Valuing Employee Stock Options

Nalin Kulatilaka and Alan J. Marcus

The Financial Accounting Standards Board (FASB) is considering a controversial proposal that will require firms to calculate and recognize as a cost of compensation the value of employee stock options at the time those options are granted. Conventional option pricing models, however, are not well suited to the valuation of employee options because such options are nontransferable and, therefore, may be exercised early when an unconstrained investor ordinarily would sell the option. Thus, conventional pricing models can substantially misvalue employee options.

The model of option valuation presented here accounts for an employee's propensity to exercise an option early. It shows that (1) employee stock options may be worth much less than would be suggested by conventional models, even using expected terms as low as one-half the stated option life; (2) the values of employee stock options are sensitive to variables that do not even appear in conventional option pricing models; (3) the values of employee stock options may fall when stock volatility rises; and (4) values of employee options can be less than the "minimum option value" that has appeared in the accounting literature.

Under current accounting rules, grants of employee stock options need never result in a charge to the firm's income statement as long as the exercise price is greater than or equal to the stock price at the time of the grant. If and when the options are exercised, the exercise price is simply added to paid-in capital and the number of outstanding shares increases. Given these rules, the firm does not need to value any options so granted.

This situation might change soon, however. The FASB is considering a controversial proposal that will require firms to calculate and recognize as a cost of compensation the value of employee stock options at the time those options are granted.[1] Under this proposal, firms will require a reasonable means to value those options.

The new rule would substantially affect reported earnings of firms that rely heavily on option-based compensation. Start-up companies commonly use employee stock options as a major part of compensation packages, and options are becoming increasingly popular among more-established firms. Some estimates show that the 1992 earnings of technology companies, in which options compose a large share of total compensation, would have been reduced up to 50 percent by the new rules.[2]

Even if FASB ultimately backs off from its proposal, financial analysts and investors still need to be able to value these options. To the extent that such options represent substantial claims against a firm, they can have a large impact on the market value of the stockholders' equity.

FASB argues in its exposure draft that standard option pricing models, such as the Black–Scholes model or binomial model, would provide reasonably simple means to value these options. In fact, dividend-adjusted versions of the Black–Scholes call option or warrant pricing models have been used widely in the accounting literature to analyze employee stock options.[3] The Black–Scholes or binomial models, however, can lead to severe noncomparabilities across firms that use even slightly different assumptions regarding stock volatility, dividend rates, and premature-job-termination rates, or different methods of accounting for delayed vesting.[4]

These models, moreover, are not well suited to deal with the issue of early exercise that arises in the context of employee options. Normally, an option holder who wants to reduce an options position would sell all or part of the holdings. Because employee stock options are not transferable, however, the only way to cash them in is to

Nalin Kulatilaka is a Professor of Finance at Boston University. Alan J. Marcus is a Professor of Finance at Boston College.

exercise them. Therefore, employee options may be exercised earlier than would seem optimal based on valuation models for unconstrained investors. Such early exercise reduces the market value of the options.

The FASB exposure draft recognizes the potential for early exercise. It suggests that employee options be valued using an option pricing model with the expected time to exercise replacing the stated maturity of the option, which in fact is the option's maximum possible life. For example, a ten-year option might be valued using the Black–Scholes formula assuming a six-year expected term.

This suggestion raises a new question. How can the expected term of the option be forecast? One might consider using historical experience to estimate average times until exercise, but this approach is subject to considerable error because exercise experience will depend in large part on stock returns during the sample period. For example, option lives will be longer when the stock performs poorly because out-of-the-money options are never exercised. Because stock market performance is so variable across time periods, estimates of average times until exercise based on sample periods of even several years can be unreliable guides to future experience. Option values are quite sensitive to assumed option life, implying that option values derived from historical exercise experience will be unreliable.

Even if one could estimate the expected term of the option, using this value in the option pricing model would not be correct for two reasons. First, because the option value is a nonlinear function of the term of the option, straightforward use of the expected term in the option pricing formula will lead to additional error in calculated option value. In addition, the time until exercise of an employee stock option generally will vary inversely with the stock price and, therefore, the option payoff. This inverse relationship implies that even if one replaces the stated term of the option with the expected time until exercise in a Black–Scholes-type model, the formula value will overstate the actual value of the option.

The model of option valuation we present here explicitly accounts for an employee's propensity to exercise the option early. Early exercise is based on portfolio diversification motives. Because employee options are not tradable, the only way for employees to diversify portfolios that are heavily dependent on the fortunes of their firms is to exercise their options, even if that action is not value maximizing according to standard valuation models.[5] We solve for optimal exercise policy, given the constraint that the options are not transferable. In this way, the exercise propensities derived are not a function of stock market performance over a particular sample period but reflect probabilistically the full range of possible outcomes for the stock.

Our focus in this study is on portfolio maldiversification as it affects option value, not on the wedge between a maldiversified employee's subjective valuation of nontradable options and the market value of those options. The market value of employee stock options is the proper concern of FASB and of investors in the firm. We show that portfolio diversification considerations affect the market, as well as the subjective value of these options, through the effect on exercise policy.

DIVERSIFICATION AND EARLY EXERCISE

In this section, we use a simple two-period example to illustrate the effect of diversification-motivated early exercise on the value of employee options. This example also highlights some of the pitfalls that may arise when these options are valued using their expected terms in a conventional valuation model. Finally, we outline how our valuation approach can be implemented in a formal model.

