



The exercise and valuation of executive stock options¹

Jennifer N. Carpenter*

Stern School of Business, New York University, New York, NY 10012-1126, USA

Received 25 March 1996; received in revised form 1 June 1997

Abstract

In theory, hedging restrictions faced by managers make executive stock options more difficult to value than ordinary options, because they imply that exercise policies of managers depend on their preferences and endowments. Using data on option exercises from 40 firms, this paper shows that a simple extension of the ordinary American option model which introduces random, exogenous exercise and forfeiture predicts actual exercise times and payoffs just as well as an elaborate utility-maximizing model that explicitly accounts for the nontransferability of options. The simpler model could therefore be more useful than the preference-based model for valuing executive options in practice. © 1998 Elsevier Science S.A. All rights reserved.

JEL classification: G13; J33; M41

Keywords: Executive stock options; Exercise policy; Option valuation; Nontransferable option; Utility maximization

1. Introduction

With the explosive growth in the use of stock options as executive compensation, investors, economists, and accountants have become increasingly concerned

* Corresponding author. Tel.: 212 998 0352; fax: 212 995 4233; e-mail: jcarpen0@stern.nyu.edu.

¹ This article is based on a chapter of my dissertation. I would like to thank my committee members, Sanford Grossman, Bruce Grundy, David Larcker, Robert Litzenberger, and Robert Stambaugh, as well as Yakov Amihud, Pierluigi Balduzzi, Giovanni Barone-Adesi, Edwin Elton, Kose John, Craig Mackinlay, Eli Ofek, Krishna Ramaswamy, Robert Reider, Matthew Richardson, Clifford Smith (the editor), Marti Subrahmanyam, Mark Vargas, David Yermack, and especially Kevin Murphy (the referee) for helpful comments and suggestions.

about the cost of these options to shareholders. Any researcher or practitioner trying to value a firm must assess the value of the claim on equity that executive options represent. Executive option valuation is also important to corporate boards and compensation consultants, and is even becoming an issue outside the U.S., in countries such as the U.K., Japan, and India.

Because these options are not transferable, their optimal exercise policy differs from that of ordinary options. This feature makes these options more difficult to value than ordinary options. This problem has thwarted the efforts of the Financial Accounting Standards Board (FASB) over the last decade to develop a standard requiring firms to deduct the cost of options from earnings. This paper shows that a simple model combining the ordinary American option exercise policy with random, exogenous early exercise and forfeiture describes exercise patterns in a sample of 40 firms just as well as an elaborate utility-maximizing model that explicitly accounts for the nontransferability of options. Because the exercise policies of executives are a crucial determinant of the cost of these options to shareholders, this result suggests that the simpler model is equally good for valuing options. Thus, while opponents of the proposed FASB standard have argued that the need to account for nontransferability makes option valuation too complex, the results here suggest that it is possible to value executive options in practice.

The nontransferability of the options means that their value to executives can be different from their cost to shareholders. The focus of this paper, like that of the FASB debate, is on the cost of the options to shareholders. This cost is the amount that an unrestricted outside investor would pay for such options. This amount is like the value of ordinary American options, with one important difference. Exercise decisions for executive options are not made by the outside investor, but rather by the executives, who cannot trade or hedge the options and therefore might not make the same exercise decisions as unconstrained option holders. For example, executives might exercise their options earlier than usual for the purpose of diversification or liquidity. They might also be forced to exercise early or forfeit options upon separation from the firm. Other factors, such as taxes or inside information, might lead to late exercise.

In order to value executive stock options, that is, in order to estimate the company's opportunity cost, we need an understanding of the exercise decisions of executives. While the effects of hedging restrictions on the exercise policies of risk-averse executives may be complex in theory, their practical impact on exercise patterns represents an empirical question. To address this question, I compare two models of a representative option holder's exercise policy. The first, a simple extension of the ordinary American option model, introduces an exogenous 'stopping state', in which the executive automatically exercises or forfeits the option. This state arrives with some fixed probability, given as the 'stopping rate', each period. The stopping state serves as a proxy for anything that causes the executive to stop the option early, including the desire for

liquidity, voluntary or involuntary employment termination, or any other event relevant to executives but not to unrestricted option holders. The model is essentially a binomial version of the continuous-time model of Jennergren and Naslund (1993). The second model assumes the executive exercises the option according to a policy that maximizes expected utility subject to hedging restrictions, as in Huddart (1994) and Marcus and Kulatilaka (1994). This utility-maximizing model not only includes a stopping state, but also includes other unobservable factors, such as the executive's risk aversion, his outside wealth, and his potential gain from voluntary separation.

If factors underlying the two models were observable, we could simply compute exercise patterns under both models and compare these patterns with actual exercises from a sample of options. Given that the factors are not observable, I start by calibrating the models, choosing factor values that make modeled exercise payoffs, times, and cancellation rates best match sample averages. Next, I examine the performance of the calibrated models in predicting actual exercise patterns for a sample of 40 firms with data from the period 1979 to 1994.

I expect the utility-maximizing model to perform better than the extended American option model because it has more flexibility and allows for richer forms of interaction between early exercises, or forfeitures, and the level of the stock price. Surprisingly, the two calibrated models perform almost identically. To be sure that the utility-maximizing model can do no better, I also examine its performance under a variety of other parametrizations. In no case does the utility-maximizing model outperform the extended American option model.

One conclusion is that the stopping rate is essentially a sufficient statistic for the utility parameters. More broadly, the results suggest that exercise patterns can be approximately replicated merely by imposing a suitable stopping rate, without the need to make assumptions about executive risk aversion, diversification, and the value of new employment. This implies that a simple extension of the usual binomial model can be adequate for valuing executive stock options. The simpler model also comes closer to meeting accounting standards of observability and verifiability than does a model that requires assumptions about the risk preferences of executives.

For the purpose of comparison, I also compute option values according to the method that the FASB recommends, which is an extension of the Black–Scholes (1973) model that replaces the stated expiration date with the option's expected life. To first order, the FASB method approximates the value that the option would have if the stopping time were independent of the stock price. The option value under the FASB method is less than either of the other two estimated option values.

The remaining sections are organized as follows. Section 2 reviews the existing literature on executive stock options. Section 3 develops a theoretical framework for option valuation, establishes the link between option valuation and the executive's exercise policy, and presents the two alternative models. Section 4 compares the ability of the models to explain actual exercise behavior

in a sample of NYSE and AMEX firms. Section 5 discusses the choice of a suitable stopping rate, and Section 6 concludes.

2. Previous research

Because of the difficulty in obtaining adequate data on option grants and exercises from public sources, little empirical research exists on employee option exercise patterns. Since the expansion of Securities Exchange Commission (SEC) disclosure requirements in 1992, firms have provided disaggregated information about option grants in proxy statements. However, before that date, proxy statements typically disclosed just the average strike price and a range of expiration dates of newly granted options. Huddart and Lang (1996) study exercise behavior in a sample of eight firms that volunteered internal records on option grants and exercises from 1982 to 1994. They find a pervasive pattern of option exercises well before expiration. They also examine the ability of different variables to predict months with intense exercise activity. For example, they find a positive relation between option exercise activity and recent stock price appreciation.

A number of other empirical papers use data on option grants to estimate the value of executive stock options using the Black–Scholes (1973) formula, as adjusted for dividends by Merton (1973). These include Antle and Smith (1985), Foster et al. (1991), and Yermack (1995). For example, Yermack (1995) reports that options represented about one-third of the average compensation of chief executive officers in 1990 and 1991, based on their Black–Scholes value.

