

EMBEDDED CALL OPTIONS AND REFUNDING EFFICIENCY

C. Douglas Howard* and Andrew J. Kalotay*

I. INTRODUCTION

A prudent liability manager will continuously monitor interest rates and will call a bond when the company's refunding rate is sufficiently low. Efficient use of embedded call options requires explicitly identifying this sufficiently low "target rate" [6]. This target rate depends on many factors, including the bond's coupon, the remaining time to maturity, the entire schedule of call prices, the perceived behavior of interest rates, tax considerations, the structure of the refunding issue, and the issuer's risk preference [1]. Refunding "efficiency" is the yardstick that measures the combined impact of these considerations. This paper focuses on the methodology underlying the computation of refunding efficiency and the target refunding rate.

*Salomon Brothers Inc.

Advances in Futures and Options Research, Vol. 3, pages 97-117.

Copyright © 1988 by JAI Press Inc.

All rights of reproduction in any form reserved.

ISBN: 0-89232-926-2

II. THE BREAK-EVEN RATE: GOOD FOR A ONE-TIME CALL

The break-even rate is a fundamental element in refunding analysis. The traditional method for evaluating the savings from a refunding is to compute the present value of the resulting cash flow savings.¹ The break-even rate is the refunding rate that will generate a present value savings of 0%.

Consider for example, a 12% 40-year bullet bond callable after 5 years at a price 110.29% of par (the "long 12s"). Assume that the issuer loses the right to refund the bond if he passes up that opportunity in year 5. The issuer, therefore, has a "one-time" call option exercisable in year 5 of the bond's life (see Figure 1, Column 2).

Figure 2 shows the present value savings as a percentage of par for a refunding in year 5 at various refunding rate levels. The break-even rate occurs at 10.5%–11%. Specifically, a present value savings of 0% is generated by a refunding at 10.85%.

Because the call option is a one-time European-type option, the issuer should exercise it if the refunding rate is below the break-even rate. In this example, the issuer has nothing to lose by exercising the option, however small the savings may be.

<i>Year</i>	<i>"One-Time" Call</i>	<i>"Two-Time" Call</i>	<i>Continuous Call</i>
0–4	NR	NR	NR
5	110.29%	110.29%	110.29%
6	NR	109.94	109.94
7	NR	NR	109.59
35	NR	NR	100.35
36–40	NR	NR	100.00

NR Not refundable

Figure 1. Schedule of optional call prices.

<i>Refunding Rate</i>	<i>PV Savings from Refunding</i>
9.00%	21.51%
9.50	15.00
10.00	9.05
10.50	3.60
11.00	-1.41

Figure 2. One-time call option.

III. A HYPOTHETICAL "TWO-TIME" CALL OPTION

In reality, most call options are exercisable continuously at any time after the refunding protection expires. This complicates the decision, because in this case, the issuer has something to lose by exercising the option — namely, the opportunity to exercise the option at a later date, possibly achieving greater savings than are currently available. This could result from a decrease in interest rates or from a declining call price schedule.

To illustrate this complication, slightly modify the option so that the bond may *also* be called in year 6 at a price of 109.94% of par (see Figure 1, Column 3). The issuer now has two alternatives:

1. It can call the bond in year 5; or
2. It can wait until year 6 and call the bond then if savings are available. If the issuer chooses to wait a year, it will, in effect, in year 6 have a "one-time" option that should be exercised only if savings are positive.

If the issuer has a strong conviction about interest rate movements over the next year, the refunding decision is clear. The issuer will wait and call in year 6 if it believes that interest rates are declining, but call the bond in year 5 if it believes that interest rates are rising. The decision is much more difficult in the absence of such a conviction. However, given a reasonable model of interest rate behavior, one can still make a rational decision.

Suppose that in year 5, the issuer could refund the long 12s with a new 10% bond. As we have seen in Figure 2, the present value savings of such a refunding would be 9.05% of the par amount refunded. Instead of calling the bond and refunding in year 5, the issuer may wait and address the question 1 year later in a different refunding rate environment.

