Estimation of Employee Stock Option Exercise Rates and Firm Cost: Methodology[1]

Jennifer N. Carpenter	Richard Stanton	Nancy Wallace
New York University	U.C. Berkeley	U.C. Berkeley

December 15, 2006

[1]Corresponding author: Richard Stanton, Haas School of Business, 545 Student Services Building #1900, University of California, Berkeley, CA 94720-1900, or via email: stanton@haas.berkeley.edu.

Estimation of Employee Stock Option Exercise Rates and Firm Cost: Methodology

Abstract

Investors have become increasingly concerned about the cost of executive stock options to shareholders. Because executives face hedging constraints, standard option theory does not apply. The valuation problem reduces to accurately characterizing the option payoff. This paper develops a methodology for estimating option exercise and cancellation rates as a function of the stock price path, time to expiration, and firm and option holder characteristics. Our estimation accounts for correlation between exercises by the same executive. Valuation proceeds by using the estimated exercise rate function to describe the option's expected payoff along each stock price path and then computing the present value of the payoff. The estimation of empirical exercise rates also allows us to test the predictions of theoretical models of option exercise behavior. The paper not only illustrates an ideal valuation method for a large dataset, but also shows how to evaluate the usefulness of some of the approximations proposed in the literature.

JEL classification: G14.

With the explosive growth of employee stock options in corporate compensation, investors, auditors, and regulators have become increasingly concerned about the cost of these options to shareholders. Recent regulation requiring firms to recognize option cost has intensified the demand for suitable valuation methods. The difficulty is that these are long-lived American options, so their value depends crucially on how employees exercise them. Yet, because employees face hedging constraints, standard option theory does not apply. For example, contrary to the predictions of standard theory, the vast majority of options on non-dividend paying stock are exercised well before expiration because option holders cannot sell the options, and this significantly reduces their present value.

In standard theory for an ordinary American call, the holder can sell the option at any time, so it is reasonable to assume that he times the option exercise in order to maximize the option's present value. The present value-maximizing exercise policy and its implications for option value are well-researched (see for example, Black and Scholes (1973), Merton (1973), or Kim (1990)).

By contrast, the holder of an executive stock option must bear the risk of the option payoff, so simply maximizing the option's present value is generally not optimal. Indeed, evidence indicates that executives systematically exercise options on non-dividend paying stocks well before expiration. These early exercises significantly reduce option value.

In principle, as long as the exercise decision generates an option payoff that is subject only to hedgeable risks, such as stock price risk, and diversifiable risks, such as uncertainties that are idiosyncratic across executives, then unconstrained market participants can replicate those payoffs and so pricing by no arbitrage is still possible. The option valuation problem then reduces to accurately characterizing the option payoffs, that is, the exercise policies of executives.

One approach that has been taken in the literature is to model the exercise decision theoretically. The executive presumably chooses an option exercise policy as part of a greater utility maximization problem that includes other decisions, such as portfolio and consumption choice and managerial strategy. The utility-maximizing exercise policy may call for early exercise for diversification purposes or forfeiture if employment terminates when the option is unvested or out of the money. Properly valuing the option from the viewpoint of shareholders requires taking these possibilities into account. Papers that develop utility-maximizing models and then calculate the implied cost of options to shareholders include Huddart (1994), Carpenter (1998), Detemple and Sundaresan (1999), Ingersoll (2006), and Carpenter et al. (2005). Another approach is to estimate the exercise policy empirically. Papers such as Carpenter (1998) and Bettis et al. (2005) calibrate utility-maximizing models to mean exercise times and stock prices in the data, and then infer option value. However, these paper provide no formal estimation and the approach relies on the validity of the utility-maximizing models used. Huddart and Lang (1996) and Heath et al. (1999) provide more flexible empirical descriptions of option exercise patterns but do not go as far as option valuation.

This paper develops a methodology for formally estimating option exercise and cancellation rates as a function of the underlying stock price path, time remaining to expiration, and firm and option holder characteristics, using a sample of CEO and top executive option exercises from 100 firms from 1990 to 2002. Theory suggests that, in addition to stock return volatility and dividend rate, variables such as executive wealth, risk aversion, and holdings of restricted stock and other options should affect exercise decisions. Our hand-constructed database carefully tracks executive stock and option grants, exercises, and salary to provide information about portfolio position and wealth over time. We will distinguish hedgeable stock price risk from unhedgeable stock price risk because they may have different effects on the executive's exercise decision. We will also examine the explanatory power of psychological variables such as functions of the stock price path that have been found significant in Heath et al. (1999).

Valuation of an option with a given set of characteristics proceeds by using the estimated hazard function to describe the option's expected payoff along each stock price path and then computing the present value of the payoff, that is, its risk-neutral expected discounted payoff, in the usual way. This approach is similar to the prepayment modeling and valuation methods developed for mortgage-backed securities (see, for example, Schwartz and Torous (1989), and Stanton (1995)).

The estimation of empirical exercise rates also allows us to test the predictions of theoretical models of option exercise behavior. For example, do larger holdings of restricted stock speed exercise? We can also weigh in on open questions such as whether increased stock volatility speeds or delays exercise and how that affects option value.

Our valuation method works best with large samples that include a wide variety of stock price paths. The Society of Actuaries has established a task force on executive stock options to begin collecting data on exercises at thousands of firms. Once the data are in place, an actuarial science for valuing compensatory stock options can develop, similar to that for pension liabilities. Our paper develops a valuation methodology that would be ideal for such a dataset.

A number of analytic methods for approximating executive stock option value have been proposed in the literature. The FASB currently permits using the Black-Scholes formula with the expiration date replaced by the option's expected life. Jennergren and Näslund (1993), Carr and Linetsky (2000)), and Cvitanić et al. (2004) derive analytic formulas for option value assuming exogenously specified exercise boundaries and hazard rates. Hull and White (2004) propose a model in which exercise occurs when the stock price reaches an exogenously specified multiple of the stock price and forfeiture occurs at an exogenous rate. Until the accuracy of these methods can be determined, the usefulness of these methods cannot be assessed. Our estimated exercise rate function not only permits us to deduce the correct option value, but also allows us to estimate the parameters, such as the option's expected life, we need to carry out a given approximation method. In this way, our paper not only illustrates an ideal valuation method for a large dataset, but also shows how to evaluate the usefulness of some of the shortcuts proposed in the literature.