The importance of executive stock options and the heat of the FASB valuation controversy have inspired a variety of theoretical papers about option valuation. Huddart (1994) and Marcus and Kulatilaka (1994) develop binomial models of the exercise policy that maximizes the expected utility of option holders when they are unable to sell or hedge their options. Other papers, such as Cuny and Jorion (1995) and Jennergren and Naslund (1993), focus instead on the impact on option value of the possibility that the executive may leave the firm, thereby forfeiting or exercising the option. Examples of this effect also appear in Rubinstein (1995). These papers all consider the value of the option from the viewpoint of the option writer. By contrast, Lambert et al. (1991) use certainty equivalents to value the option in a utility-based framework from the option holder's point of view. In earlier work, Smith and Zimmerman (1976) provide bounds on executive stock option value.

3. Executive stock option valuation

Standard American option pricing theory assumes that holders of options can trade freely. This assumption implies that the option holder will exercise the

option according to a strategy that maximizes its market value. Therefore, the value of the ordinary option does not depend on the risk preferences of either the option holder or the option writer.

Executive stock options are nontransferable, and section 16-c of the Securities Exchange Act prohibits insiders from selling their firm's stock short. These restrictions mean that executives holding such options can neither sell nor hedge their positions. Consequently, they might exercise options that would have more value if left unexercised, for the purpose of portfolio diversification, consumption, or employment termination.

Hedging restrictions imply that executives' personal valuations of their options depend on their risk preferences and endowments, as Section 3.1 describes. Hedging restrictions also imply that the cost of options to shareholders can depend on the personal characteristics of the company's executives, even though the shareholders, the option writers, face no restrictions in trading or hedging their short position. To see why, think of the cost of the option to shareholders as the amount an unconstrained investor would pay for the option, with the understanding that executives make the exercise decisions. Equivalently, this is the amount shareholders would have to pay to an unconstrained outside institution or investor to assume their short position. This option cost can still depend on the characteristics of executives because they control the timing of option exercises, and thus control the option payoffs.

To examine the practical effects of the nontransferability of options on the exercise policies of executives, this paper compares a model in which early exercises and forfeitures arise exogenously to a model in which such early option terminations result from a fully specified utility-maximization problem. This section develops the two models in detail. The section begins by drawing a distinction between the value of options to executives and the cost of options to shareholders. Section 3.2 sketches a theoretical framework for pricing the risks associated with executive stock options, which explains why valuing options from the writer's viewpoint becomes a matter of determining the exercise policies of the option holders. Sections 3.3 and 3.4 develop alternative models of executive exercise policies. The exercise policy of a given model not only determines the option cost, but also provides average values of observable variables, such as the times and payoffs of option exercises. On the basis of these forecasts, in Section 4, I compare each model's ability to fit the data.

3.1. Distinction between the value to executives and the cost to shareholders

Because they cannot sell or hedge their options, executives value their options subjectively. One measure of the value of the option to an executive is the dollar amount of a cash bonus that would make him equally happy. As Lambert et al. (1991) show, the personal valuations executives make would depend on their

risk preferences, endowments, constraints on their portfolio or mobility, and other details of their personal lives.

The following special cases illustrate the distinction between the value of options to executives and their cost to shareholders. First, suppose the options are European, that is, not exercisable before the expiration date. Suppose further that there is no requirement that executives forfeit options if they leave the firm. Then, the cost of executive options is that of ordinary European options, because their payoffs are exactly the same. Yet the value to executives is less than that of ordinary options because they cannot sell or hedge them. This discrepancy between the executive's valuation and the cost to shareholders might suggest that the option is an inefficient form of compensation. For example, no such discrepancy would exist with a cash payment. Of course, the reason for using an option as compensation instead of a cash payment is that it might yield superior managerial performance. The problem of how best to compensate managers given the benefits of improved incentives and the costs of inefficient risk-sharing is the subject of a large literature on the principal-agent problem. According to Holmstrom and Hart (1987), option-like sharing rules do not always represent the best compensation strategy. The optimal shape of the sharing rule can be arbitrary. Nevertheless, the widespread use of options as compensation suggests that firms find their benefits outweigh their costs. For example, Brickley et al. (1985) report evidence that a firm's stock price responds favorably to an announcement of the adoption of an option compensation plan.

Now consider American options, those that are exercisable any time prior to expiration, in the case that the stock pays no dividends. Then the options are worth more to executives than if they were European. However, any exercise by the executives prior to expiration reduces the option cost to shareholders. If the options are American and the stock pays dividends, the options are still worth more to executives than if they were European, but the difference in the cost to shareholders depends how executives use the right to exercise early. The fact that executives must forfeit options if they leave the firm when options are unvested or when the stock price is below the strike price reduces the option value from the viewpoint of both executives and shareholders.

3.2. Pricing nonmarket risks in the option payoff

How should shareholders or any other unrestricted investors value stock options controlled by executives? If the option payoffs could be replicated with a trading strategy using the company's stock, then the value of the options would be the cost of setting up that trading strategy. This replication cost could be represented using a mathematical construct known as the stock's 'risk neutral' probability measure. In particular, the cost of the replicating stock portfolio would be the expected value, under the risk-neutral measure, of the

option payoffs discounted at the riskless rate. See, for example, Cox and Ross (1976) or Harrison and Kreps (1979).

This contingent-claim approach might not be appropriate for valuing executive option payoffs, because it might not be possible to replicate those payoffs with the underlying stock. Unlike the exercise decision for an ordinary tradeable option, which depends only on the stock price path, the time at which executives exercise or cancel options might also depend on nonmarket risks, such as whether executives suffer personal liquidity shocks, or leave the firm.

I assume that any risks in option payoffs that shareholders cannot hedge with the stock are idiosyncratic across different executives. It is feasible to hold a diversified portfolio of executive stock options by holding a portfolio of stocks. In a diversified portfolio, these idiosyncratic risks become trivial. Therefore, shareholders should value options at their expected value with respect to idiosyncratic risks. On the other hand, a diversified portfolio of options remains subject to market risks, and is therefore not worth its expected value, but instead has the same value as a replicating portfolio of other assets. Therefore, even though a given executive stock option is not strictly a contingent claim on traded assets, the cost of the option to shareholders is still

$$E(\zeta_{\tau}(S_{\tau} - S_0)^+ 1_{\{\tau \geq t_v\}}), \quad (1)$$

where S is the stock price, τ is the random time at which the executive stops the option through exercise or cancellation, t_v is the vesting date, and ζ is the pricing kernel appropriate for valuing ordinary options. The random time τ is essentially a plan that specifies when the executive will stop the option for every possible sequence of future events. With this approach, valuing executive stock options becomes a matter of determining the exercise policies of executives.

3.3. *The exercise policy in the extended American option model*

This section presents a binomial model of a representative option holder's exercise rule, which extends standard theory simply by introducing random, exogenous early exercises and forfeitures. In each period, there is some probability, q , that a stopping state will occur. The occurrence of this state is independent of the path of stock prices. In a stopping state, the executive exercises the option if it is in the money and vested, or else forfeits the option. Otherwise, the executive acts according to standard American option theory, whereby he exercises the option or leaves it alive, depending on which action gives the option greater value.

This extended American option model, based on the continuous-time model of Jennergren and Naslund (1993), requires only one parameter more than the standard model. This additional parameter is the stopping rate q . The aim of the model is to capture as simply as possible the fact that executives may exercise options earlier than standard theory predicts, and that option forfeitures can

take place. The stopping event serves as a proxy for anything that might cause the executive to deviate from the usual exercise policy by stopping the option early, such as a liquidity shock, a desire to diversify, employment termination, or any forced exercise or forfeiture.