An Example

Most of the issues governing early exercise of employee stock options can be illustrated using a simple two-period binomial model. Consider an employee with two sources of wealth: a single stock option and other financial assets, assumed to be invested in risk-free assets with an interest rate of zero. The option may not be sold, but it may be exercised early and the proceeds placed in the risk-free portfolio.[6]

Figure 1 depicts the possible evolution of stock price over two periods. In each period, the price can rise by 20 percent with probability 0.55 or fall by 20 percent with probability 0.45. The current stock price is \$100, and the exercise price of the employee stock option also is \$100. The option has a two-period maturity. Because the stock pays no dividends, an unconstrained investor who could sell the option would never choose to exercise it early; the option could always be sold for more than its intrinsic value. Early exercise may be rational, however, when the employee stock option cannot be sold.

At the maturity of the option, the exercise rule

Figure 1. Stock Price Evolution, Two-Period Example

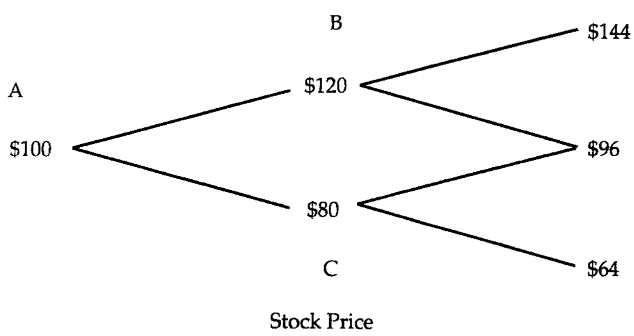

is trivial: Exercise if the option is in the money. Exercise in the Figure 1 example would occur only at the upper node, where the price equals \$144. One period earlier, however, the exercise rule is more complex. At Node B, where the stock price is \$120, the investor may exercise the option and receive a payoff of \$120 − \$100 = \$20 or may choose to hold the option until the next period. If the option is exercised early, the proceeds are added to the risk-free portfolio and grow until the option maturity date.[7] If the option is held, its payoff in the next period will be either \$44 (if the stock's price rises to \$144) or zero (if the price falls to \$96). Because the probability of a stock price increase is 0.55, the expected payoff in the next period is \$24.20, greater than the exercise value of the option at Node B.

The exercise rule at Node B will depend on the investor's risk tolerance. Clearly, a risk-neutral investor will choose to hold the option until the next period. Early exercise will lead to final wealth of \$20 (because the interest rate at which the proceeds of early exercise can be reinvested is zero), and the expected payoff of the option next period is \$24.20. In contrast, a sufficiently risk-averse investor will choose to exercise early. It would not take much risk aversion to prefer a sure payoff of \$20 to a gamble with possible payoffs of \$44 or zero. The willingness to take the gamble also depends on the amount of nonoption wealth the employee holds. When the option payoff is a small part of total wealth, the option holder will be more willing to accept the gamble.

To examine the exercise choice formally, call the investor's utility of final wealth $U(W)$ and call nonoption wealth W_0. The choice to exercise at Node B is determined by comparing utility of final wealth given early exercise, $U(W_0 + 20)$, to expected utility of final wealth with no early exercise, $E[U(W_0 + x)]$, where x is a random variable that can take on a value of either 44 or 0. For example, if utility of final wealth exhibits constant relative risk aversion, γ, so that

$$U(W) = \frac{W^{1-\gamma}}{1-\gamma}, \tag{1}$$

then the investor will choose not to exercise at Node B if γ is close enough to zero (risk neutrality) but will exercise for higher values of γ. For example, if $\gamma = 2$ (which is broadly consistent with some empirical work[8]) and $W_0 = 30$, then maturity-date utility with exercise at Node B will be

$$U(W_0 + 20) = U(50) = -0.02,$$

and expected utility with no exercise will be

$$\begin{aligned} E[U(W_0 + x)] &= [0.45 \times U(30) + 0.55 \times U(74)] \\ &= [0.45 \times (-0.0333) + 0.55 \times (-0.0135)] \\ &= -0.0224. \end{aligned}$$

In this case, exercise would be optimal.[9] If W_0 is very high, say $W_0 = 200$, the uncertainty associated with the option is less compelling and expected utility is higher without early exercise. This example suggests that the tendency to exercise early will increase with both risk aversion and the fraction of wealth the option represents.[10]

The exercise decision at Node C is simple. Because the option is out of the money at that node, no one will exercise. Similarly, at Node A, the option is at the money, so no exercise would occur.

More generally, the exercise decision will be made according to the standard rules of dynamic programming under uncertainty. At each node, the investor will compare utility at the maturity date, given exercise, with expected utility, given no exercise, and choose the action with the higher expected utility.

Given the exercise policy dictated by the optimal program, the market value of the option can be determined using standard methods. The basic approach is to calculate the market value of the cash flow stream that the option will produce, given the possible stock price paths and conditioning on the exercise rules that will be followed.[11]

In the two-period binomial example, the market value of the option can be found by calculating the present value of the option payoffs using "risk-adjusted probabilities" for the stock price process. In our example, the price can increase by 20 percent or fall by 20 percent; thus,

Amount of up-movement		$u = 1.2$
Amount of down-movement		$d = 0.8$
1 + risk-free rate		$R = 1.0$

The risk-adjusted probability of an up-movement is

$$p = \frac{(R - d)}{(u - d)} = \frac{0.2}{0.4} = 0.5.$$

The value tree for the option if it could be sold freely is given in the top panel of Figure 2. The value at each node is simply the discounted expected value of the option across the two nodes to which the stock price can jump in the following period, calculated using the risk-adjusted probability of 0.5 that the stock price will increase and the assumed risk-free interest rate of zero.[12] Because the \$22 option value at Node B is greater than \$20, an unconstrained investor who wishes to unload the option position would choose to sell it rather than to exercise at this node.