1 Option Cost to the Firm

To motivate the importance of estimating exercise rates from data, we first describe their role in option valuation. Standard American option theory assumes the option holder can trade or hedge the option, so his exercise policy is to maximize the option's present value. Employee stock option holders can neither sell their options nor short sell the stock, so in order to monetize the value of the option they may be forced to exercise it in a valuedestroying way. For example, the vast majority of options on non-dividend paying stocks are exercised early. These early exercises may arise because the option holder needs liquidity or diversification away from stock price risk. In addition, the holder might be forced to exercise or forfeit the option if he leaves the firm. Accurate valuation of employee stock options must take the possibility of these early exercises and forfeitures into account.

Recent theoretical models of employee stock options derive the exercise policy as the solution to the employee's utility maximization in the presence of departure risk and constraints on sales of options and stock. These models are useful for predicting which variables are likely to affect exercise decisions and for illustrating implications for option cost, but these models are likely to be too simplified and inflexible to describe adequately the actual exercise patterns observed in the data.

If enough data are available, a better approach is to estimate the exercise policy, i.e., the option payoff, directly from data on option outcomes. To allow for decisions that may depend on more than just the stock price, we describe the exercise policy in terms of a hazard function that describes the conditional probability of option exercise or forfeiture at each time and state, and then estimate the hazard function from the data. The hazard function could accurately describe utility-maximizing policies, if those are generating the data. But it also allows for the possibility of more general forms of exercise behavior. The hazard function could depend not only on the stock price path and time remaining to expiration, but also on any other variables that affect the probability of an option exercise. These could include characteristics of the firm, such as its volatility, and characteristics of the option holder, such as wealth and portfolio holdings.

For an individual option, the hazard function describes the probability that the option is exercised at a given time and state, conditional on having survived to that point. If the event that the option is actually exercised is sufficiently independent across option holders with identical hazard functions, conditional on the given time and state, then in a large enough pool of such option holders, the hazard function describes the fraction of options remaining in the pool that actually exercise at that time and state. We assume that such diversification is possible, or, more generally, that the conditional variance in the number of options actually exercised or forfeited around the expected value is not a priced risk in the market, so that option valuation proceeds as if perfect diversification were possible.

For each possible stock price path, we use the hazard function to calculate the fraction of remaining options that get exercised at each point along that path in a hypothetical pool of identical options. In this way, the hazard function generates a stream of cash flows from the pool along every stock price path. The present value of these cash flows can be represented as their expected discounted payoff, using so-called "risk-neutral probabilities" to average across stock price paths. This present value, or replication cost, is the cost of the option to the firm.

Various methods for approximating option value have been proposed in the literature. One approximation method currently permitted by the FASB is to compute the Black Scholes value of the option using the option's expected life in place of its maturity. Our estimated hazard function can be used to compute option expected life, and thus the option's FASB value. We could then compare this with the correct option value and evaluate the usefulness of this approximation.

2 Data

Our data cover option grants, exercises, and terminations for the top five executives at 106 high technology firms. These include option grant dates, vesting periods, strike prices, exercise dates, executive names and rank. We also track the executives' salary, bonus, and holdings of restricted stock. The sample firms are all those listed in *Hoover's Guide to Computer Companies 1995* that were publicly traded in the US and operated primarily in the computer industry. Using annual proxy statements, we compile the history of the contractual structure of the executive compensation packages as far back as 1992 and as far forward as 2003. We get the data on option exercises from the SEC's *Transactions and Holdings Information*, which includes Form 4: *Statement of Changes of Beneficial Ownership of Securities* and Form 5: *Annual Statement of Beneficial Ownership of Securities.* These documents are obtained from Lancer Analytics, a subsidiary of Thompson Financial. Data on termination is also available from Hoover's. Tables 1 through 14 and Figures 1 through 7 describe the data on CEO salary and option grants and exercises from 1992 to 2002.

2.1 Illustrative Data Analysis for Selected High Tech Firms

We use two additional data sets to illustrate the applicability of fractional logit estimators for modeling ESO exercise transitions at daily frequencies.[1] The first data set includes the employee stock option grants and option exercise histories for all employees that work for a large publicly traded semiconductor firm, Semiconductor Firm A. The data set includes 55,523 ESO grants to about seventy five hundred employees over the period 1982 through 2005. About 83% of Semiconductor Firm A's ESOs were granted in 1997 (24.89% of total), 2000 (35.3% of total), and 2004 (23.2% of total). In all three of these years, the vesting structure is twenty five percent per year over four years and the option maturities are six years. There were no recorded exercises of options granted later than 1998, and about 96% of the employee exercise events involved the exercise of 1997 grants. For that reason, our illustrative discussion will focus on the 14,506 option exercise experiences of 1997 ESO grants that Semiconductor Firm A made to its employees.

We supplement the ESO performance data from Semiconductor Firm A with a data set developed by the authors that includes the option grants and exercise histories for 100 CEOs in the high technology computer industry (semiconductor, computer, and software/services firms) over the period 1992 through 2004 (about 1,600 grants). These data allow us to focus on the option exercise dynamics of large option positions and to contrast these with the performance of the smaller grants that appear in Semiconductor A data. The average number of options granted per grant was 4,058 in the data from Semiconductor Firm A, whereas, in the high tech CEO data, the average number of options granted per grant was 287,753.

We merge each grant to an appropriate path of stock prices and stock splits using the initial granting date, the vesting structure and the ESO's maturity. This merge generates a panel data set of daily time varying covariates for each grant. These covariates include the split-adjusted stock prices, exercise prices, outstanding vested options, and alternative

[1] We thank Terrence Adamson at AON Consulting for these data.

termination events (due to death, retirement, and separations among others). Our unit of analysis is the fraction of each grant exercised by each employee each period and our period is defined as a stock trading day. The resulting data sets are quite large.[2] We provide summary statistics for these grant-level panels in Tables 1 through 3 and Figures 1 though 8. The primary objective of these summaries is to illustrate the time series and cross-sectional structure of ESO performance data and to highlight the superiority of the fractional logit class of estimators over the more standard logit and proportional hazard estimators.

In Figures 1 and 2, we graph the experience paths for two different employees who received relatively small ESO grants from Semiconductor Firm A on October 27, 1997. As shown in Figure 1 and 2, the stock price and strike price paths are the same given the common grant date. The exercise experience and the consequent dynamics of the options outstanding are quite different for the two employees. The average fraction of vested options that were exercised by the employee shown in Figure 1 is 56%, whereas, as reported in Figure 2, the average fraction of vested options that were exercised in each of the five exercise events by the second employee is 71%. The employees also differ in average number of options exercised and the time elapsed between the first vesting period and the option exercise event. The Figure 1 employee always has vested options outstanding after the first vest through the last full exercise, however, Figure 2 employee frequently fully exercises each block of vested options and then uses fractional exercises of the outstanding stock later in the holding period. These different patterns of fractional exercise behavior and the path persistence of the behavior cannot be successfully handled in either a logit or proportional hazard framework, but they are estimable in a fractional logit framework.