The remainder of this section provides details of the construction of the executive's exercise policy. The stock price process in the model is a standard binomial, multiplicative random walk, and the interest rate r is constant. The strike price of the option is equal to the stock price on the grant date. Without loss of generality, I set the strike price and the stock price at the option grant date equal to one. In each period, the stock price can move up by a factor u or down by a factor d with equal probability. To construct a stock price process with mean return μ , volatility σ , and dividend rate δ , let

$$u = e^{(\hat{\mu} - \delta)h - \log(\cosh \sigma\sqrt{h}) + \sigma\sqrt{h}}, \quad (2)$$

$$d = e^{(\hat{\mu} - \delta)h - \log(\cosh \sigma\sqrt{h}) - \sigma\sqrt{h}}, \quad (3)$$

where $\hat{\mu} = \log(1 + \mu)$, n is the number of periods per year, and $h = 1/n$ is the length of each period in years. The probability that the stock moves up under the risk-neutral measure is

$$\tilde{p} \equiv (e^{(r - \delta)h} - d)/(u - d). \quad (4)$$

The exercise policy of the executive maximizes the option value, or cost to the option writer, just as in the ordinary American option model. The policy is essentially a list of the optimal exercise decisions at every possible decision state. A decision state for the executive is represented by a time, a stock price, and an indication of whether the executive is in a stopping state. When he is not in an automatic stopping state, the executive decides to exercise or continue according to which action maximizes the option's value. The value of the option, if left alive, depends on future exercise decisions. Therefore, determining the exercise policy requires working backward from the expiration date, recording the exercise decision and the resulting maximized option value at each possible state. In a stopping state, the executive automatically either exercises the option if it is vested and in the money, or else forfeits the option. The option value in that state is either its exercise value or zero. Otherwise, the executive only exercises the option if its exercise value exceeds the value the option would have if not exercised. The unexercised option can end up in one of four possible states the following date, because the stock price can move up or down, and the stopping event can arrive or not. The value of the unexercised option is thus the discounted probability-weighted average of the four possible option values at the next date, using the true probability of a stopping state, and the risk-neutral probability of a given stock price move. This valuation method is consistent with the theoretical

framework outlined in Section 3.2. The appendix contains formal definitions of the decision function and value function for this maximization problem.

3.4. *The exercise policy in the utility-maximizing model*

This section describes a model of an exercise rule that maximizes the expected utility of terminal wealth of a representative option holder. It builds on prior work by Huddart (1994) and Marcus and Kulatilaka (1994). In the extended American option model presented above, deviations from the standard option policy arise exogenously. In this model, such deviations are the outcome of an optimization problem that explicitly accounts for restrictions on the manager's ability to hedge the option.

The optimization problem is based on the following assumptions. The executive holds options and outside wealth x . He invests the outside wealth as well as the proceeds of any early option exercise until the option expiration date. The executive has constant relative risk aversion, with coefficient A . He chooses an exercise policy that maximizes the expected utility of his wealth at the option expiration date.

To accommodate the possibility of option forfeiture or an early exercise caused by a nonmarket event, I introduce a stopping state in this model as well. However, the executive does not automatically stop in the stopping state. Instead, I make the decision to stop an endogenous feature of the model by supposing that, in each period, there is some probability q that the executive is offered a monetary payoff y to leave the firm. Leaving means stopping the option, either through exercise or forfeiture. Depending on the size of the payoff, the vesting status of the option, and the level of the stock price, the executive may decline the payoff in order to keep his option alive. Introducing this payoff parameter allows for various forms of dependence between the stock price and the risk of forfeiture or early exercise. When y is a large sum, the executive always accepts the offer, but when y is small, he tends to decline when the option is not vested or when the stock price is near the strike price.

The utility value of continuing in any state depends on future decisions to stop or continue, so the exercise rule must be determined by working backward from the expiration date. This backward recursion is possible as long as the executive's decision depends only on the prevailing level of the stock price and whether or not he is in a stopping state, not on the past stock price path. Therefore, I assume that if the executive exercises the options, he will do so all at once.

The executive invests his initial outside wealth, and the proceeds of any early option exercise, in the constant proportion portfolio of the stock and bond that would be optimal in the absence of the option and the possibility of receiving y . The value of this portfolio is path independent. In any state, the portfolio value is only a function of the time and the prevailing stock price. This portfolio is

a binomial version of the continuous-time portfolio developed in Merton (1969, 1971). By contrast, in the models of Huddart (1994) and Marcus and Kulatilaka (1994), nonoption wealth is invested in the riskless asset. Their assumption places an artificial constraint on the portfolio choice of the executive after the option is exercised, which distorts the exercise decision. Investing nonoption wealth in the Merton portfolio is more appealing, although not fully optimal in the presence of the option. Full optimality would allow the executive to choose investment and exercise strategies simultaneously. This scenario is intractable because the nonnegativity constraint on the stock holdings would become binding along some stock price paths, but not along other paths. Under these conditions, the optimal portfolio value would be a path-dependent function of the stock price, and backward recursion would be impossible.

Because of the shape of the utility function, rescaling all payoffs leaves the executive's optimal policy unchanged. Therefore, without loss of generality, the number of options is one and the initial stock price is one. Initial outside wealth, x , represents dollar wealth divided by the initial value of the shares underlying the option. Similarly, y represents the payoff for leaving, divided by the initial value of shares under option. A formal presentation of the optimal exercise policy appears in the appendix.

The utility-maximizing model abstracts from a number of aspects of the executive's situation that complicate the optimization problem. First, the option holder has some control over the underlying stock price process. Indeed, Agrawal and Mandelker (1987), Lambert et al. (1989), and DeFusco et al. (1990) find evidence that option-compensated managers increase the asset variance and leverage of their firms, and reduce dividends to shareholders. My valuation approach essentially treats the underlying stock price process as the one that already incorporates any changes in managerial strategy due to the option's incentive effects, ignoring the potential interaction between managerial policy and exercise policy. Thus, I do not try to quantify the incentive effects of the option.

In addition, the model presented here does not account for the fact that the firm's decisions about the executive's future compensation mix may depend on the state of existing options, and knowledge of this dependence may affect the executive's exercise policy. This model also leaves out the fact that option holders may have private information about the future path of the stock price. There is no published evidence that option exercises by insiders are followed by significant abnormal returns. See Seyhun (1992) (footnote 20). However, before 1991, the SEC's restriction on the resale of shares acquired through the exercise of options may have made option exercise an ineffective way for insiders to act on private information. Now that the SEC restriction has been lifted, it may be easier for insiders to incorporate private information in their exercise policies. Finally, the model ignores the presence of taxes.

4. Empirical study of option exercises

Because the utility-maximizing model provides a richer description of the executive's situation, it may have more theoretical appeal than the extended American option model. An important question, though, is whether it is better for valuing options in practice. The question cannot be answered directly, because executive stock option values are not observable. Instead, I address the question of whether the utility-maximizing model can better explain actual exercise patterns. This question bears directly on the issue of valuation.

Section 4.1 describes my sample of option exercises from 40 firms. Section 4.2 explains how I use the exercise policy from a given model to forecast observable variables, including the times and payoffs of option exercises, and the annual rate at which options are canceled. Section 4.3 presents the selection of base case parameters for each of the two models that best fit a representative firm constructed from the sample. In Section 4.3, I also list a variety of other parametrizations and examine the option value and forecasted exercise variables implied by each parametrization. In Section 4.4, I test and reject the null hypothesis that the more flexible utility-maximizing model can explain cross-sectional variation in exercise times and payoffs better than the extended American option model.