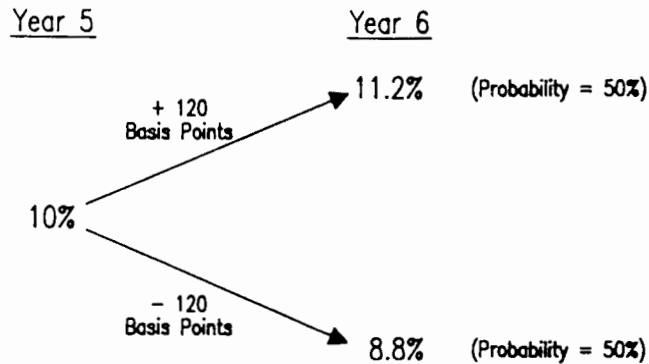


Figure 3. A simple model of interest rate behavior.

For simplicity, assume that only two equally likely outcomes are possible. From year 5 to year 6, rates will rise or fall by 120 basis points from their current level of 10% (see Figure 3). This equates to an annual "volatility" of 1.2%/10%, or 12%.² If the refunding rate increases by 120 basis points to 11.2%, the savings decline to 0%. This is because 11.2% is above the break-even rate, and the option would not be exercised. If the refunding rate declines to 8.8%, the savings rise to 22.20% of par.³ Because these two events are equally likely, the "expected value of waiting" to year 6 is $0.5 \times 0\% + 0.5 \times 22.20\%$, or 11.1%. The theoretical value of the call option in year 5 is equal to the *greater* of the savings that are available from a current exercise (9.05%) and the expected value of waiting (11.1%), or 11.1%.

Figure 4 shows current savings (in year 5), the expected value of waiting

Refunding Rate In Year 5	PV Savings of a Call In Year 5	Expected PVS of Waiting To Year 6	Theoretical Value of Call Option
9.00%	21.51%	21.22%	21.51%
9.50	15.00	15.21	15.21
10.00	9.05	11.10	11.10
10.50	3.60	8.37	8.37
11.00	-1.41	5.87	5.87

Figure 4. Two-time call option.

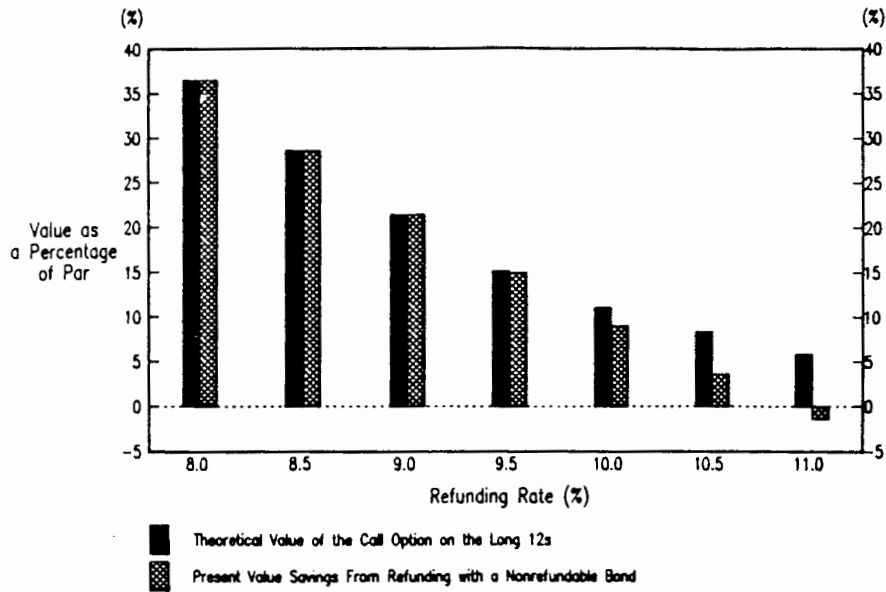


Figure 5. Option valuation applied to the long 12s (“two-time” call option).

and the theoretical option value over a range of refunding rate environments. In each refunding rate environment, the theoretical option value is the greater of current savings or the expected value of waiting. Figure 5 illustrates the option value and the savings of a current call. In refunding rate environments below roughly 9.5%, the savings generated by a refunding match the value of the call option. In contrast, if the refunding rate is roughly 10% or higher, the option value greatly exceeds the savings of a current refunding. The prudent liability manager will compare the savings generated by a refunding with the value of the option forfeited by the exercise of the call and will refund the bond only if the savings represent a substantial fraction of the option’s value.