Figure 2. Option Valuation, Freely Traded and Nontradable Options, Two-Period Example

American Call, $X = \$100$ and $r_f = 0$

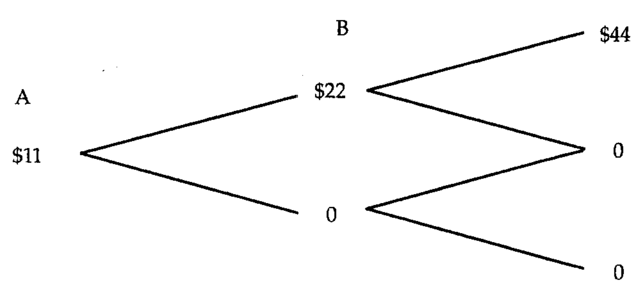

Employee Stock Option, $X = \$100$ and $r_f = 0$

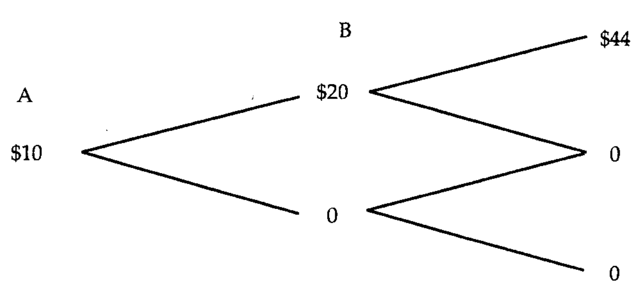

The lower panel of Figure 2 is the value tree for the option assuming that it is currently held by an employee who must exercise to unload the position. We assume that risk aversion is sufficient to induce early exercise at Node B. The value of the option at that node is \$20. As a result of this "suboptimal" exercise, the current value of the option (at Node A) also is reduced. The value at Node A is the probability-weighted average (using risk-neutral probabilities) of \$20 and \$0 rather than the probability-weighted average of \$22 and \$0 dollars.

To the extent that employees exercise suboptimally, the firm's contingent liability is not as valuable as one would predict from an option pricing model predicated on unconstrained rational behavior. In this example, risk aversion reduces the value of the firm's obligation and the value of the employee's compensation from a potential value of \$11 to an actual value of \$10. Employee risk aversion affects the market value of the options through the effect on early exercise.

Average Versus Effective Maturity

Historical data rather than a maximizing model might appear to be an easier way to account for the effect of early exercise of employee stock options. Unfortunately, historical data are of limited use for this purpose. Refer again to Figure 1 and note that, although risk-averse employees would exercise at Node B, no one would exercise at Node C because the option is out of the money at that point. As stock prices increase, the fraction of total wealth represented by the option holdings will rise as well, which in turn will engender higher rates of early exercise as diversification motives become more telling.

This simple example illustrates that the incidence of early exercise will depend on the path of stock prices subsequent to the grant of the options. Therefore, historical rates of early exercise will be of little value in forecasting future exercise rates. Historical rates are tied to stock market performance during the relevant periods and are not necessarily predictive of future rates.

The binomial model can be used to calculate the *ex ante* expected time until exercise, given the probability distribution of future stock prices. For example, the option in Figure 2 would be exercised at the end of the first period if the stock price increases in that period, but the option will be held until maturity if the stock price decreases. Therefore, the expected term of the option is E(Term) = 0.55×1 period $+ 0.45 \times 2$ periods $= 1.45$ periods.

More generally, with a richer stock price tree, we could take a probability-weighted average of the exercise date across the entire range of possible dates. This approach results in an expected time until exercise consistent with the full probability distribution of stock prices. Unlike historical-average exercise rates, the result will not be affected by the incidence of bull or bear markets in the sample period.

Even this simple example reveals the potential bias in FASB's proposal to use expected term in place of stated term. Although the expected term is 1.45 periods, the effective term of the option is better characterized as only one period. By this, we mean that a binomial model solved assuming the true option maturity is only one period would result in precisely the same estimate of option value as the actual value obtained using the true two-period term but allowing for early exercise. The discrepancy between effective and expected term arises because in the states for which the option is kept alive for its two-period life, the option payoff is zero. The extra life of the option in these states contributes nothing to value.

Such an extreme result is an artifact of the two-period example, but the qualitative result carries forward to more realistic cases. Because options will not be exercised when payoffs are low or zero, expected time to maturity is less than effective maturity. For this reason, use of the expected option term can impart substantial upward bias to the estimate of option value.

A Formal Model

The two-period example is a simple illustration of a dynamic programming approach to option valuation. To allow for realistic stock price dynamics and option exercise decisions in our actual valuation model, we divided the option life into 1,000 periods. The investor at each date, t, chooses whether to exercise or continue to hold the option so as to maximize the expected value of utility at the option maturity date, T.[13] Wealth at the option maturity date, W_T, will depend on both the stock price, S_T, and the investor's nonoption wealth; utility at this date is denoted $U(W_T)$. If nonoption wealth at times prior to T is denoted F_t, then the expected value of an investor's time-T utility can be written as a function of the stock price and nonoption wealth; that is, $J(S_t, F_t)$. Therefore, the exercise decision at t is made to maximize $J(S_t, F_t)$.