4.1. Sample of option exercises and cancellations

The sample variables relevant to these models include average times to exercise, stock prices at the time of exercise, and vesting periods of ten-year nonqualified or incentive stock options for 40 firms on the NYSE or AMEX. I begin with a collection of 70 firms for which I have option grant information that includes specific grant dates and exercise prices. Nearly half of these come from a proprietary database of large firms constructed by Mark Vargus at Wharton, which is described in Vargus (1994). The database uses information from a variety of corporate filings, including proxies, Forms 10-K and Forms S-8. I augment this database to include smaller firms.

Form S-8, the option plan prospectus, is one of the only public documents to give disaggregated information about grant dates and strike prices of outstanding options. This information is essential for determining the age of an exercised option. Prior to May 1991, insider filings of option exercises contain the strike price and exercise date, from which one can determine the prevailing level of the stock price at the time of exercise. However, these insider filings do not contain the grant date of the option.

Starting with firms in the smallest size decile on the Center for Research in Security Prices (CRSP) database, I search option plan prospectuses for explicit information about grant dates, strike prices, and vesting periods of ten-year options. For each firm, I select one option with a strike price that is distinct from

all other options granted by that firm in the database and equal to the stock's market price on the grant date. If more than one option is available, I select the last one expiring before the end of 1992. I then search all option exercises filed with the SEC by insiders at that firm for exercises with a matching strike price, adjusted for stock splits and stock dividends. Although this procedure leaves open the possibility of selecting option exercises from grants that are not in the database, this is unlikely. For only two of the 70 firms do I find exercises with strikes that match a selected option's split-adjusted strike price, but are not in the time range when the option's price was in effect. I eliminate these option grants. If there are no exercises reported from that grant, then I select the next earlier grant. If no grants before 1982 are available with corresponding exercises filed with the SEC, I consider grants from 1983 to September 1984.² I find at-the-money option grants followed by at least one exercise for 40 firms. One reason I am unable to find matching exercises for some of the grants may be that the options had tandem stock appreciation rights, allowing for cash settlement of the options. I also eliminate four firms that merged.

Firms in the sample tend to be large manufacturing firms. Based on their market capitalization at the option grant date, 63% are in CRSP size deciles 8–10 and 25% are in deciles 5–7. Of the sample firms, 85% were in the Manufacturing Division of the Standard Industry Classification at their grant date, while this division contains only about half the firms in the CRSP database that existed in 1982.

For each of the 40 options in the sample, I compute the average time of exercise, τ , and ratio of the stock price at exercise to the strikeprice, s_τ , across exercises from that option. I weight the average by the split-adjusted number of shares in the transaction. The options generally vest according to a schedule, such as a quarter of the grant per year over the first four years. I approximate a single average vesting date, t_v , for each grant.

I estimate the stock return volatility, σ , and the dividend rate, δ , for each firm using monthly data from CRSP. Volatilities are estimated over the five years prior the grant date and dividend rates are estimated over the ten years from the grant date to the expiration date. I also record the stock price at the expiration date, s_{10} , normalized by the stock price at the grant date, to get a measure of the overall performance of the stock over the potential life of the option.

Table 1 presents summary statistics for these data. As Table 1 shows, the options were exercised after an average of 5.8 years and the stock price at the time of exercise was 2.8 times the strike price at the time of exercise. The average volatility of firms in the sample was 31% and the average dividend rate was 3%.

² I prefer options that expired before the end of 1992, to avoid early exercises triggered by anticipation of the 1993 tax increase. Tax cuts in 1981 and 1986 would seem to alter option exercise strategies only by delaying exercise from the time the cut is anticipated until it is enacted, whereas a tax hike might cause an exercise to occur several years before it would otherwise have taken place.

Table 1
Summary statistics of firms with option exercises

Panel A displays descriptive statistics for a sample of option exercises at 40 firms from 1979 to 1994. For each firm, the variable s_t is the average of the ratio of stock price at the time of option exercise to the option strike price, the variable τ is the average time to option exercise in years, and the variable t_v is the average vesting date of the options in years. The variables σ and δ are the estimated annualized stock return volatility and dividend rate for each firm, respectively. The variable s_{10} is the ratio of the stock price at the option expiration date to the strike price. Panel B presents the correlation matrix for these variables.

	s_t	τ	t_v	σ	δ	s_{10}
<i>Panel A: Descriptive statistics</i>						
Minimum	1.15	1.15	0.00	0.19	0.00	0.07
Maximum	8.32	9.48	4.41	0.57	0.07	9.12
Average	2.75	5.83	1.96	0.31	0.03	3.27
Median	2.47	6.08	2.00	0.31	0.03	2.75
Standard deviation	1.42	2.25	1.03	0.10	0.02	2.25
<i>Panel B: Correlation matrix</i>						
τ	0.14					
t_v	0.43	0.42				
σ	-0.19	-0.02	-0.01			
δ	-0.05	-0.08	-0.18	-0.33		
s_{10}	0.60	0.04	0.22	-0.31	0.12	

By contrast, among the five firms in Huddart and Lang (1995) that granted ten-year options, the average time of exercise is 3.4 years, the average volatility is 34%, and the average dividend rate is 5%. Across all exercises in their sample, the average ratio of the price of the stock at the time of exercise to the strike price is 2.2. Exercises at firms in their sample tend to be earlier and at lower stock prices, which is consistent with the high average dividend rate. The earlier exercises could also stem from the fact that the option holders in their sample represent employees at all levels of the firm, not just top executives. They might be less affluent than executives, more risk averse, and less inclined to hold options for strategic reasons.

Panel B of Table 1 also provides a correlation matrix for the firm variables. Options with longer vesting periods tend to be exercised later and deeper in the money. As Panel B shows the correlation of the vesting period, t_v , with the time to exercise, τ , is 42%, and the correlation of t_v with the level of the stock price at exercise, s_t , is 43%. Also, as expected, at firms with strong overall stock price performance, options are exercised deeper in the money. For these firms, the correlation between s_t and s_{10} is 60%. Recalling results from standard American option theory, such as Kim (1990), the exercise policy for an option on a stock with a proportional dividend requires the option to be exercised once the stock

price reaches a critical boundary. That critical point is higher the longer the time left to expiration, the higher the volatility of the stock, and the lower the dividend rate. These relations are not apparent in my sample. The correlation between the stock price at the time of exercise, s_τ , and the dividend rate, δ , is negligible. The correlation between s_τ and volatility σ , -19% , has the wrong sign. However, I believe this is due to an irregularity in the data. The small, higher volatility firms in the sample tend to have poorer stock price performance during the time period of this study, which includes the crash of 1987. Indeed, the correlation between volatility and terminal stock price is -31% . Thus, it may not be volatility, but rather poor performance, that explains why exercises at these smaller firms are not very deep in the money. The correlation between s_τ and τ of 14% is only slight and also has the wrong sign.

Given the underlying question of option valuation, I attempt to gain information about options that result in a zero payoff as well as those that result in an exercise. The event of zero payoff, or cancellation, occurs if an option is forfeited or if it expires. The SEC does not require insiders to file information about canceled options. However, in their annual reports, firms often give an inventory of their options, listing the number of options outstanding at the beginning of the year, options granted, options exercised, and options canceled. I construct a sample of cancellation rates for 52 of the original 70 firms. I define the cancellation rate as the average fraction of outstanding options forfeited or expired per year. To measure the cancellation rate for a given firm, I take up to ten years of option inventories from annual reports and compute the average ratio of the number of options canceled to the sum of the number of options outstanding plus half the number of options granted. In some cases, annual reports combine ten-year options and options with terms other than ten years in the same inventory, or they indicate that tandem stock appreciation rights are outstanding but do not make clear whether their exercise counts as an option exercise or cancellation. I include the firm in the sample if I find at least three years of data that do not suffer from these problems. I use up to ten years of data for each firm, from 1984 to 1993. Where four of the firms reduced the strike prices on their options, I treat this as a cancellation of the original option and a grant of a new option. The cancellation rates of the 52 firms range from 0.7% to 34.3% with a mean of 7.3% , a median of 4.5% , and a standard deviation of 7.1% .