IV. EFFICIENCY

As a function of the refunding rate environment, the percentage of the option value that is captured by a refunding is shown in Figure 6. This is known as the “efficiency” of the refunding. For example, a refunding rate of 10% will generate 9.05% savings while the option value is 11.1%. The efficiency of this refunding would, therefore, be 81.5%.




Figure 6. Efficiency of refunding the long 12s in year five ("two-time" call option).

Each individual issuer must decide what degree of efficiency is sufficient to trigger a refunding. This decision should be made in the context of risk versus potential reward. A game of chance analogy illustrates the point. Suppose a coin is flipped, and if heads comes up, the player is paid \$200. If tails comes up, the player is paid nothing. The "theoretical" value of the game to the player is $\$200 \times 0.5 + \0×0.5 , or \$100. Now suppose the player is offered money to *not* play the game. How much would the player settle for? A risk-neutral player would settle for nothing less than the full \$100 — the theoretical value of the game. However, a player with some degree of risk aversion would settle for less than the value of \$100 just to avoid the possibility of ending up with nothing. The greater the degree of risk aversion, the less the player would settle for.

The debt issuer must make the same decision about the call option. Consider again the long 12s in a 10% refunding rate environment. The issuer can refund in year 5, settling for a certain 9.05% savings, or it can "flip a coin" and wait until year 6. In year 6, the issuer will either get 22.2% savings or nothing at all. Waiting is riskier, but the "expected" payoff of 11.1% savings ($0.5 \times 0\% + 0.5 \times 22.2\%$) is greater.

For a risk-averse issuer, there is some level of efficiency below 100% where the "expected" benefit of deferring a refunding will not justify the additional risk. In the absence of institutional, political, and regulatory considerations that dictate otherwise, we would recommend 85% efficiency

as a reasonable balance between reward and risk. In the preceding coin-tossing analogy, this is equivalent to settling for \$85 rather than playing the game.

For the long 12s with the hypothetical “two-time” call option, a refunding rate of 9.94% will lock in present value savings equal to 85% of the value of the call option (see Figure 6). This 85% efficiency “target refunding rate” means that the issuer would call the bond in year 5 only if the refunding rate is at or below 9.94%.

V. A CONTINUOUS OPTION

We simplified the preceding analysis of the valuation of the call option in several respects. First, we assumed that the issuer could exercise the option at only two moments of the bond’s life — the end of years 5 and 6. In reality, an embedded call option is usually exercisable at any time after the refunding protection expires. Figure 1 (Column 4) shows a realistic set of call prices declining to par 5 years prior to the maturity of the bond. Second, our model of refunding interest rate behavior (up or down from year to year) is simplistic. The “binomial” model can be improved by dividing a year into several steps. With each step, the refunding rate moves up or down. Figure 7 illustrates this process with four quarterly steps of 60 basis points up or down per year. Clearly this more closely approximates the actual continuous movement of interest rates. The magnitude of the quarterly movement is smaller than the annual movement in Figure 3 (60




Figure 7. A more reasonable model of interest rate behavior.




Figure 8. Efficiency of refunding the long 12s in year five (continuous option).

basis points versus 120 basis points) to preserve the *annual* volatility of 12%.⁴

Figure 8 shows the efficiency of calling the long 12s if the option is exercisable at *any time* after year 5. For computational purposes, each year is divided into a sufficient number of discrete periods to very closely approximate the continuous nature of the option and of interest rate behavior.