An investor who exercises early and places the proceeds in risk-free assets will arrive at time T with certain wealth:

$$W^* = [S_t - X + F_t]e^{r(T-t)},$$

where X is the strike price of the option and r is the risk-free interest rate. Once the investor exercises, time-T utility is certain and will equal $U(W^*)$. If instead, the investor foregoes early exercise and holds the option for an additional time interval of Δt, then expected utility $J(S_t, F_t)$ can be derived from the recursive formula

$$J(S_t, F_t) = E[J(S_{t+\Delta t}, e^{r\Delta t}F_t)],$$

where $E(.)$ is the expectations operator. The boundary condition for this recursive relationship is that, at option maturity,

$$J(S_T, F_T) = U[\max(S_T - X, 0) + F_T].$$

Therefore, at each date, the investor exercises if the certain utility that results from such exercise is greater than the expected utility from continuing to hold the option for another small time interval.

Notice that we use true probability distribution of stock price evolution, not risk-neutral probabilities, in calculating the expected utility conditional on not exercising. Therefore, the expected rate of return on the stock, and not just the risk-free rate, will determine exercise policy and hence the value of the option.[14]

VALUATION RESULTS

In this section, we examine the impact of early exercise using realistic parameters. In the base case, the life of the option is ten years, the short-term risk-free interest rate is 2.5 percent, the expected rate of return on the stock is 10 percent, and the stock price return has an annual standard deviation of 35 percent.[15] We assume initially that the stock pays no dividends. This assumption simplifies the analysis of the impact of early exercise on the value of nonmarketable employee stock options, because the benchmark unrestricted stock option would never be exercised early.

Figure 3 illustrates the effect of risk aversion and nonoption wealth on the value of an employee stock option. We normalized by assuming that the employee is granted one option with an exercise price of \$1 on a stock initially selling for \$1 a share. The value of the option (given the employee's exercise policy) is plotted as a function of the risk-aversion parameter, γ. Each curve in Figure 3 was plotted for a different level of nonoption wealth.[16] Because of the normalization, this value can be interpreted as the ratio of nonoption wealth to the current value of the stock for which the option can

Figure 3. Sensitivity of Option Value to Risk Aversion

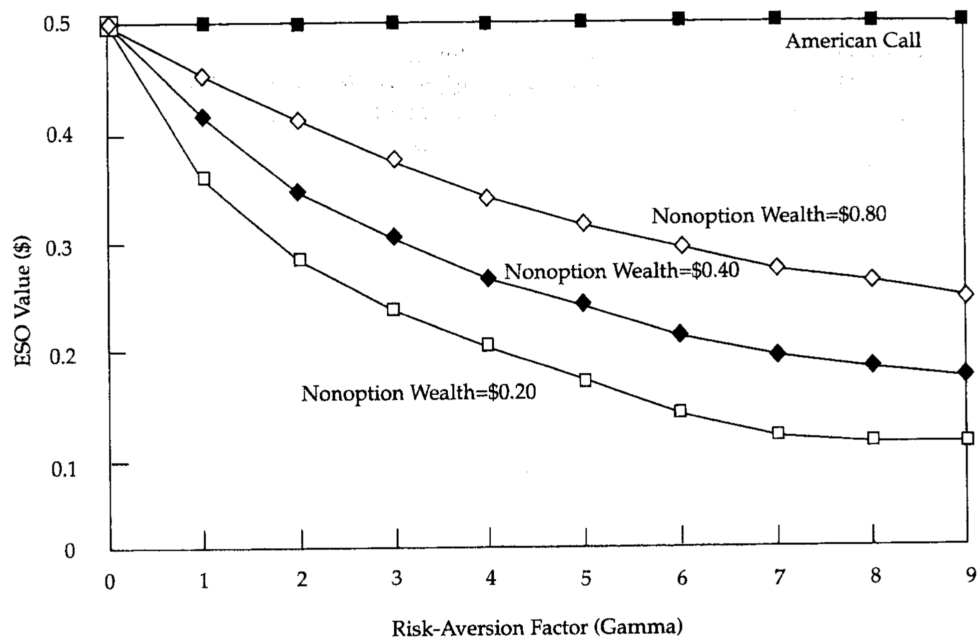

be exercised. The option value is a function of this multiple because of its effect on exercise policy.

The vertical intercept in Figure 3 is the value of the employee stock option to a risk-neutral investor. Risk-neutral employees would never exercise early for diversification reasons. For these employees, the option value equals the value obtained from an unconstrained option valuation model. The vertical intercept, therefore, provides a benchmark equal to the value the option would have if it were freely marketable. The actual option-value profiles all lie below this benchmark because of the impact of early exercise. All option-value curves radiate from the same vertical intercept because for risk-neutral employees, the option value equals the unconstrained benchmark regardless of nonoption wealth.

Figure 3 demonstrates that as risk aversion rises, the value of the option falls rapidly, because more risk-averse employees will exercise more readily than others. For even moderately risk-averse employees (say with $\gamma = 3$) with nonoption wealth of 40 cents, the option is worth 30 cents, which is about 60 percent of the unconstrained binomial model value of the option (the top line in Figure 3).