4.2. Model forecasts of exercise and cancellation variables

The exercise policy prescribed by a given model determines the mean values of the two exercise variables, the level of the normalized stock price at the time of exercise, s_τ , and the time of exercise, τ . The mean values of interest are

$$\hat{s}_\tau = E(s_\tau | \text{option is exercised}, t_v, r, \mu, \sigma, \delta, s_{10}, \theta), \quad (5)$$

$$\hat{\tau} = E(\tau | \text{option is exercised}, t_v, r, \mu, \sigma, \delta, s_{10}, \theta), \quad (6)$$

where t_v is the vesting date, μ and σ are the mean and volatility of the stock return, δ is the dividend rate, and θ is the set of unknown parameters. For the extended American option model, $\theta = q$. For the utility-maximizing model, $\theta = \{A, x, y, q\}$. Note that the predictions \hat{s}_t and $\hat{\tau}$ are conditioned on the fact that the option is exercised and are conditioned on the terminal level of the stock price. In other words, these predictions are the average of the outcomes of the random variables s_t and τ across all stock price paths that result in an exercise and terminate at the level s_{10} , weighted by their conditional probability, assuming the model is true. By conditioning on the overall performance of the stock over the ten years from grant to expiration, the predictions account for the effects of a bull or bear stock market.

Each model also implies an average value for the cancellation rate, cr , at the firm under the assumption that the firm grants an equal number of options to identical executives every year, each following the model's prescribed exercise policy:

$$c\hat{r} = E(cr|t_v, r, \mu, \sigma, \delta, \theta). \quad (7)$$

Thus, $c\hat{r}$ is the average ratio of the number of options canceled during a year, through forfeiture or expiration, to the number of options outstanding at the beginning of that year. The computation takes into account the unconditional distribution of the ages of the options still outstanding in any year, and the likelihood of a cancellation given that age.

In counting cancellations, the data in annual reports combine forfeitures and expirations. The strike prices of canceled options are not publicly available to help distinguish forfeitures from expirations. Therefore, the modeled cancellation rate combines forfeitures and expirations as well. The mean cancellation rate, $c\hat{r}$, may be interpreted as the probability that an option that is outstanding at the beginning of the year gets canceled during that year. Note that this is much lower than the probability that the option is ever canceled throughout its ten-year life.

Finally, to distinguish between the stopping rate q and the cancellation rate, cr , note that the stopping rate is a model parameter or an input to constructing an exercise policy, while the cancellation rate is determined by an exercise policy. The stopping rate governs the frequency with which both nonmarket-driven exercises and forfeitures occur prior to expiration. The cancellation rate is the annual rate at which options are canceled through forfeiture or expiration. The cancellation rate at an actual firm depends on the exercise policies of its executives. Similarly, the cancellation rate of a given model depends on the exercise policy implied by that model, and therefore depends on the stopping rate as well as other model inputs.

4.3. Parameter selection

In Section 4.4, I test whether the utility-maximizing model can forecast the payoffs and times of option exercises at the 40 firms better than the extended

American option model. To do so, I must choose values of the unknown parameters, θ , for each model. As a starting point, I select base case parametrizations for each model that match model predictions of the observable variables for a representative firm to sample averages of those variables. It turns out that the base case parametrizations of the two models generate almost identical forecasts. To provide further evidence that the utility-maximizing model cannot outperform the simpler model, without relying on the validity of the base case parametrizations, I examine the performance of the utility-maximizing model in all other regions of the parameter space. This section describes the various parametrizations, and illustrates their implications for option value and characteristics of the exercise policy.

4.3.1. Base case parametrizations

If executive stock option values were directly observable, I would choose values of the unknown parameters to make model option values best match observed prices. In the absence of observable option prices, I instead calibrate the model to match observable features of the exercise policy. These observable features are the level of the stock price at exercise, the time of exercise, and the cancellation rate. Fortunately, these variables relate directly to the size and timing of nonzero option payoffs, as well as the frequency of zero payoffs, so they are fundamental to option value.

To speed computation time, I construct a representative firm whose vesting date, stock return volatility, dividend rate, and terminal stock price are equal to the sample average values. I set the riskless rate equal to 7%, roughly the average Treasury Bill rate over the time the options were alive. I set the mean annual stock return equal to 15.5%, the sum of the average riskless rate over the time the options were alive plus the average equity premium from 1926 to 1975. For the utility-maximizing model, I fix $A = 2$, because the model is relatively insensitive to the risk-aversion coefficient with outside wealth in the Merton portfolio. I then choose values for q in the extended American option model, and for x , y , and q in the utility-maximizing model, to minimize

$$\frac{(\bar{s}_\tau - \hat{s}_{\tau,0})^2}{S^2(s_\tau)/40} + \frac{(\bar{\tau} - \hat{\tau}_0)^2}{S^2(\tau)/40} + \frac{(c\bar{r} - c\hat{r}_0)^2}{S^2(cr)/52}, \quad (8)$$

where \bar{s}_τ , $\bar{\tau}$ and $c\bar{r}$ are, respectively, the sample averages of the normalized stock price at exercise, time of exercise, and cancellation rate, and $S^2(s_\tau)$, $S^2(\tau)$ and $S^2(cr)$ are their respective sample variances. The quantities $\hat{s}_{\tau,0}$, $\hat{\tau}_0$, and $c\hat{r}_0$ are the predicted values of these variables for the representative firm:

$$\hat{s}_{\tau,0} = E(s_\tau | \text{exercise}, t_v = \bar{t}_v, r = 0.07, \mu = 0.155, \sigma = \bar{\sigma}, \delta = \bar{\delta}, s_{10} = \bar{s}_{10}), \quad (9)$$

$$\hat{\tau}_0 = E(\tau | \text{exercise}, t_v = \bar{t}_v, r = 0.07, \mu = 0.155, \sigma = \bar{\sigma}, \delta = \bar{\delta}, s_{10} = \bar{s}_{10}), \quad (10)$$

$$c\hat{r}_0 = E(cr | t_v = \bar{t}_v, r = 0.07, \mu = 0.155, \sigma = \bar{\sigma}, \delta = \bar{\delta}). \quad (11)$$

All forecasts are generated using a monthly stock price tree and annual decision dates. The exercise variables and the cancellation rate describe the events of positive option payoff and zero option payoff, respectively. All of these variables contain important information about option value and are therefore included in the calibration. Note that the calibration matches the modeled values of the exercise variables, which are conditional on the outcome of exercise, to their average in a sample of exercised options. By contrast, the modeled cancellation rate, which is unconditional, is matched to its average in a sample of all firms with valid cancellation rates regardless of whether the firm was included in the exercise sample or not.

Table 2 contains the parameter values resulting from this procedure as well as model forecasts of the exercise and cancellation variables for the representative firm. Table 2 also lists option values under the different models. The second column from the right, labeled ESO value, gives the theoretically correct value for each option, based on the framework presented in Section 3. The last column on the right, labeled FASB value, gives option values using the method required by FASB (1995), which I discuss below.