At any given refunding rate level, the efficiency of calling the long 12s in year 5 is lower if the option is continuously exercisable than if it is exercisable only in years 5 and 6. This is because the issuer, when exercising and calling the bond, is parting with a much more valuable option. The theoretical option value now reflects the issuer's right to refund the bond at any opportune moment starting in year 5 — certainly a greater value than the "two-time" right to refund only in years 5 and 6. Since the option value is the denominator of the efficiency calculation, the higher option value translates into a lower efficiency. Therefore, a lower refunding rate environment is necessary to generate a specified level of efficiency. In fact, if the option is continuously exercisable, the 85% efficiency target refunding rate in year 5 is 9.37%, almost 60 basis points lower than the corresponding target rate on the "two-time" option.

VI. THE IMPORTANCE OF VOLATILITY

Remarkably, the only numerical assumption underlying the option valuation model is the volatility of the refunding rate, chosen in the preceding examples to be 12% annually. Because the option value is driven by prospective volatility, rather than historical volatility, the assumption is, at best, an estimate. Therefore, it is important to quantify the sensitivity of the decision to refund (or not refund) to the volatility assumption.

Figure 9 shows the efficiency of refunding the long 12s given three different interest rate volatilities. The larger the volatility assumption, the larger the theoretical option value and the lower the refunding efficiency for any given refunding rate environment.

This reflects the enhanced prospects of the refunding rate eventually drifting low enough to generate more savings from a later refunding. This effect translates into a lower target rate given a higher volatility. The 85% efficiency target rate declines to 9.07% given 15% volatility, and rises to 9.70% given a 9% volatility.

Figure 10 plots the 85% efficiency and 100% efficiency target rates as functions of the volatility assumption. Lower efficiency generates a target rate that is less sensitive to the volatility assumption. Fortunately, the ten-basis-point increment of the 85% efficiency target rate for a 1% change in volatility is not substantial. In addition, despite occasional peaks and




Figure 9. Refunding efficiency under several volatility assumptions.




Figure 10. The sensitivity of target rates to the volatility assumption.

valleys, the volatility of long-term interest rates in recent years has generally remained in a range of 9–15% (see Figure 11). The sensitivity of the target rate in this example is increased by the 40-year term of the long 12s. Shorter term bonds, or bonds with less time to maturity, display less sensitivity to the volatility assumption.




Figure 11. Recent historical volatility of long-term treasury bonds.

VII. THE STRUCTURE OF THE REFUNDING ISSUE

We have assumed that the long 12s will be refunded with a *nonrefundable* bond. This assumption is manifested in the nature of the present value savings calculation. This calculation is based on a matched-maturity refunding. The differential between the 12% coupon on the outstanding bond and the refunding issue coupon is present valued through the maturity of the outstanding issue. The call premium is subtracted from this number to obtain a net present value savings. The cash flows are deterministic, and no consideration is given to possible subsequent refundings of today's refunding issues.

If the refunding is effected with a *refundable* bond, this calculation understates the value of refunding by ignoring the potential for additional refundings. A simple two-step calculation is the remedy. First, compute the deterministic present value savings as previously described, pretending that the refunding issue is nonrefundable. Second, compute the value of a new call option on the refunding issue using theoretical option valuation. This option value by definition exactly captures the value of the potential for additional refundings. Therefore, when this new option value is added to the deterministic present value savings, the result is the total benefit of