As nonoption wealth increases, the options tend to be worth more. The higher profiles in Figure 3 correspond to the higher values of nonoption wealth. If the option constitutes a small share of overall wealth, employees will be less sensitive to the option's risk. Indeed, if nonoption wealth were extremely large, the value curve for the restricted option would flatten out at the value of an otherwise identical unrestricted option, indicating that premature exercise would all but disappear.

Figure 4 plots the value of employee stock options for a risk-aversion level of 2 and nonoption wealth of 40 cents, as well as the value of unconstrained American options, as a function of the dividend rate on the firm's stock. Although both value profiles are downward sloping, the relationship between the values is fairly stable. The employee option is worth about 75 percent of the unconstrained option for dividend yields in the neighborhood of 2.5 percent.

Early exercise of employee stock options gives rise to an interesting pricing anomaly. Although the values of tradable options rise with the volatility of the returns on the underlying stock, the values of restricted employee options can actually fall as volatility increases. When volatility increases, a risk-averse option holder is more apt to exercise early, which reduces the value of the option. This effect may dominate the tendency that volatility otherwise would have to increase the option value. Figure 5 shows the effect of return

Figure 4. Sensitivity of Option Value to Dividend Rate

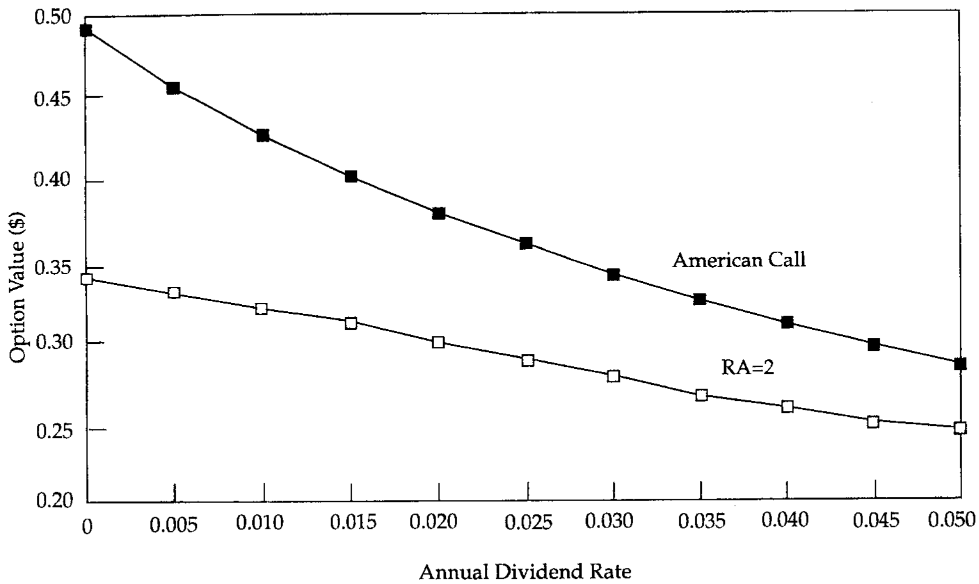

volatility on the value of employee stock options. For levels of risk aversion above 2, beyond some threshold level of volatility, the option value is a decreasing function of volatility.

This point is important for investors and FASB to consider when employee options are valued, because it shows that stock volatility and the expected term of the option cannot be chosen independently. In particular, consider the inconsistency that might arise if a firm uses historical data to estimate the expected term of the option and then independently forecasts stock price volatility. The historical exercise experience will be a poor guide to future experience if the firm expects volatility to change either for exogenous reasons or because of changes in the firm's business policies.

Figure 5. Sensitivity of Option Value to Volatility

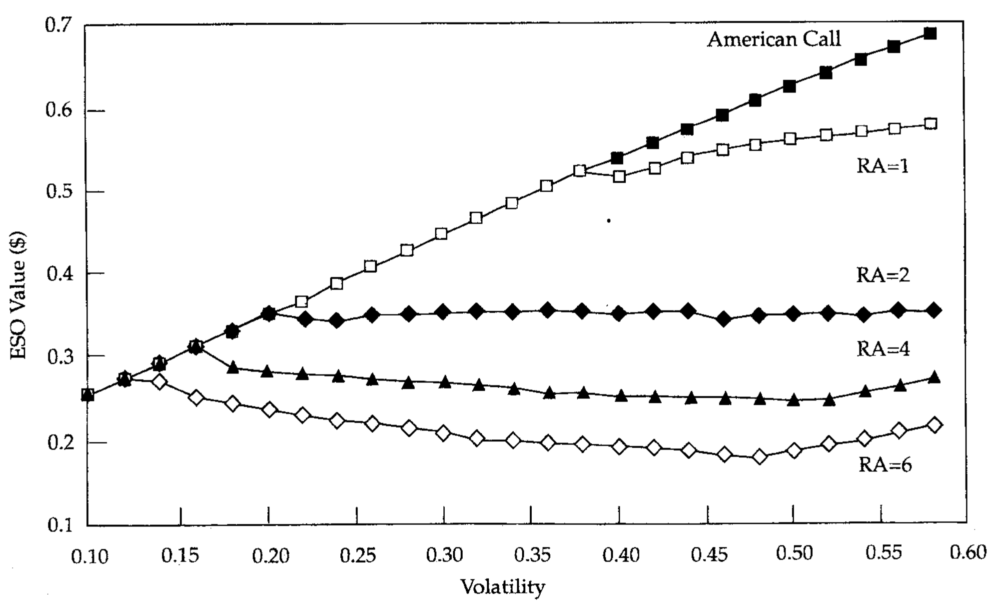

A conventional model calculated using any fixed expected term will imply that value increases with volatility. In reality, however, volatility and expected term will interplay in such a way that value eventually must fall with volatility.