Tables 1 through 3 summarize the exercise experience of all the exercised option grants for about 6,879 employees at Semiconductor Firm A. As reported in Table 1, the average run-up in the ratio of the stock price to the strike price was slightly more that 300%, although for employees that exercised at the height of the stock market peak the maximal increase was 1,304%. The average number of days from first vest was 958.78 and the percentage of vested options that were exercised in any one exercise event was 91%. The average number of options exercised was 2,361 at each exercise and the maximal is 1.6 million options.

Table 2 presents a summary of the time series of exercise events for the 1997 option grants awarded by Semiconductor Firm A. The average cumulative change in the price to strike ratio rises to a high as expected in 2000 at the stock market peak and then the ratio decreases. The average percentage of vested options exercised ranges from a high of 98% per exercise experience in 1998 and to a low of 89% per exercise experience in 2000 when most option holdings were fully exercised at each event.

In Table 3, we compare the characteristics of exercise events for large blocks of options to those for small blocks of options. We expect that the larger grants belong to senior managers, however, this expectation is not possible to verify because we lack information on individual employment status. In the upper panel of Table 3, we report the performance characteristics for ESO exercises of greater than 10,000 options and, in the lower panel, we report the characteristics of ESO exercises involving fewer than 1,000 options. We find that exercise of large block of options is usually not complete and ranges from 68% to 87% of the

[2]The are over twelve million records generated for the Semiconductor Firm A data and over five million records for the CEO data

vested holdings. The exercise of smaller blocks of vested options is usually associated with the full exercise of the outstanding stock of options and the percentages range from 79% to 99%. It is interesting to note that the highest exercise percentages are not associated with the 2000 stock price run-up.

These summary data indicate apparent differences in the exercise behavior of employees with large option holdings and employees that do not have such holdings. In addition, fractional exercise of vested options is a common feature of large exercise events. Since logit and hazard estimators do not control for the size of transition events these important features of ESO exercise decisions would be inaccurately modeled. The fractional logit specification, however, properly controls the relative fraction of the exercise events and the effects of important time varying covariates on the level of these exercises.

In Figure 3 and 4, we plot the ESO exercise experience for the CEO of Semiconductor Firm A. As shown, although the CEO has a large option position, he always exercises 100% of the outstanding vested shares when he exercises. Interestingly, the CEO appears highly time series consistent in his exercise decisions and exercises at about the same period in the option's life. A contrasting picture of a CEO's exercise strategy is plotted in Figure 5 and 6 for another semiconductor firm: Semiconductor Firm B. This CEO again exhibits time series consistent ESO exercise strategies and he prefers to initially exercise a small fraction of his vested options and then wait to fully exercise the remainder after about three years. Finally, in Figure 7 and 8, we present a more extreme form of fractional exercise by the CEO of a large publicly traded software firm: Software Firm. This CEO exercises small percentages of his total vested options toward the end of the ESO holding period. These exercises are substantially less than one in all but the last exercise which is 100%.

To conclude, our summaries of large panels of cross sections of ESO performance histories indicates that there is important heterogeneity in the exercise strategies of employees. Fractional exercises of vested option holdings is quite common among employees that are CEOs and for large option holders. The CEO of Semiconductor Firm A is somewhat outside this general conclusion, although he does exhibit strong time series consistency in his pattern of exercises before the option expiration date. In contrast, smaller vested option positions tend to wholly exercise. These findings strongly support the use of an estimator that controls for fractional exercise while explicitly handling the correlation between option exercises within and between different grants held by the same employee. The fractional logit estimator is suitable for cross sections of panel data, such as our grant-day performance data for ESOs, and allows for fixed effects for employees and controls for time series persistence. Neither of the standard alternative estimators, such as logit or proportional hazard models, allow for partial or fractional transitions.

3 Specification and Estimation of Exercise Behavior

As discussed above, we need to estimate the hazard rate governing the likelihood of a given option's being exercised at any given time and state. There is a large literature on hazard rates and their estimation, but primarily in settings where the occurrence of the event being modeled is uncorrelated across different individuals, once dependence on a set of underlying covariates has been taken into account.[3] This is very different from the setting we face.

As is clear from Figures 1–8, the exercise of one option in a given grant held by an individual is extremely highly correlated with the exercise of another option in the same grant held by the same individual. It is also highly (though less so) correlated with the exercise of options in other grants held by the same individual. Using standard econometric techniques to estimate hazard rates based on exercise at the individual option level would result in very misleading results. It would treat each option as an independent observation (thus adding new information to the estimation), whereas in fact each individual option is so highly correlated with other options that, by itself, it adds almost no additional information. As a result, the estimation would report wildly optimistic confidence intervals, suggesting a much more precise knowledge of the true hazard rate than we actually have.

This issue does arise in other settings, such as in modeling corporate bond default, where it has been observed that defaults cluster more than a pure hazard rate story would suggest. This implies that, after controlling for the hazard rate, the event of default is correlated across firms. One popular solution, when the number of firms involved is small, has been to use "copula functions", which explicitly model this correlation.[4] However, this is not feasible in our case, since the number of options (and hence the number of correlation coefficients to consider) is too high.

Instead of looking at individual options, we directly model the *fraction* of each grant exercised each period. Other authors have previously looked at fractions of grants exercised, including Heath et al. (1999), who regress the fraction of each grant exercised against various explanatory variables. However, this simple regression approach has several problems. First, it does not force the expected proportion of options exercised in a given month to be between zero and one, which causes problems when we try to use the parameters for valuation. Second, aggregating exercises across different individuals throws away a lot of potentially important information about the differences in exercise behavior across individuals.

In this paper, we model the fraction of each grant exercised by each holder each period, in a manner that allows us to obtain hazard rates that must be positive, while explicitly handling the correlation between option exercises within and between different grants held by the same individual.

3.1 The covariates

Even a parsimonious description of the exercise rate must allow for dependence on the level of the stock price and the time to expiration, since these variables determine whether or not the option is vested or in the money. A richer description of the hazard function would also allow the hazard rate to depend on other variables that theory predicts would be important, such as the wealth of the option holder and his holdings of other options and restricted stock, as well as the volatility of the underlying stock price and portion of risk that is hedgeable using other assets. It might also be of interest to include executive rank as an explanatory variable.

One of the most difficult problems is how to describe the portfolio of the option holder.