Table 2 contains forecasts of the observable variables and option values for the representative firm generated by various parametrizations of the two models. For reference, the last row of Table 2 contains the sample average values of the observable variables. For the sake of comparison, the first model presented in Table 2 is the ordinary American option model with no exercise prior to the vesting date. This is just the extended model with the probability of a stopping state set to zero. Note that under this standard exercise policy, options would be exercised much deeper in-the-money and much later and would have much lower cancellation rates than is typical in our sample. The average value of the stock price at the exercise date would be 3.3 times the strike price, and the average option age at exercise would be 7.6 years. With the exercise policy prescribed by the standard American option model, the option would be worth \$0.39 per dollar of initial stock price.

The second row of Table 2 illustrates that simply introducing a stopping rate of 11% in the value-maximizing model brings the fitted values of the observable variables much closer to the average actual values. It also substantially reduces option value. Under this base case calibration of the extended American option model, option value is only \$0.29 per dollar of initial stock price.

The third row of Table 2 presents the base case utility-maximizing model. To make a good fit, the values of x and y must be quite large to make both the exercise variables and cancellation rate high enough to match the sample

Table 2

Model forecasts and sample averages for exercise variables and cancellation rates

Model forecasts and sample averages for option exercises at 40 firms from 1979 to 1994, and cancellation rates for a sample of 52 firms from 1984 to 1993. The American option model assumes the executive follows a market value-maximizing exercise policy. The extended American option model assumes the executive follows a market value-maximizing exercise policy subject to the possibility that he exercises or forfeits the option with an annual probability q . The utility-maximizing model assumes the executive maximizes constant relative risk-averse utility with coefficient 2. The initial level of his nonoption wealth is x and each period he is offered a payoff y to leave the firm with annualized probability q , where x and y are multiples of the initial value of shares under option. $\hat{s}_{\tau,0}$ is the predicted value of the stock price at the time of exercise, divided by the strike price, $\hat{\tau}_0$ is the predicted time of exercise in years, and $c\hat{r}_0$ is the predicted rate of annual cancellations. ESO value is the market value of the option for the representative firm. FASB value is the probability that the option vests, times the option value under the Black-Scholes model, adjusted for proportional dividends with the expiration date set equal to the option's expected life, given that it vests. Both of these option values normalize the initial stock price to one. Finally, sample averages for the predicted variables are displayed.

Model	Parametrization			Model forecasts and implied option values				
	x	y	q	$\hat{s}_{\tau,0}$	$\hat{\tau}_0$	$c\hat{r}_0$	ESO value	FASB value
American			0	3.33	7.57	0.03	0.39	0.36
Extended American			0.11	2.65	5.77	0.07	0.29	0.29
Utility-maximizing	342	132	0.12	2.67	5.87	0.07	0.29	0.28
	5	10	0.11	2.53	5.55	0.07	0.29	0.29
	1	10	0.05	2.12	4.51	0.05	0.33	0.30
	0.1	10	0.06	1.68	3.09	0.06	0.27	0.26
	4.67	0.15	0.11	2.65	5.93	0.04	0.38	0.34
	3.00	0	0	2.54	5.75	0.03	0.39	0.35
	8.18	0.30	0.20	2.49	5.39	0.06	0.35	0.32
Sample averages				2.75	5.83	0.07		

average values. The payoff for leaving is so large that the executive always takes it, so the departure decision is independent of the stock price. This base case utility-maximizing model is almost identical to the base case extended American option model.

One may argue that the high stopping rates introduce noise that masks the nuances of difference between the two models. Yet, this is essentially a result of the paper. Simply adding random, nonmarket-driven exercise and forfeiture to the standard American option model goes surprisingly far toward bringing the implied exercise patterns in line with the data. As the remainder of this paper will show, the additional effects of hedging restrictions on the executive's exercise policy are too subtle to make an incremental contribution to the description of actual option exercises.

Again, when the utility-maximizing model is parametrized to match actual exercise patterns, the typical option is worth only three quarters of its fully tradeable, American call option model value. It is also worth about three quarters of the dividend-adjusted Black–Scholes value of \$0.37. Thus, while Yermack (1995) reports that options represented about a third of the value of the average compensation of chief executive officers in 1990 and 1991, based on their Black–Scholes value, the base case models here suggest adjusting this fraction to one quarter, which is still a substantial component of total compensation.

4.3.2. *The representative option under other parametrizations*

The base case parametrizations serve as a natural starting point for comparing the two models. However, the calibration presented above incorporates some simplifications. In particular, the calibration matches each model's mean cancellation rate, across all stock price paths, to the sample average cancellation rate at 52 firms over the period 1984–1993. If the stock price performance at those firms over the period is atypical, the sample average cancellation rate may be a biased estimate of the unconditional mean cancellation rate. To the extent that this leads to incorrect parametrizations, it may not be fair to compare the models under the base case parametrization alone.

Firm-specific forecasts of the exercise variables in Section 4.4 reveal that the base case parametrizations of the utility-maximizing and extended American option models are almost indistinguishable. To be sure that the utility-maximizing model cannot generate superior forecasts under another parametrization, I set a variety of other parametrizations of the utility-maximizing model against the base case extended American option model to try to detect the potential for improvement. To demonstrate how these other parametrizations alter the predictions of the utility-maximizing model, I present the characteristics of the representative option with these alternative choices of θ in the remaining rows of Table 2.

The first three of these alternative parametrizations, in rows 4–6 of Table 2, force x to take ever smaller values with y fixed at 10 and q chosen to minimize Eq. (8). These constraints tend to make exercises occur earlier and less deep in the money. The next alternative parametrization, in the seventh row of Table 2, fixes $y = 0.15$, small enough to introduce a dependency in the departure decision, and optimizes over x and q . This change reduces the option cancellation rate.

The third to last row of Table 2 forces $y = q = 0$ and optimizes over x , so that deviations from the standard option exercise policy result solely from the hedging restrictions, with no risk of departure. Note that while $x = 3.00$ makes the forecasts of the exercise variables close to the sample averages, the cancellation rate, 0.034, is much lower than the sample average of 0.073. The next to the last row of Table 2 optimizes over x and y holding q fixed at a high value of 0.2, which slightly reduces exercise times, payoffs, and cancellation rates.

4.3.3. The FASB valuation method

In addition to generating option values and forecasts of observable variables, the models also provide the inputs necessary to implement the option valuation method required in footnotes to financial statements by FASB (1995). Because this method has come under attack, I take the opportunity now to interpret the method in the context of the framework sketched in Section 3 and demonstrate that the FASB method need not overstate option value, as some researchers and practitioners have argued.

According to the new accounting standard, firms must disclose an estimate of the value of outstanding options in financial statements. The FASB recommends measuring option value using the Black–Scholes model adjusted for proportional dividends, with the expiration date of the option set equal to its expected life. This value should then be multiplied by the fraction of granted options expected to vest. The expected option life indicated by the FASB appears to mean the expected option life, conditional on the option vesting.

Suppose the option stopping time τ is independent of the stock price path in the sense that the distribution of S_τ , given that $\tau = t$, is just the unconditional distribution of S_t . Let $c(t)$ be the value of an ordinary European call on the same stock with expiration date t : $c(t) = E(\zeta_t(S_t - S_0)^+)$. Then the executive option value, $E(\zeta_\tau(S_\tau - S_0)^+ 1_{\{\tau \geq t_v\}})$, reduces to $E(c(\tau)|\tau \geq t_v)P(\tau \geq t_v)$. By contrast, the FASB value is $c(E(\tau|\tau \geq t_v))P(\tau \geq t_v)$, which switches the order of the expectation and the call function operators. The FASB calculation differs from the correct value only because the function c is nonlinear in t . Thus, up to nonlinearity in the call price as a function of time to expiration, the FASB's method is correct if the stopping time is independent of the stock price path.