In fact, it is easy to derive cases for which the employee stock option has a lower value than the "minimum value" of the option, defined as the maximum of zero or the stock price minus the present value of the exercise price and dividends to be paid prior to option maturity. At high levels of volatility and high risk aversion, the option is exercised so readily that its value is reduced below this benchmark. For example, with a risk-free rate of 2.5 percent and an option term of ten years, the minimum value of the option (on a non-dividend-paying stock) in our example is $\$1 - \$1/(1.025)^{10} =$ 0.219 cents. Figure 5, however, shows that for sufficiently high values of risk aversion (e.g., > 7), the option value drops below 20 cents for stock volatilities above 40 percent. This happens because risk-averse investors exercise their options on these very volatile stocks so readily that effective terms are close to zero.

Another important difference in the pricing of employee versus tradable stock options is that the value of employee options will depend on the expected rate of return of the underlying stock. For a given level of price volatility, employees will be less apt to exercise options on stocks with higher expected returns, and therefore such options will have a higher market valuation.[17] Figure 6 demonstrates that this effect can be economically significant.

Just as the value of an option is sensitive to risk aversion, so is the expected term of the option. The expected term can be derived from the binomial model by averaging the exercise date across all possible price paths. Table 1 presents the expected term for risk aversion of 2 and various values of nonoption wealth. As nonoption wealth becomes very large, the expected term approaches the stated term of the option—ten years—because diversification-motivated early exercise diminishes. At low levels of nonoption wealth, expected option life falls below four years.

Table 1. Expected Option Term as a Function of Nonoption Wealth

Expected Term (years)	Nonoption Wealth ($)
3.41	0.25
5.66	0.50
6.57	1.00
7.86	5.00
10.00	10.00

Note: Relative risk aversion = 2.

EXTENSIONS

In this section, we consider several dimensions along which the valuation model can be extended.

Figure 6. Sensitivity of Option Value to Return Premium

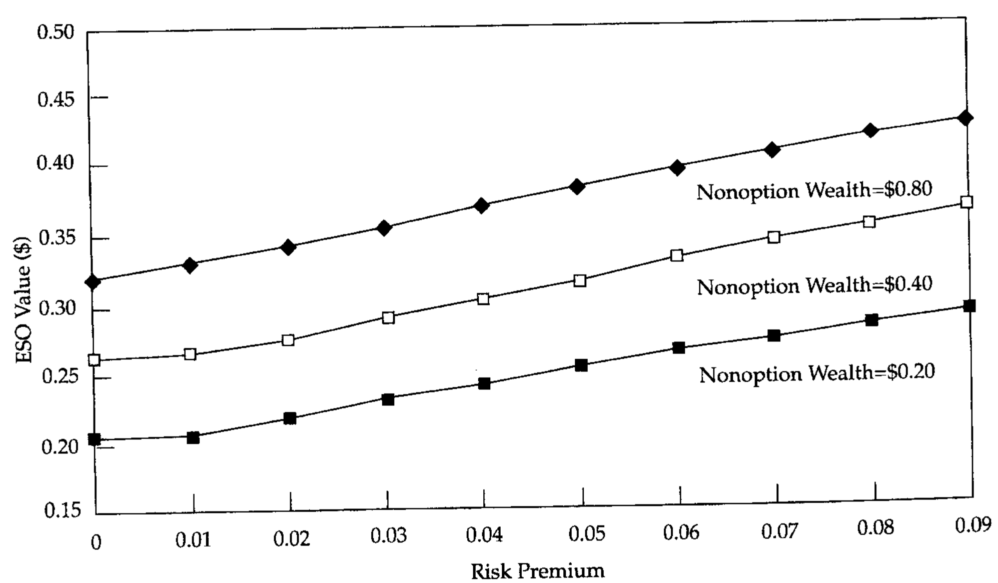

We treat the effects of vesting requirements, exogenous exercise, and partial exercise of the options.

Vesting and Exogenous Exercise

In our model, option exercise is always rational, at least given the constraints imposed by nonmarketability. In practice, some vested options may be exercised for exogenous reasons—because an employee leaves the firm, for example. Such exercise will reduce the value of employee stock options even farther below the unconstrained value.

The nonsystematic uncertainty surrounding exogenous exercise is easy to treat within our valuation framework. First, value the option using the dynamic programming approach outlined above, but replace the actual option maturity with a possible exogenous termination date. Calculate the option value in this manner for each possible termination date. Then, compute the weighted average of these option values using actuarial probabilities of exogenous termination at each date.

Forfeiture of nonvested options can be treated similarly. If employees who leave the firm forfeit their options, simply set the option value for termination dates prior to the vesting date equal to zero when calculating the average option value across possible termination dates.