[3]For good introductions to hazard rate analysis, see Cox (1972) or Kalbfleisch and Prentice (1980).

[4]See, for example, Li (2001).

One issue is that theory would suggest that total wealth is what matters but only data on compensation is available. Salary may be a reasonable proxy for wealth. A more complicated problem is summarizing the stock and option holdings of the executive. Formal theory of optimal exercise of multiple executive stock option grants has not yet developed. Intuition suggests that the greater the option holder's total forced exposure to the stock risk, the greater the exercise rate. This might be measured as the dollar variance of total or unhedgeable stock risk represented by the stock and option positions, using Black-Scholes option deltas to quantify the size of those positions.

Another variable of importance is a measure of the attractiveness of exercising the given option in question. Some candidates include the market-to-strike-price ratio or the Heath et al. (1999) ratio of option exercise value to option value under a standard model. Heath et al. (1999) also find that so-called psychological variables such as recent lagged stock returns and stock price relative to its maximum have predictive power.

Estimation of a hazard function that depends on option holder characteristics such as portfolio holdings would allow for tests of theory about optimal executive stock option exercise. However, estimation with a dataset that only includes exercise data would require a more parsimonious specification of the hazard function that would depend only on the stock price process. In the more parsimonious specification, variation in the characteristics of the option holders is essentially treated as another diversifiable risk that can be priced at its expected value. In any case, we can use the estimation to address important open questions such as how stock return volatility affects exercise rates and option value.

3.2 Estimating the Model

This section describes the details of our estimation procedure, which is based on the fractional logistic approach of Papke and Wooldridge (1996). Let y_{ijt} be the fraction exercised at time t of grant j held by individual i, and assume that we can write

$$y_{ijt} = G(X_{ijt}\beta) + \epsilon_{ijt}, \tag{1}$$

where X_t is some set of covariates in I_t, the information set at date t, where G, the expected fraction exercised at date t, is a function satisfying $0 < G(z) < 1$, and where

$$\begin{aligned} E(\epsilon_{ijt} \mid I_t) &= 0, \\ E(\epsilon_{ijt}\,\epsilon_{i'j't'}) &= 0 \quad \text{if } i \neq i' \text{ or } t \neq t'. \end{aligned}$$

In application, we'll use the logistic function,

$$G(X_{ijt}\beta) = \frac{\exp(X_{ijt}\beta)}{1 + \exp(X_{ijt}\beta)}.$$

Note that, unlike a linear specification, the predicted fraction exercised each period must always be between 0 and 1. Note also that, while we are assuming the residuals ϵ_{ijt} are uncorrelated between individuals and across time periods, we are allowing for ϵ_{ijt} to be arbitrarily correlated between different grants held by the same individual at a given point in time, and we are not making any further assumptions about the exact distribution of ϵ_{ijt}, or even about its variance.[5]

As in Papke and Wooldridge (1996), we estimate the parameter vector β using quasi-maximum likelihood (see Gouriéroux et al. (1984)) with the Bernoulli log-likelihood function,

$$l_{ijt}(\beta) = y_{ijt} \log\left[G(X_{ijt}\beta)\right] + (1 - y_{ijt}) \log\left[1 - G(X_{ijt}\beta)\right]. \tag{2}$$

Estimation involves solving

$$\max_{\beta} \sum_{i,j,t} l_{ijt}(\beta).$$

The K first order conditions, corresponding to the K elements of β, are given by

$$\begin{aligned}
\sum_{i,j,t} \frac{dl_{ijt}(\beta)}{d\beta} &= \sum_{i,j,t} X_{ijt} \left[G'(X_{ijt}\beta)\left(\frac{y_{ijt}}{G(X_{ijt}\beta)} - \frac{1 - y_{ijt}}{1 - G(X_{ijt}\beta)}\right)\right] \\
&= \sum_{i,j,t} \left(\frac{X_{ijt}G'(X_{ijt}\beta)}{G(X_{ijt}\beta)(1 - G(X_{ijt}\beta))}\right)(y_{ijt} - G(X_{ijt}\beta)), \qquad (3) \\
&= 0.
\end{aligned}$$

Equation (1) implies (using iterated expectations) that the population expectation of these first order conditions is zero, hence this QML estimator, $\widehat{\beta}$, is a (consistent) GMM estimator of β, with no assumptions other than Equation (1). Following the notation in Papke and Wooldridge (1996), define the residual

$$\widehat{u}_{ijt} \equiv y_{ijt} - G(X_{ijt}\widehat{\beta}),$$

and define

$$\begin{aligned}
\widehat{G}_{ijt} &\equiv G(X_{ijt}\widehat{\beta}), \\
&\equiv \widehat{y}_{ijt}, \\
\widehat{g}_{ijt} &\equiv G'(X_{ijt}\widehat{\beta}).
\end{aligned}$$

With uncorrelated residuals, the asymptotic covariance matrix of $\widehat{\beta}$ takes the usual "sandwich" form (see Gouriéroux et al. (1984)),

$$\widehat{\mathbf{A}}^{-1}\widehat{\mathbf{B}}\widehat{\mathbf{A}}^{-1},$$

[5]In particular, unlike assuming (say) a beta distribution for y_{ijt} (see, for example, Mullahy (1990) or Ferrari and Cribari-Neto (2004)), we are allowing for a positive probability that y_{ijt} takes on the extreme values zero and one.

where

$$\begin{aligned}
\widehat{\mathbf{A}} &= \sum_{i,j,t} \frac{\partial^2 l_{ijt}(\widehat{\beta})}{\partial\beta\partial\beta'} \\
&= \sum_{i,j,t} \frac{\widehat{g}_{ijt}^2 X_{ijt} X_{ijt}'}{\widehat{G}_{ijt}(1-\widehat{G}_{ijt})}, &(4) \\
\widehat{\mathbf{B}} &= \sum_{i,j,t} \left(\frac{\partial l_{ijt}(\widehat{\beta})}{\partial\beta}\right)\left(\frac{\partial l_{ijt}(\widehat{\beta})}{\partial\beta'}\right) \\
&= \sum_{i,j,t} \frac{\widehat{u}_{ijt}^2 \widehat{g}_{ijt}^2 X_{ijt} X_{ijt}'}{\left[\widehat{G}_{ijt}(1-\widehat{G}_{ijt})\right]^2}. &(5)
\end{aligned}$$