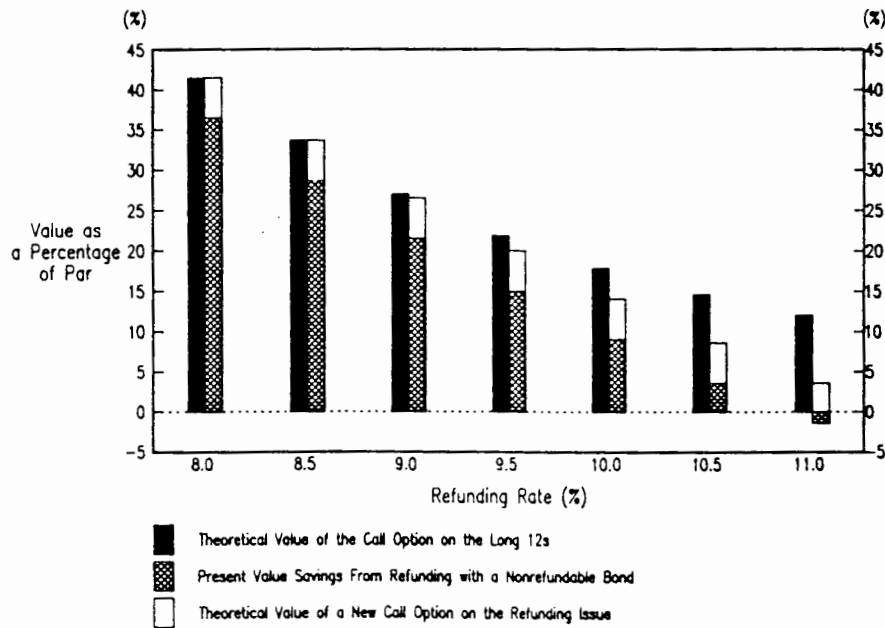


Figure 12. Option valuation with a refundable refunding issue.

refunding with a refundable bond. Figure 12 illustrates this as applied to the long 12s.

Figure 12 warrants several observations. Surprisingly, the value of the option on the *refunding* issue is virtually independent of the refunding rate. Moving right to left from an 11% rate environment to an 8% rate environment, one might expect this option value to increase as refunding rates decline. However, the coupon on the bond in which this option is embedded simultaneously declines from 11 to 8%. In other words, the relative degree to which this new option is "in the money" does not change. In contrast, the value of the option on the *refunded* issues (the long 12s) exhibits the expected behavior and substantially increases as the refunding rate declines from 11 to 8%. This reflects the fact that this option is embedded in a fixed-coupon (12%) bond, while the refunding rate environment is changed.

In addition, the value of the option on the refunding bond is much lower than the value of the option on the long 12s. Two factors contribute to this. First, in rate environments below the break-even rate of 10.85%, the option on the long 12s is "in the money." Second, the option on the long 12s is currently exercisable. The option on the refunding issue is not in the money and cannot be exercised until refunding protection expires in 5 years.

VIII. TAX CONSIDERATIONS

Although we have disregarded taxes for illustrative purposes, in actual applications, refunding efficiency must be calculated on an after-tax basis. Taxes affect cash flows, the discount rate, and the value of the call option. Because of the reasons described, refunding efficiency on an after-tax basis is usually higher than on a pretax basis.

The most obvious tax benefit is that the entire premium above the issue's tax basis — which is usually at or close to face value — is immediately recognized as a current expense. This is particularly important in the case of repurchasing currently nonrefundable debt, because in that case, the premium can be considerably above the initial call price. Because the premium over the call price essentially represents the present value of incremental interest payments until the call date, its immediate deductibility is desirable. This effect is particularly important if the issuer's marginal tax rate is expected to decline in the future.

Figure 13 shows the refunding efficiency for the sample 12s assuming 0 and 34% marginal tax rates. It illustrates that the efficiency increases with higher tax rates.

The effect of taxes on refunding efficiency far exceeds the deductibility of




Figure 13. Efficiency in different tax environments.

the premium. The manner in which taxes affect the value of the call option on the refunding issue is much more subtle, but equally important. Boyce and Kalotay [2] showed that the expected value of this option is higher on an after-tax basis than on a pretax basis. For example, suppose that the market, which consists of essentially nontaxable institutions, demands 50 basis points in exchange for a call feature. Even if this is a fair price on a pretax basis, a taxable issuer should be willing to pay a higher coupon spread, perhaps as much as 80 basis points for the call feature. Because of this tax advantage, refunding with a refundable bond will generate a greater efficiency than refunding with a nonrefundable bond.

IX. TARGET RATES CAN BE CALCULATED AT ANY TIME

For a specific bond, the target rate depends only on two inputs: the assumed interest rate volatility and the required level of efficiency. Once these inputs are specified, the target rate can be computed over the entire life of the issue. In fact, this calculation can be done at the time of issuance. Figure 14 displays the target rates for the long 12s beginning in year 5, when refunding protection expires.




