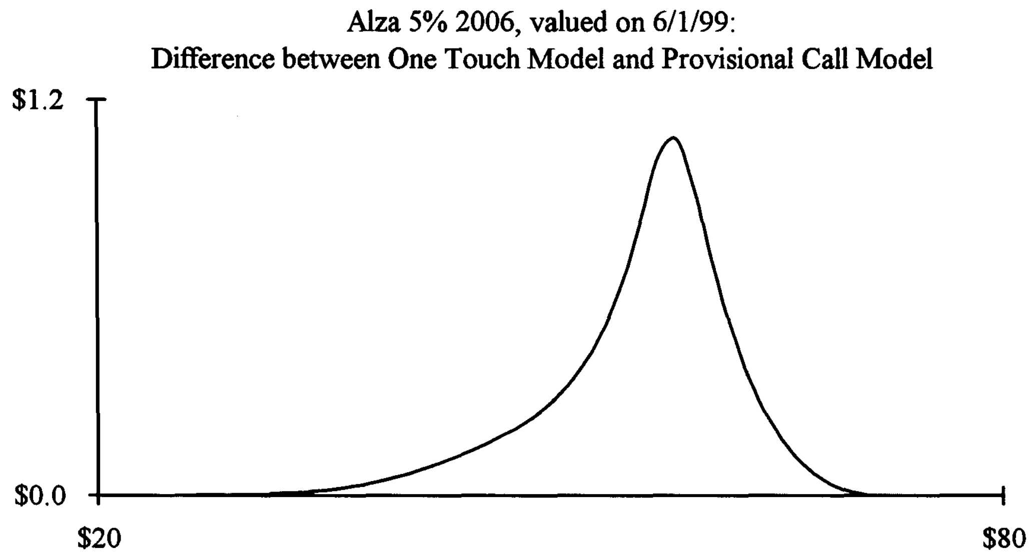

Alza 5% 2006, valued on 6/1/99:
Difference between One Touch Model and Provisional Call Model
$1.2
$0.0
$20
$80

Fig. 8. *(continued) - Delta*

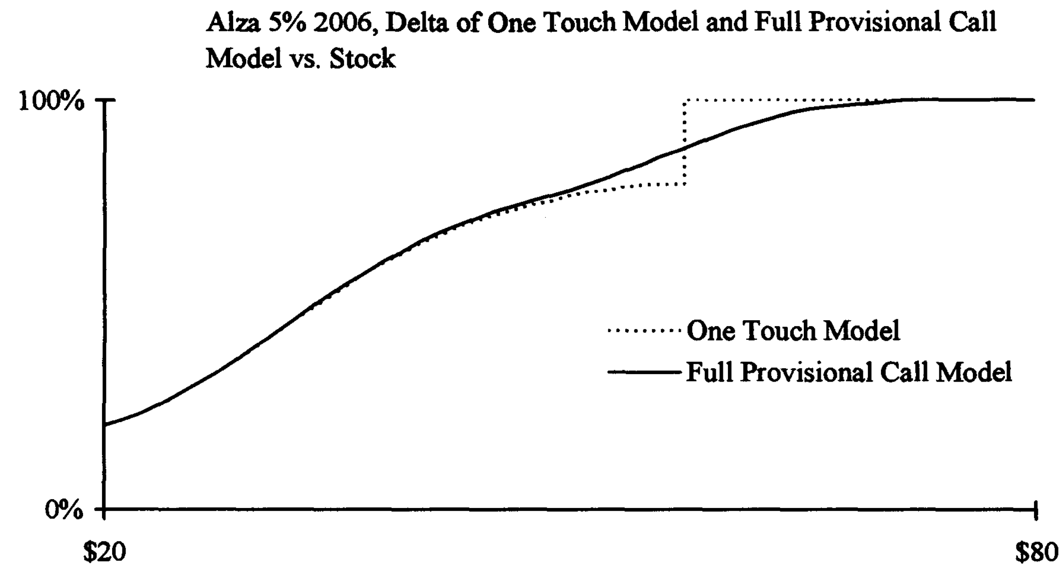

Fig. 8. *(continued) - Gamma*

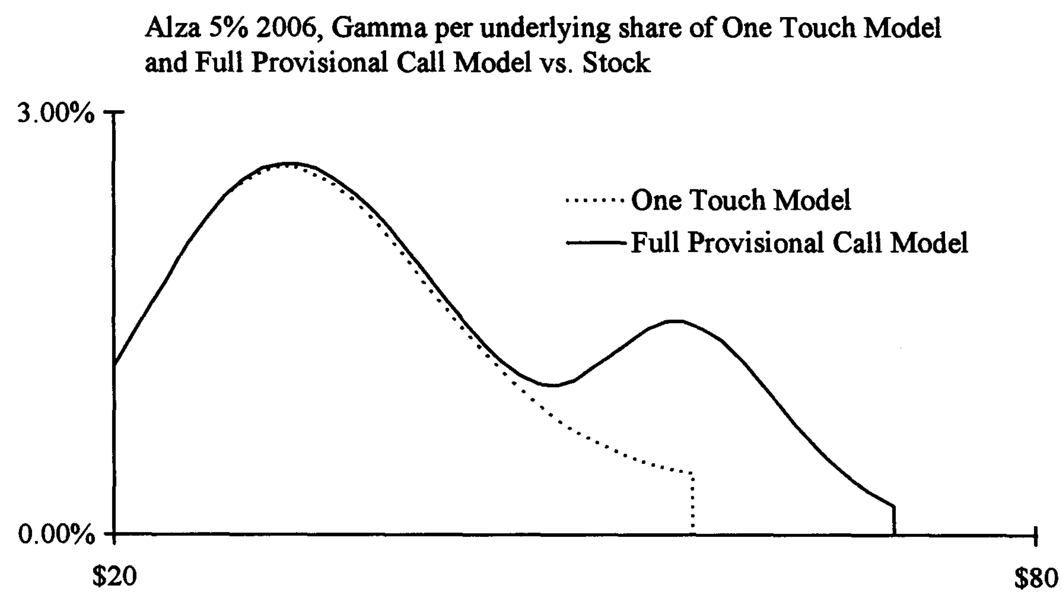

Conclusion

We have outlined a practical solution to the problem of modeling the soft-call feature of many convertible bonds. The method may be implemented by simple alteration of existing one-touch models of provisional call, and it is a good approximation to the exact solution for such a model. It passes most qualitative tests that are relevant except that it is not sensitive to calculations that differ by higher orders as defined above. The grid will be slowed down as it is necessary to run with at least a few time steps per day, but this is merely a reflection of the fact that we are valuing a short dated option embedded in a long dated instrument.

Acknowledgments

Thanks to Arkady Goldrin, at Highbridge Capital Management, LLC, who wrote or assisted with most of the computer codes used in this project, and also to Robert Wong, at Highbridge Capital Management, LLC, who prepared draft versions of this paper.