There is no reason to believe that an executive's optimal stopping time should be independent of the stock price path, but it is also not clear that this value is biased. Marcus and Kulatilaka (1994) claim that the tendency for earlier exercises to take place at higher stock prices causes the FASB value to overstate true option value. Clearly this is potentially false for a dividend-paying stock. For example, in the case of constant volatility, dividend rate, and interest rate, the standard, value-maximizing exercise policy for an American option prescribes exercising the option if the stock price rises above a critical stock price, and that critical level decreases as expiration approaches. Thus, earlier exercises will be at higher stock prices under this policy, and, since it is the value-maximizing policy, the correct option value under this policy exceeds the option value with any deterministic stopping date. In the standard case, therefore, the FASB value would understate true option value. Even if the stock pays no dividends, it is easy to construct examples in which the correct option value exceeds the FASB value and the option is exercised according to a decreasing, time-dependent boundary of critical stock prices.

To show how the FASB value compares to the theoretically correct option value under the exercise policies prescribed by the utility-maximizing and extended American option models, I use each model in turn to determine the option's expected life given vesting, and the probability that the option vests. Because these two inputs vary with the exercise policy, the FASB option value varies with the exercise policy as well. The FASB option values under the different exercise policies appear in the last column of Table 2. Note that, in general, the FASB value is quite close to the correct option value under either model. The FASB values are slightly less, suggesting that exercise times under both models are not independent of the stock price, but instead, depend on the stock price in a way that increases the option value.

4.4. Comparison of the utility-maximizing and extended American option model forecasts

In this section, I compare each of the parametrizations of the utility-maximizing model from the last section and the base case parametrization of the extended American option model. Given a model parametrization θ , I generate forecasts for each of the 40 firms that incorporate specific information about each firm: volatility, dividend rate, vesting date, and terminal stock price. To be precise, the forecasted level of the stock price at exercise and the time of exercise for firm i , are

$$\begin{aligned} \hat{s}_{\tau,i} &= E(s_{\tau} | \text{exercise}, t_v = t_{v,i}, r = 0.07, \mu = 0.155, \\ &\sigma = \sigma_i, \delta = \delta_i, s_{10} = s_{10,i}(\theta)), \end{aligned} \quad (12)$$

and

$$\begin{aligned} \hat{\tau}_i &= E(\tau | \text{exercise}, t_v = t_{v,i}, r = 0.07, \mu = 0.155, \sigma = \sigma_i, \\ &\delta = \delta_i, s_{10} = s_{10,i}(\theta)). \end{aligned} \quad (13)$$

Fig. 1 plots time-stock price pairs for the data and the calibrated models.

4.4.1. Hypotheses

The utility-maximizing model explicitly takes into account the risk preferences and endowments of the executive, and the hedging restrictions he faces in determining his optimal policy for exercising the option. In addition, the utility-maximizing model has more parameters than the extended American option model to accommodate patterns in the data. Therefore, it is natural to expect that at least some parametrizations of the utility-maximizing model will explain more of the variation in the actual exercise variables than the extended model. I use the following criteria to assess the ability of a given model to




Fig. 1. Stock price at exercise vs. time of exercise: Data and base case model forecasts. Scatter plots of pairs (τ, s_τ) , where τ is the time of option exercise and s_τ is the stock price at exercise. The circles represent actual values of these two variables for options from each of 40 firms. The black disks represent forecasts of the variables for each firm from the extended American option model. The gray disks represent forecasts of the variables from the utility-maximizing model. The extended American option model assumes the executive follows an exercise policy that maximizes the option's market value subject to the possibility that he automatically exercises or forfeits the option with some fixed probability. The utility-maximizing model assumes the executive follows an exercise policy that maximizes his expected utility. The executive has constant relative risk aversion. He invests his nonoption wealth in a constant proportion portfolio of the stock and riskless asset. Each period, with some fixed probability, he is offered a payoff to leave the firm. The parameters for each model are chosen to make the model forecasts of these two exercise variables, as well as the annual option cancellation rate, as close as possible to the sample averages of these variables. The sample includes exercise variables at 40 firms from 1979 to 1994, and cancellation rates at 52 firms from 1984 to 1993.

explain cross-sectional variation in the exercise variables:

1. Root mean squared error in the model prediction of the level of the stock price at exercise and the time of exercise:

$$\left(\sum_{i=1}^N (s_{\tau,i} - \hat{s}_{\tau,i})^2 / N \right)^{1/2}, \quad (14)$$

$$\left(\sum_{i=1}^N (\tau_i - \hat{\tau}_i)^2 / N \right)^{1/2}, \quad (15)$$

2. R^2 in the following regressions of the actual exercise variable on the model forecast:

$$s_{\tau,i} = \alpha + \beta \hat{s}_{\tau,i} + \varepsilon_i, \quad (16)$$

$$\tau_i = \alpha + \beta \hat{\tau}_i + \varepsilon_i. \quad (17)$$

I hypothesize that the utility-maximizing model will outperform the extended American option model in two ways. First, the utility-maximizing model should be capable of producing significantly lower mean squared prediction errors than the extended American option model. Second, in a regression of the actual exercise variable on the model forecast, the utility-maximizing model should be capable of producing significantly higher R^2 's.

4.4.2. Results

Tables 3 and 4 contain measures of performance for the model forecasts of the stock price at exercise and the time of exercise, respectively. In both tables, Panel 1 presents the ordinary American option model, Panel 2 presents the extended American option model with the base case parametrization, and Panel 3 presents the various parametrizations of the utility-maximizing model listed in Table 2. The tables also present measures of model bias and results of regressions from Eq. (16) and Eq. (17). In particular, column 2 of Table 3 gives the mean error and percentage error in the model forecast of the market price at exercise, as follows:

$$\sum_{i=1}^N (s_{\tau,i} - \hat{s}_{\tau,i})/N, \quad (18)$$

$$\sum_{i=1}^N ((s_{\tau,i} - \hat{s}_{\tau,i})/\hat{s}_{\tau,i})/N. \quad (19)$$

Column 3 of Table 3 gives the mean absolute error and percentage error:

$$\sum_{i=1}^N |s_{\tau,i} - \hat{s}_{\tau,i}|/N, \quad (20)$$

$$\sum_{i=1}^N (|s_{\tau,i} - \hat{s}_{\tau,i}|/\hat{s}_{\tau,i})/N. \quad (21)$$

Column 4 of Table 3 gives the square root of the mean squared error defined by Eq. (14) as well as the square root of the mean squared percentage error:

$$\left(\sum_{i=1}^N ((s_{\tau,i} - \hat{s}_{\tau,i})/\hat{s}_{\tau,i})^2/N \right)^{1/2}. \quad (22)$$

Columns 2–4 of Table 4 give the same summary statistics of the forecast errors for the time of option exercise.

Contradicting the first hypothesis, the root mean squared errors of the extended American option model shown in Tables 3 and 4 are actually among the smallest of any of the parametrizations of the models considered. For the stock price at exercise in Table 3, the root mean squared error under the extended American option model in Panel 2 is 1.19, about the same as that of 1.17 under the base case parametrization of the utility-maximizing model in the

Table 3

Model forecasts of the stock price at exercise

Mean, mean absolute, and root mean squared values errors for alternative models in forecasting the ratio of the stock price at exercise to the strike price, and results of a regression of the actual normalized stock price at exercise on the model forecast. The forecast error for firm i is $s_{\tau,i} - \hat{s}_{\tau,i}$ and the regression equation is $\hat{s}_{\tau,i} = \alpha + \beta \hat{s}_{\tau,i} + \varepsilon_i$, where $s_{\tau,i}$ is the actual average normalized stock price at exercise for firm i , and $\hat{s}_{\tau,i}$ is the model forecast given firm i 's stock return volatility, dividend rate, vesting date, and terminal stock price. The American option model assumes the executive follows a market value-maximizing exercise policy. The extended American option model assumes the executive follows a market value-maximizing exercise policy subject to the possibility that he exercises or forfeits the option with an annual probability q . The utility-maximizing model assumes the executive maximizes constant relative risk averse utility with coefficient 2. The initial level of his nonoption wealth is x and each period he is offered a payoff y to leave the firm with the annualized probability q , where x and y are multiples of the initial value of shares under option.