Figure 14. Advance calculation of target rates for the long 12s.

X. SINKING-FUND ISSUES ARE DELICATE⁵

Each sinking-fund payment should be considered as an individual bullet bond, with its own maturity and call price schedule. The original issue can be considered as a *portfolio* of bonds — one for each sinking-fund payment. At any time, the call price of each outstanding sinking-fund payment is the same. We can then determine the efficiency curve for each bond in this portfolio (that is, each sinking-fund payment). This is illustrated in Figure 15 for the long 12s, assuming a sinking-fund commencing in year 10. From these efficiency curves, we can determine the target refunding rate for each payment (see Figure 16).

Depending on its due date, each sinking fund has its own refunding rate, as determined by the issuer's yield curve. If this yield curve is flat (see Figure 17), the issuer should first call the balloon and then work backward, if appropriate. Conversely, if the yield curve is steeply upward sloping (see Figure 18), calling only some of the nearby payments may be optimal. As a general rule, each payment must be considered individually, and the optimum policy may entail calling only part of the issue.

Finally, two further considerations may be relevant in the analysis of sinking-fund issues: the acceleration provision such as a "double-up" and the delivery option. The former entitles the issuer to call at par some multiple of the mandatory sinking fund; the latter enables the issuer to deliver actual securities instead of cash. The delivery option is most valuable




Figure 15. Efficiencies of several sinking-fund payments.

when the issue is selling at a discount.⁶ If a bond is called, it cannot be doubled up against a subsequent sinking-fund payment or purchased at a price below par in the open market if interest rates rise. Because these options favor deferring redemption, their incorporation into the analysis will reduce current refunding efficiency.




Figure 16. Year-five target rates for each sinking-fund payment.




Figure 17. Sinking-fund strategy given a flat yield curve.




Figure 18. Sinking-fund strategy given a steep yield curve.

XI. EFFICIENCY HAS MORE GENERAL APPLICATIONS

Refunding efficiency is motivated by the idea that the interest savings gained from calling a bond should be compared with the theoretical value of the call option, which reflects current and *potential* opportunities. This concept readily applies to many situations other than a standard call that the issuer may wish to undertake to eliminate high-coupon debt and, thus, improve its debt service. Two obvious examples are repurchase programs for currently nonrefundable debt, either through tender offers or open market operations, and advance refunding of municipal issues. Another possible way of capturing the value of the call option is by utilizing the emerging market of options on interest rate swaps. There are various strategies for receiving, in the form of a cash payment, a high percentage of the theoretical value of the call option on the outstanding bond.

Although the call option is not involved directly, it plays a crucial role in these transactions. In the case of a repurchase, the bonds tend to be relatively cheap as they trade on a yield-to-call basis. This phenomenon is even more apparent in the case of "high to low" municipal refundings, in which the "effective purchase price" (that is, the size of the escrow account) is explicitly determined by the issuer's refunding rate and the initial call price of the outstanding issue.⁷ The saving is essentially the difference between pricing the bond on a yield-to-maturity basis versus on a yield-to-call basis. This simple functional relationship between refunding rate and effective purchase price enables us to compute straightforwardly the efficiency of an advance refunding: The only determinant of future effective purchase price is the issuer's future refunding rate [3]. In contrast, the future price and refunding efficiency of a tender/repurchase candidate depend on the issuer's funding rate of matching maturity, the interest rate to call, and, even more importantly, supply/demand conditions.

XII. CONCLUSION

Prior to the advent of embedded option valuation and refunding efficiency, the refunding decision was usually based on a simple present value savings calculation. The issuer decided by "gut feel" if savings were sufficient in the prevailing interest rate environment to refund the bond. Unfortunately, the appropriate level of savings depends heavily on the specifics of the bond. For example, savings of 5% would probably not warrant refunding a long telephone bond but probably would be more than enough to refund an intermediate-term bond maturing in 2 or 3 years. Furthermore, the savings calculation considers only one side of an equation. Although this calculation values the benefit of a refunding, it entirely ignores the hidden cost

of the refunding — namely the value of the call option that the issuer relinquishes to effect the refunding.