If we allow for correlation between the residuals, assuming that

$$\mathrm{var}(u) = \Omega = \begin{pmatrix} \Sigma_1 & & \ldots & & 0 \\ & \ddots & & & \\ \vdots & & \Sigma_i & & \vdots \\ & & & \ddots & \\ 0 & & \ldots & & \Sigma_I \end{pmatrix},$$

where each Σ block corresponds to all of the option grants held by a given individual on a particular date, then $\widehat{\mathbf{B}}$ takes the slightly more complex form

$$\widehat{\mathbf{B}} = \widehat{S} = \left(\frac{\widehat{g}}{\widehat{G}(1-\widehat{G})} \otimes \mathbf{X}\right)' \widehat{\Omega} \left(\frac{\widehat{g}}{\widehat{G}(1-\widehat{G})} \otimes \mathbf{X}\right),$$

where $\widehat{g}/[\widehat{G}(1-\widehat{G})]$ is a vector containing the stacked values of $\widehat{g}_{ijt}/[\widehat{G}_{ijt}(1-\widehat{G}_{ijt})]$, $\mathbf{X}$ is a matrix containing all of the stacked X_{ijt} values,

$$\widehat{\Omega} = \begin{pmatrix} \widehat{\Sigma}_1 & & \ldots & & 0 \\ & \ddots & & & \\ \vdots & & \widehat{\Sigma}_i & & \vdots \\ & & & \ddots & \\ 0 & & \ldots & & \widehat{\Sigma}_I \end{pmatrix}$$

and

$$\widehat{\Sigma}_i = \widehat{u}_i \widehat{u}_i'.$$

This covariance matrix (which reduces to Equation (5) if Ω is diagonal) is robust both to arbitrary heteroscedasticity and to arbitrary correlation between the residuals in a given block. For further discussion of calculating standard errors in the presence of clustering, see Rogers (1993), Baum et al. (2003), Wooldridge (2003) and Petersen (2005).

4 Evaluating Alternative Valuation Methods

A number of simpler valuation methods that yield analytic expressions for option value have been proposed in the literature. For example, the FASB allows firms to value options at their Black-Scholes value using the option's expected life as the expiration date. Hull and White (2004) propose a model in which the option is exercised once the stock price rises above a prespecified multiple of the strike price.

We can also use the estimated hazard function to determine the option's expected life and then compute the FASB approximation for executive stock option value. Similarly, we can use the estimated hazard function to determine the expected level of the stock price multiple at exercise and then compute the option's Hull and White value. We can then compare these to the correct option value and evaluate the usefulness of these approximations.

5 Conclusions

This paper develops a methodology for estimating option exercise and cancellation rates as a function of the stock price path, time to expiration, and firm and option holder characteristics. Our estimation is based on a fractional logit approach and accounts for correlation between exercises by the same executive. Valuation proceeds by using the estimated exercise rate function to describe the option's expected payoff along each stock price path and then computing the present value of the payoff. The estimation of empirical exercise rates also allows us to test the predictions of theoretical models of option exercise behavior. The paper not only illustrates an ideal valuation method for a large dataset, but also shows how to evaluate the usefulness of some of the approximations proposed in the literature. Finally, the paper describes our sample of option exercises for the top five executives at 106 firms. We will apply this methodology using our sample of option exercises in future work.

References

Baum, C. F., M. E. Schaffer, and S. Stillman, 2003, Instrumental variables and GMM: Estimation and testing, *Stata Journal* 3, 1–31.

Bettis, J. C., J. M. Bizjak, and M. L. Lemmon, 2005, Exercise behavior, valuation, and the incentive effects of employee stock options, *Journal of Financial Economics* 76, 445–470.

Black, F., and M. S. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–659.

Carpenter, J. N., 1998, The exercise and valuation of executive stock options, *Journal of Financial Economics* 48, 127–158.

Carpenter, J., R. Stanton, and N. Wallace, 2005, Optimal exercise of executive stock options and implications for valuation, Working paper, New York University and U.C. Berkeley.

Carr, P., and V. Linetsky, 2000, The valuation of executive stock options in an intensity based framework, *European Financial Review* 4, 211–230.

Cox, D. R., 1972, Regression models and life tables (with discussion), *Journal of the Royal Statistical Society* 34, 187–220.

Cvitanić, J., Z. Wiener, and F. Zapatero, 2004, Analytic pricing of executive stock options, Working paper, University of Southern California.

Detemple, J., and S. Sundaresan, 1999, Nontraded asset valuation with portfolio constraints: A binomial approach, *Review of Financial Studies* 12, 835–872.

Ferrari, S. L., and F. Cribari-Neto, 2004, Beta regression for modeling rates and proportions, *Journal of Applied Statistics* 31, 799–815.

Gouriéroux, C., A. Monfort, and A. Trognon, 1984, Pseudo-maximum likelihood methods: Theory, *Econometrica* 52, 681–700.

Heath, C., S. Huddart, and M. Lang, 1999, Psychological factors and stock option exercise, *Quarterly Journal of Economics* 114, 601–627.

Huddart, S., and M. Lang, 1996, Employee stock option exercises: An empirical analysis, *Journal of Accounting and Economics* 21, 5–43.

Huddart, S., 1994, Employee stock options, *Journal of Accounting and Economics* 18, 207–231.

Hull, J., and A. White, 2004, How to value employee stock options, *Financial Analysts Journal* 60, 114–119.

Ingersoll, J., 2006, The subjective and objective evaluation of incentive stock options, Working paper, Yale University, forthcoming *Journal of Business*.

Jennergren, L. P., and B. Näslund, 1993, A comment on "Valuation of executive stock options and the FASB proposal", *Accounting Review* 68, 179–183.

Kalbfleisch, J. D., and R. L. Prentice, 1980, *The Statistical Analysis of Failure Time Data* (John Wiley, New York).

Kim, I. J., 1990, The analytic valuation of American options, *Review of Financial Studies* 3, 547–572.

Li, D. X., 2001, On default correlation: A copula function approach, Working paper, RiskMetrics Group.

Merton, R. C., 1973, Theory of rational option pricing, *Bell Journal of Economics and Management Science* 4, 141–183.

Mullahy, J., 1990, Regression models and transformations for beta-distributed outcomes, Working paper, Trinity College.

Papke, L. E., and J. M. Wooldridge, 1996, Econometric methods for fractional response variables with an application to 401(k) plan participation rates., *Journal of Applied Econometrics* 11, 619–632.

Petersen, M. A., 2005, Estimating standard errors in finance panel data sets: Comparing approaches, Working paper, Northwestern University.

Rogers, W., 1993, Regression standard errors in clustered samples, Stata Technical Bulletin 13, pp. 19–23.

Schwartz, E. S., and W. N. Torous, 1989, Prepayment and the valuation of mortgage-backed securities, *Journal of Finance* 44, 375–392.