To the extent that vesting requirements prohibit early exercise of options that eventually will become vested, they increase value. For example, a typical vesting period of three years means that even highly risk-averse employees cannot exercise the option before that date. Figure 7 shows that the value of the employee stock option increases as the vesting period increases. In the limit, as the vesting period approaches the stated life of the option, the option value approaches the value of a tradable option, at least in the case that the stock pays no dividends and early exercise is not optimal. This property leads to another paradox compared with the normal properties of tradable options: A European option that rules out early exercise will be worth more than an American one.[18]

Partial Exercise

The valuation model treats exercise as an all-or-nothing decision; it does not allow for exercise of a fraction of an employee's options. Partial exercise leads to a path-dependency problem, because the number of options held at any date can depend on the path of the stock's price prior to that date. For example, consider an employee stock option issued when the stock is selling at $50. Suppose the price first rises to $80 but has now declined to $55. In this case, the employee might have few options left in the portfolio because most would have been exercised at the earlier date when the stock price surged to $80. Given the relatively few options left in the portfolio, exercise might not be optimal.

Figure 7. Sensitivity of Option Value to Vesting Period Restriction

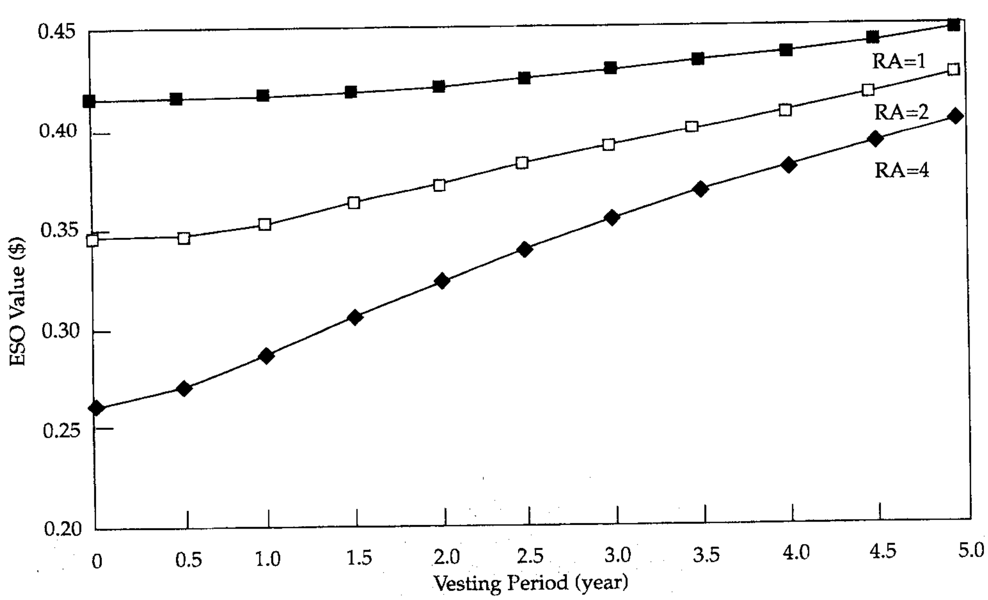

In contrast, if the price first fell to $40 and has now recovered to $55, the employee still would have most or all of the options originally granted because there would have been little opportunity to exercise until this date. Given the greater number of options still held, exercise of some options would be more likely. Clearly, exercise policy will depend on the path of stock prices.

What is not obvious is whether all-or-nothing exercise leads to overstatement or understatement of the effect of early exercise on value. If partial exercise is allowed, an employee will be more apt to exercise at least part of the stock options earlier than if all the options must be exercised. For example, under partial exercise, an employee might exercise 20 percent of the granted options in each of Years 6 through 10. Under all-or-nothing exercise, the employee might exercise all the options in Year 7. Which strategy will result in greater market value is not a priori obvious.

CONCLUSION

Because they cannot be traded, employee stock options may be exercised earlier than conventional option valuation models would predict. This early exercise reduces the value of the options. Early exercise will be more pronounced and the value of the options will be lower when employees are more risk averse and when a larger fraction of total compensation is in the form of options. The quantitative significance of early exercise can be substantial. For reasonable parameters, early exercise reduces the expected time until exercise and the option value by more than half.

Can a model such as the one described in this article actually be used to value employee stock options? Probably not. The model would require as inputs variables that are difficult to measure or observe—employee risk aversion, for example, and nonoption wealth invested in the firm as well as in other assets. Moreover, as Mark Rubinstein has pointed out, even small changes in the "usual" inputs into an option valuation model, such as stock volatility or dividend yield, would have major impacts on the values of these long-lived options.[19] Nevertheless, the model demonstrates clearly that option values will be sensitive to such variables.

Our conclusion is that employee stock options will need to be valued in a manner similar to that used for mortgage-backed securities. Instead of valuing each mortgage in a pool separately—which would require detailed knowledge about the homeowner's age, mobility, and so forth—mortgage-pricing models take a statistical approach in which the characteristics of the mortgage pool guide the prediction of prepayment rates in various economic environments. Similarly, a pool of employee stock options can be subjected to statistical analysis and valued using Monte Carlo simulations like those used widely in the mortgage industry.

A natural pool of options would be the employees at one firm. Liquidity-motivated early exercise of employee options is analogous to mortgage prepayment arising from largely exogenous factors such as home changes (as opposed to interest-rate-driven refinancings). An early-exercise function could be estimated from historical experience, as mortgage prepayment functions are now estimated. Early-exercise behavior would be specified to depend on stock price movements, as well as the demographic characteristics of the employees in the pool. Given an early-exercise function, employee options could be valued by simulating the early-exercise experience across several hundred or thousand stock price paths, just as mortgage analysts examine prepayment experience across interest rate paths.