The efficiency benchmark overcomes both of these problems. The appropriate degree of efficiency (whether 85%, 90%, or even 100%) is related to the risk preference of the issuer and not to the specifics of the refunding candidate. If 85% efficiency is a good benchmark for a long bond, then it is also a good benchmark for an intermediate bond. More importantly, the efficiency calculation directly compares the value obtained from a refunding with the value of the forfeited option.

Today's volatile interest rate environment poses quite a challenge and opportunity for the liability manager to effectively and efficiently realize the full value of embedded call options.

APPENDIX

This appendix illustrates an intuitively appealing definition of interest rate volatility. However volatility is defined mathematically, conceptually it is intended to measure the degree of magnitude of future interest rate uncertainty. If interest rates are "highly volatile," one can estimate with only minimal confidence what interest rates will be after 1 year, for example. Conversely, low volatility implies that next year's rates can be predicted with a high degree of confidence.




Figure A1. Volatility given a normal distribution.




Figure A2. Normal distributions with 9 and 15% volatility.

Suppose, for example, that the refunding rate for the issuer of the long 12s is currently 10%. The issuer cannot know what the refunding rate will be after 1 year, but assume that it is known that the rate is normally distributed with a mean of 10% and a standard deviation of 1.2% (see Figure A1). Volatility is defined to be the standard deviation of 1.2%




Figure A3. Discrete distribution with one step per year.




Figure A4. Discrete distribution with four steps per year.

divided by next year's "expected" rate level (which is also the current rate level) — expressed as a percentage. In this example, interest rate volatility is $1.2\%/10\% \times 100\%$, or 12% .

The standard deviation (and, therefore, volatility) is a direct measure of the issuer's degree of uncertainty about next year's refunding rate. Figure A2 illustrates the link between volatility and future interest rate uncertainty. In this exhibit, the probability distributions of next year's refunding rate are shown assuming 9 and 15% volatility and a normal distribution. The 95% confidence interval for each distribution shows the interval centred at 10% in which one can be 95% certain that next year's interest rate will fall. A higher volatility gives rise to a larger interval and, therefore, greater interest rate uncertainty.

This definition of volatility works for any type of future interest rate probability distribution — not just the normal distribution. Figures A3 and A4 show the discrete probability distributions resulting from the interest rate behavior illustrated in Figures 3 and 7, respectively. In both cases, the standard deviation around the mean of 10% is 1.2%. This equates to a volatility of $1.2\%/10\%$, or 12% .

NOTES

1. For a discussion about alternative approaches to this problem, with regard to the specification of the cash flows and the choice of discount rate, see Kalotay [7].

2. See Appendix.
3. To be comparable with the 9.05% savings in year 5, the savings available in year 6 is discounted back to year 5.
4. See Appendix.
5. For further information, see Kalotay and McIntyre [8].
6. See Kalotay [5].
7. Advance refunding is mathematically equivalent to the index call. See Leibowitz [4].

REFERENCES

1. Boyce, W. M., and A. J. Kalotay, "Optimum Bond Calling and Refunding," *Interfaces* (November 1979).
2. Boyce, W. M., and A. J. Kalotay, "Tax Differentials and Callable Bonds," *Journal of Finance* (September 1979).
3. Gurwitz, A. S., A. J. Kalotay, and D. Howard, *Call Provisions and Efficient Refunding in Municipal Finance*. Salomon Brothers Inc, March 1987.
4. Leibowitz, M. L., *The Index Call*, Salomon Brothers Inc, April 1981.
5. Kalotay, A. J., *The Effect of Sinking Funds on the Cost of Debt*. Salomon Brothers Inc, September 1981.
6. Kalotay, A. J., *The Interest Rate Risk of Callable Bonds*. Salomon Brothers Inc, September 1982.
7. Kalotay, A. J., "On the Structure and Valuation of Debt Refundings," *Financial Management* (Spring 1982).
8. Kalotay, A. J., and P. McIntyre, *Sinking Fund Issues: A Challenge for the Active Manager*. Salomon Brothers Inc, November 1984.