Stanton, R., 1995, Rational prepayment and the value of mortgage-backed securities, *Review of Financial Studies* 8, 677–708.

Wooldridge, J. M., 2003, Cluster-sample methods in applied econometrics, *American Economic Review* 93, 133–138.

Table 1: Semiconductor Firm A: Summary Statistics for Employee Option Exercise Experiences

This table provides summary statistics for 14,506 option exercise experiences for options granted in 1997 to 6,879 employees in a large publicly traded semiconductor firm.

	Mean	Standard Dev.	Maximum	Minimum
Cumulative change in stock price/strike price	3.04	1.14	13.04	1.03
Number of days since first vest	958.78	631.21	2216.00	15.00
Percent of vested options exercised	0.91	0.21	1.00	.01
Options exercised	2361.48	21,175.55	1,600,000.00	4.00
Stock Price/Strike Price	3.04	1.14	10.04	1.03

Table 2: Semiconductor Firm A: Summary Statistics for Employee Option Exercise Experiences by Year of Exercise

This table provides summary statistics for 14,506 option exercise events for options granted in 1997 to 6,879 employees in a large publicly traded semiconductor firm. The exercise experiences are summarized for the year in which the options were exercised.

	Mean	Standard Dev.	Maximum	Minimum	Exercise Experiences
Cumulative change in stock price/strike price					
1998	1.39	0.13	1.96	1.03	218
1999	2.80	0.73	4.99	1.16	4334
2000	4.97	1.29	13.03	1.53	2268
2001	2.95	0.44	5.45	1.54	2973
2002	2.78	0.59	4.83	1.32	1524
2003	2.31	0.52	4.85	1.09	1574
2004	2.36	0.39	4.33	1.27	1615
Number of days since first vest					
1998	77.82	46.18	295.00	15.00	218
1999	319.45	120.96	700.00	94.00	4334
2000	589.15	117.72	991.00	458.00	2268
2001	1012.61	121.49	1400.00	823.00	2973
2002	1314.12	112.50	1760.00	1188.00	1524
2003	1738.02	114.16	2128.00	1553.00	1574
2004	2121.37	89.01	2216.00	1918.00	1615
Percent of vested options exercised					
1998	0.98	0.10	1.00	0.18	218
1999	0.93	0.18	1.00	0.07	4334
2000	0.89	0.23	1.00	0.01	2268
2001	0.90	0.22	1.00	0.03	2973
2002	0.90	0.23	1.00	0.01	1524
2003	0.92	0.21	1.00	0.01	1574
2004	0.91	0.22	1.00	0.01	1615
Options exercised					
1998	1289.74	4331.76	40000.00	128.00	218
1999	590.57	3573.27	100000.00	16.00	4334
2000	1503.72	9558.97	259600.00	8.00	2268
2001	1317.29	7665.33	288000.00	20.00	2973
2002	2785.44	18156.44	480000.00	40.00	1524
2003	6050.03	35798.28	1120000.00	4.00	1574
2004	6390.33	46423.38	1600000.00	40.00	1615

Table 3: Semiconductor Firm A: Comparison of the Percentage of Vested Options that are Exercised When the Number of Options Exercised is Large (Greater than 10,000 options exercised) and When the Number of Options Exercised is Small (Less than 1,000 options exercised) and for Exercise Events by Year of Exercise

This table provides a comparison of the percentage of vested options that are exercised when the number of options that are exercised is large and when the number that are exercised is small. The exercise experiences are summarized for the year in which the options were exercised.

	Mean	Standard Dev.	Maximum	Minimum	Exercise Experiences
Exercise Events when the Number of Options that are Exercised is $\geq$ 10,000					
Cumulative change in stock price/strike price					
1998	1.61	0.18	1.89	1.42	32
1999	2.58	0.86	4.91	1.21	134
2000	5.86	2.27	10.04	1.53	268
2001	2.99	0.96	4.95	1.54	725
2002	2.81	0.82	4.65	1.32	673
2003	2.48	0.91	4.34	1.09	1268
2004	2.23	0.94	4.33	1.27	1369
Percent of vested options exercised					
1998	0.68	0.43	1.00	0.18	5
1999	0.80	0.28	1.00	0.15	47
2000	0.75	0.34	1.00	0.06	48
2001	0.87	0.26	1.00	0.07	47
2002	0.79	0.32	1.00	0.04	41
2003	0.68	0.38	1.00	0.08	84
2004	0.72	0.33	1.00	0.09	94
Exercise Events when the Number of Options that are Exercised $\leq$ 1,000					
Cumulative change in stock price/strike price					
1998	1.38	0.10	1.96	1.17	186
1999	2.50	0.71	4.68	1.32	4201
2000	4.95	1.81	8.85	2.28	2000
2001	2.94	0.40	5.45	1.77	2260
2002	2.94	0.46	3.58	1.34	850
2003	2.38	0.48	3.30	1.55	310
2004	2.41	0.34	3.18	2.01	256
Percent of vested options exercised					
1998	0.99	0.07	1.00	0.22	186
1999	0.93	0.17	1.00	0.07	4201
2000	0.90	0.22	1.00	0.02	2000
2001	0.88	0.24	1.00	0.03	2260
2002	0.86	0.25	1.00	0.03	850
2003	0.84	0.29	1.00	0.03	310
2004	0.79	0.31	1.00	0.01	256

Figure 1: Fractional Exercise Experience for a Small Option Grant

The figure illustrates stock price dynamics, vesting structure, strike price dynamics, and option exercise events for a small option grant. The employee received the option grant in 1997 and the firm is a large semi-conductor company.

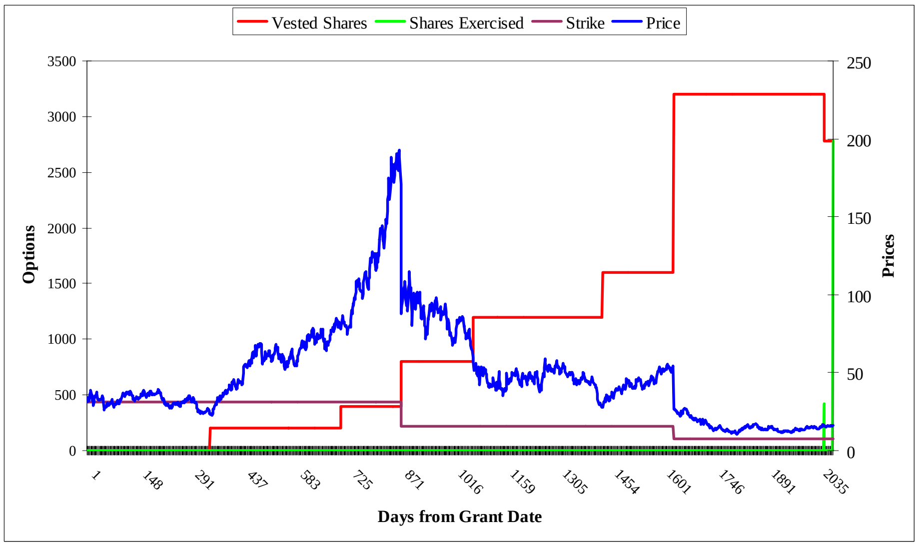

Figure 2: Semiconductor Company A: Fractional Exercise Experience for a Small Option Grant

The figure illustrates stock price dynamics, vesting structure, strike price dynamics, and option exercise events for a small option grant. The employee received the option grant in 1997 and the firm is a large publicly semi-conductor company.