As experience with stock option programs accumulates, FASB might consider a "generic" model designed to enhance comparability across firms. Such a model would take as inputs demographic features of an employee group, as well as features of the firm and its option program, such as stock price volatility. Most firms could then use this standard model to value the option program. This procedure would not eliminate the problems highlighted by Rubinstein but would enhance comparability across firms.[20]

FOOTNOTES

1. See Financial Accounting Standards Board, "Proposed Statement of Financial Accounting Standards: Accounting for Stock-based Compensation," *Financial Accounting Series No. 127-C* (June 30, 1993).
2. According to studies by Bear Stearns & Co. and R.G. Associates, quoted in L. Berton, "New Payroll Rule May Eat into 1997 Profits," *Wall Street Journal* (June 3, 1993):C1.
3. See, for example, T.W. Foster, P.R. Koogler, and D. Vick-

rey, "Valuation of Executive Stock Options and the FASB Proposal," *The Accounting Review*, vol. 66, no. 3 (July 1991):595–610, and the references cited therein.

4. This point is emphasized in M. Rubinstein, "On the Accounting Valuation of Employee Stock Options" (Berkeley, California: University of California, November 1994).
5. Employees cannot shed this risk by shorting stock in the firm. Aside from the reputational cost that such short sales would entail, such "short-swing profits" would be considered a violation of Section 16-b of the 1934 Securities Act. Any profits generated by short selling stock in the employee's firm that are realized within a six-month period must be returned to the firm.
6. Employee stock options are in fact warrants, not options. We will ignore the difference in this paper. The potential dilution from such options would typically be so small that option and warrant models would result in virtually identical valuations.
7. By carrying all wealth forward until the option maturity date, we abstract from consumption until after the option matures. This simplification is innocuous as long as employees' consumption programs are not constrained by a lack of marketable financial wealth (as opposed to human capital) before the options mature. Of course, if such constraints were binding, the incentive to exercise early would be even greater and the options would be worth even less than the values derived here.
8. See, for example, I. Friend and M. Blume, "The Demand for Risky Assets," *American Economic Review*, vol. 65, no. 5 (December 1974):900–23; or S.J. Grossman and R.J. Shiller, "The Determinants of the Variability of Stock Market Prices," *American Economic Review*. vol. 71, no. 2 (May 1981): 222–27.
9. An unconstrained investor with the same risk aversion would not exercise the option, however. Instead, such an investor would choose to sell the option.
10. The fraction of the employee's total wealth invested in the option should be interpreted broadly to include other forms of wealth tied to the fortunes of the firm, such as stock grants or human capital. Any dependence on the firm will result in a diversification motive to exercise the options early.
11. A standard reference for the binomial model is J.C. Cox and M. Rubinstein, *Options Markets*, (Englewood Cliffs, N.J.: Prentice-Hall, 1985).
12. Notice that we used the true probabilities of a stock-price increase to solve for the investor's decision to exercise or hold the option. The investor maximizes expected utility in a nonmarket setting. To find the market value of the option, however, we calculate the expected cash flows that will accrue to the option holder using risk-neutral probabilities. The reason is that the market value of the prospective stream of cash flows—the value the firm's shareholders will place on the options granted—is determined by unconstrained investors. The use of risk-neutral probabilities may be viewed as a means of obtaining the certainty-equivalent values of the cash flows.
13. We continue to assume that the utility function for end-of-period wealth exhibits constant relative risk aversion, γ, as specified in Equation 1.
14. See footnotes 12 and 15 and Figure 6.
15. The actual rate of return on a stock usually is unimportant in option valuation. The actual return affects exercise propensities, however, and therefore does affect the value of employee stock options. For example, in the two-period example above, the true probability of a stock price increase affected the exercise decision. A model similar to ours is developed independently in S. Huddart, "Employee Stock Options," *Journal of Accounting and Economics*, forthcoming. Huddart failed to recognize this point and solved for the value of the option under the assumption that the rate of return on the stock equals the risk-free interest rate. This assumption is unnecessary and overly restrictive. Moreover, as we demonstrate in Figure 6, the dependence of the value of the option on the risk premium of the stock is economically significant.
16. To the extent that the employee's nonoption wealth actually is held in restricted stock of the firm, the diversification problem is even more severe and the option will be exercised even more readily than in our simulations. If the employee's portfolio is broadly diversified, however, and can be adjusted to mitigate the exposure to the firm, the diversification problem is less severe. Even with a large portfolio, the employee can hedge out only the component of firm risk that is correlated with other stocks in the portfolio. Thus, although "market risk" can be hedged, "firm-specific risk" cannot.
17. Of course, higher expected return implies higher systematic risk. Figure 6 is drawn on the assumption that total risk is independent of expected return, so implicitly, the comparative-static assumption is that firm-specific risk adjusts to maintain total risk unchanged.
18. Vesting restrictions reduce the subjective value that an employee places on the option because the employee is made worse off by the inability to exercise and diversify. The market value of the contingent claim against the firm increases, however, because the vesting restriction forces the employee to follow the no-exercise strategy that willingly would be adopted by traders who could sell the option.
19. Rubinstein, "On the Accounting Valuation of Stock Options."
20. We are grateful to Pakorn Peetathawatchai for able research assistance.