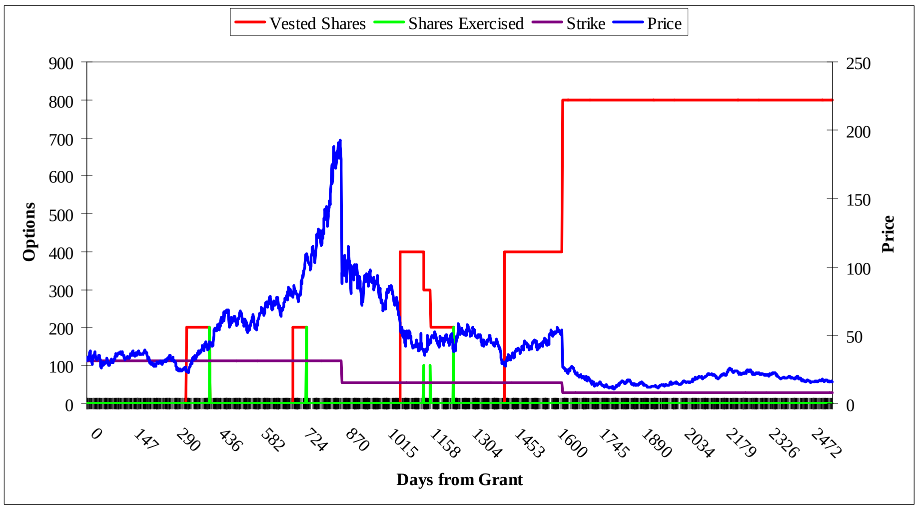

Figure 3: Semiconductor Company A: Exercise Experience for a Large Option Grant to the CEO

The figure illustrates stock price dynamics, vesting structure, strike price dynamics, and option exercise events for a large option grant. The employee is the CEO of a large publicly traded semiconductor company and the options were granted in 1996.

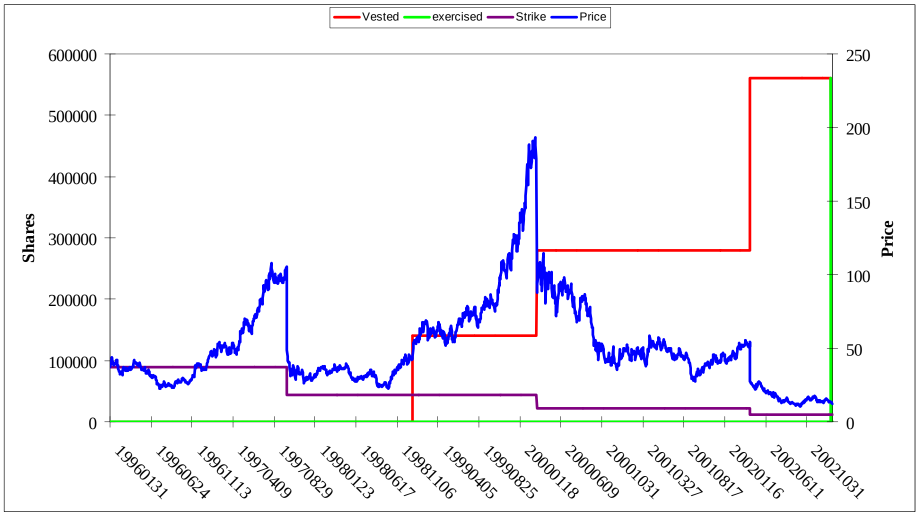

Figure 4: Semiconductor Company A: Exercise Experience for a Large Option Grant to the CEO

The figure illustrates stock price dynamics, vesting structure, strike price dynamics, and option exercise events for a large option grant. The employee is the CEO of a large publicly traded semiconductor company and the options were granted in 1998.

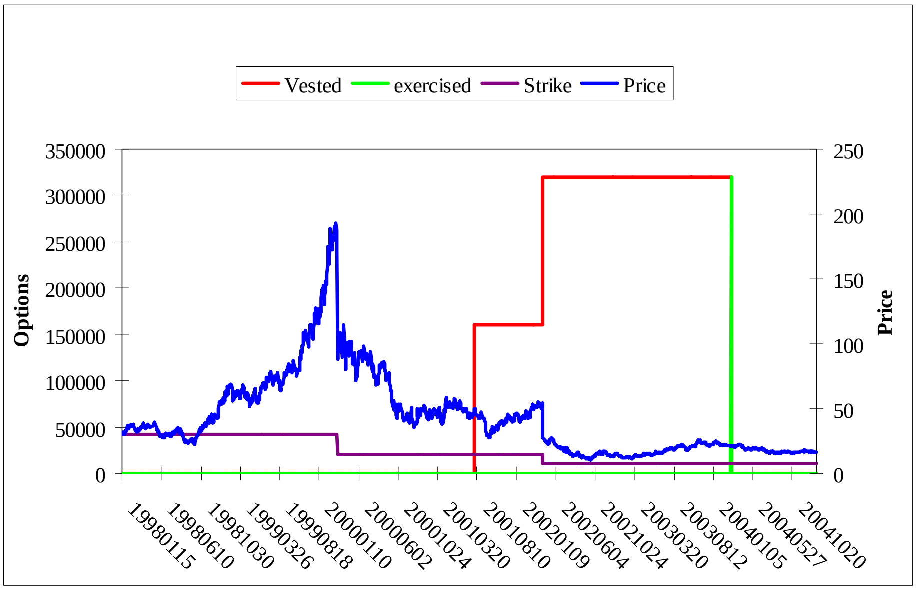

Figure 5: Semiconductor Company B: Exercise Experience for a Large Option Grant to the CEO

The figure illustrates stock price dynamics, vesting structure, strike price dynamics, and option exercise events for a large option grant. The employee is the CEO of a large publicly traded semiconductor company and the options were granted in September, 1998.

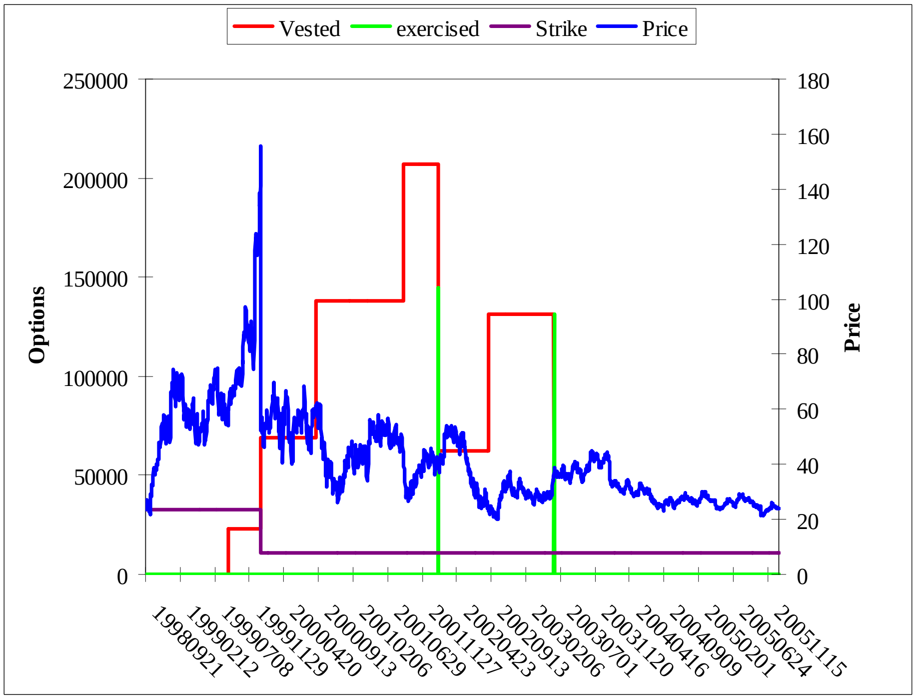

Figure 6: Semiconductor Company B: Exercise Experience for a Large Option Grant to the CEO

The figure illustrates stock price dynamics, vesting structure, strike price dynamics, and option exercise events for a large option grant. The employee is the CEO of a large publicly traded semiconductor company and the options were granted in July, 1998.

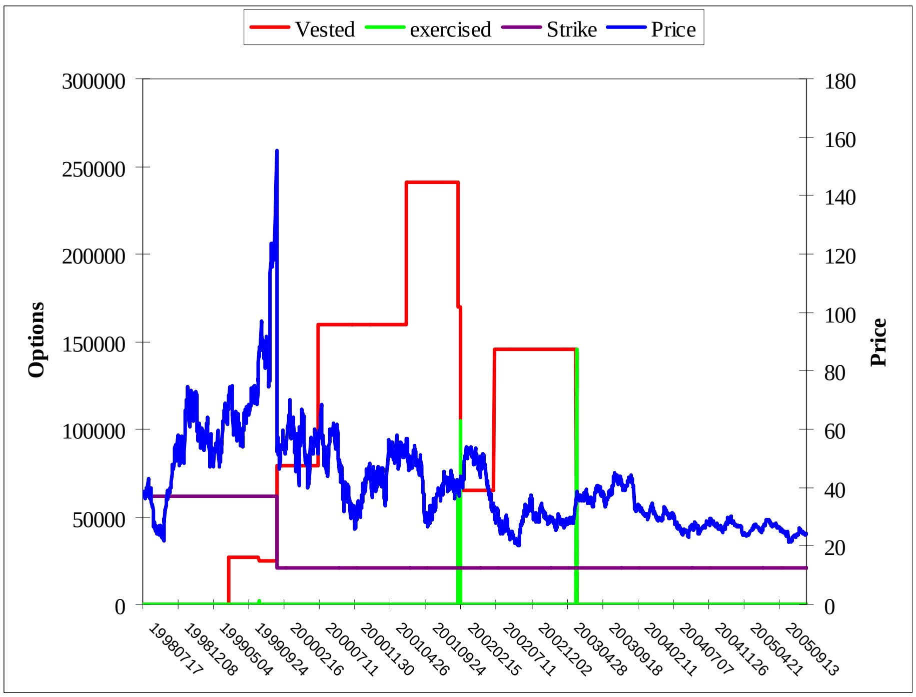

Figure 7: Software Company: Exercise Experience for a Large Option Grant to the CEO

The figure illustrates stock price dynamics, vesting structure, strike price dynamics, and option exercise events for a large option grant. The employee is the CEO of a large publicly traded software company and the options were granted in 1996.

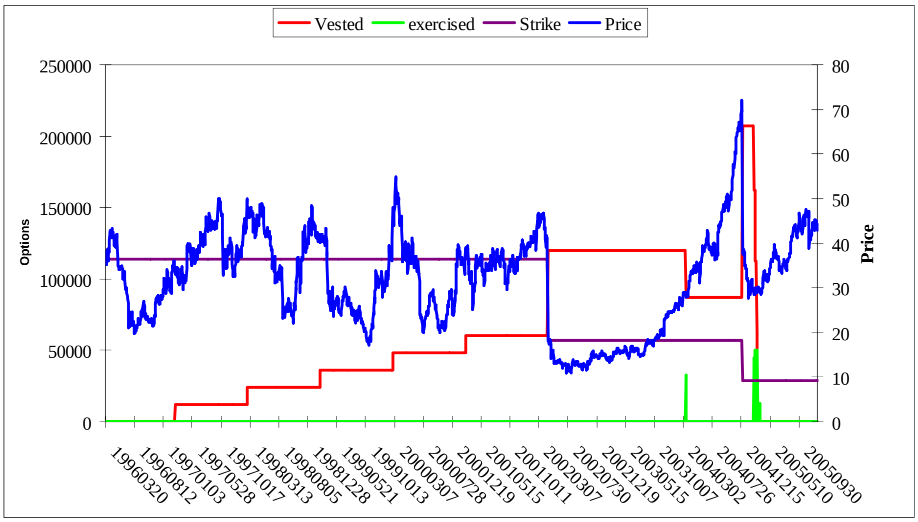

Figure 8: Software Company: Exercise Experience for a Large Option Grant to the CEO
The figure illustrates stock price dynamics, vesting structure, strike price dynamics, and option exercise events for a large option grant. The employee is the CEO of a large publicly traded software company and the options were granted in 1996.