Parameter setting	Forecast errors (percentage errors)			Regression coefficients (standard errors)		
	Mean	Mean absolute	Root mean squared	α	β	R^2
<i>Panel 1: American option model</i>						
$q = 0$	-0.26 (0.00)	1.16 (0.36)	1.71 (0.47)	2.04 (0.58)	0.24 (0.18)	0.04
<i>Panel 2: Extended American option model</i>						
$q = 0.11$	0.42 (0.19)	0.76 (0.34)	1.19 (0.50)	0.02 (0.59)	1.18 (0.24)	0.38
<i>Panel 3: Utility-maximizing model</i>						
$x = 342, y = 132, q = 0.12$	0.42 (0.19)	0.75 (0.33)	1.17 (0.49)	-0.04 (0.58)	1.20 (0.24)	0.40
$x = 5, y = 10, q = 0.11$	0.40 (0.19)	0.76 (0.33)	1.20 (0.50)	-0.20 (0.66)	1.26 (0.27)	0.36
$x = 1, y = 10, q = 0.05$	0.57 (0.27)	0.87 (0.40)	1.34 (0.60)	0.33 (0.71)	1.11 (0.31)	0.25
$x = 0.1, y = 10, q = 0.06$	0.99 (0.54)	1.12 (0.61)	1.56 (0.81)	-1.18 (0.85)	2.23 (0.47)	0.37
$x = 4.67, y = 0.15, q = 0.11$	0.20 (0.09)	0.84 (0.32)	1.27 (0.46)	0.27 (0.80)	0.97 (0.30)	0.22
$x = 3.00, y = 0, q = 0$	0.13 (0.06)	0.88 (0.32)	1.30 (0.45)	0.62 (0.78)	0.81 (0.29)	0.18
$x = 8.18, y = 0.30, q = 0.20$	0.33 (0.13)	0.83 (0.33)	1.26 (0.48)	-0.12 (0.82)	1.18 (0.33)	0.26

first row of Panel 3, and lower than that under any other parametrization. Although the significance of the difference between 1.19 and 1.17 is not formally tested, it is clear that this difference is not meaningful. The root mean squared errors for the time of exercise in Table 4 tell the same story. None of the

Table 4

Model forecasts of the time of exercise

Mean, mean absolute, and root mean squared values errors for alternative models in forecasting the time of exercise, and results of a regression of the actual time of exercise on the model forecast. The forecast error for firm i is $\tau_i - \hat{\tau}_i$, and the regression equation is $\tau_i = \alpha + \beta \hat{\tau}_i + \varepsilon_i$, where τ_i is the actual average time of exercise for firm i and $\hat{\tau}_i$ is the model forecast given firm i 's stock return volatility, dividend rate, vesting date, and terminal stock price. The American option model assumes the executive follows a market value-maximizing exercise policy. The extended American option model assumes the executive follows a market value-maximizing exercise policy subject to the possibility that he exercises or forfeits the option with an annual probability q . The utility-maximizing model assumes the executive maximizes constant relative risk averse utility with coefficient 2. The initial level of his nonoption wealth is x and each period he is offered a payoff y to leave the firm with the annualized probability q , where x and y are multiples of the initial value of shares under option.

Parameter setting	Forecast errors (percentage errors)			Regression coefficients (standard errors)		
	Mean	Mean absolute	Root mean squared	α	β	R^2
<i>Panel 1: American option model</i>						
$q = 0$	-1.23 (-0.12)	2.37 (0.34)	2.97 (0.40)	4.28 (1.29)	0.22 (0.18)	0.04
<i>Panel 2: Extended American option model</i>						
$q = 0.11$	0.23 (0.11)	1.95 (0.40)	2.29 (0.54)	3.33 (1.26)	0.45 (0.22)	0.10
<i>Panel 3: Utility-maximizing model</i>						
$x = 342, y = 132, q = 0.12$	0.23 (0.11)	1.92 (0.39)	2.26 (0.53)	3.19 (1.30)	0.47 (0.22)	0.10
$x = 5, y = 10, q = 0.11$	0.46 (0.14)	1.91 (0.41)	2.22 (0.54)	2.85 (1.37)	0.56 (0.25)	0.12
$x = 1, y = 10, q = 0.05$	0.89 (0.25)	2.01 (0.46)	2.39 (0.61)	3.39 (1.19)	0.49 (0.23)	0.11
$x = 0.1, y = 10, q = 0.06$	2.25 (0.71)	2.57 (0.80)	3.04 (0.99)	3.03 (1.08)	0.78 (0.29)	0.16
$x = 4.67, y = 0.15, q = 0.11$	-0.07 (0.04)	2.09 (0.37)	2.41 (0.43)	3.79 (1.32)	0.35 (0.21)	0.07
$x = 3.00, y = 0, q = 0$	-0.20 (0.02)	2.12 (0.37)	2.42 (0.43)	3.74 (1.33)	0.35 (0.21)	0.07
$x = 8.18, y = 0.30, q = 0.20$	0.39 (0.12)	2.05 (0.41)	2.34 (0.48)	3.51 (1.32)	0.43 (0.23)	0.08

calibrations of the utility-maximizing model give forecast errors that are markedly smaller than those of the extended model.

Columns 5–7 of Tables 3 and 4 contain results of the cross-sectional regressions given in Eq. (16) and Eq. (17) for the different calibrations of the utility-maximizing and extended American option models. The standard errors of the

coefficients in the tables are estimated from the cross-sectional regression and do not take into account uncertainty in the estimates of the option holder and stock return parameters.

The results of the regressions fail to support the second hypothesis. For the stock price at exercise in Table 3, the R^2 of the extended American option model is 38%, among the highest of all models. The regression for the extended American option model is also closest to a 45° line with α near zero and β near one. In terms of the regression coefficients and the R^2 's, none of the regressions with the utility-maximizing model in Panel 3 look better. For the time of exercise in Table 4, the regression lines are too flat and the R^2 's are low under all models. In terms of the regression, the utility-maximizing model with outside wealth equal to 0.1 looks better than the extended model. The regression line is steeper and the R^2 is 16%, compared with only 10% for the extended model. However, the forecast errors under this calibration of the utility-maximizing model are larger than those under the extended model, and the bias is a full two years greater.

In summary, Tables 3 and 4 demonstrate that the utility-maximizing models show virtually no improvement over the extended American option model in terms of either the regression or the size of the forecast errors. Based on these results, I conclude that the utility-maximizing model performs no better than the extended American option model. Despite the additional parameters and flexibility incorporated into the model, the extended American option model fits the data at least as well, and sometimes better. Based on comparisons of the ordinary and extended American option models in the first two panels of Tables 3 and 4, the extended model appears to offer a clear improvement over the ordinary American option model with no stopping state.

5. Choosing the stopping rate

The main contribution of this paper is to dispel the misconception that a preference-based model is necessary for valuing executive stock options. In making this point, the paper reveals that a simple extension of the ordinary American option model might be adequate.

Implementing the extended American option model involves selecting an appropriate stopping rate. The base case parametrization of Section 4.3.1 illustrates one method for choosing this rate. The purpose of this method is to calibrate the model to data on the outcomes of both exercised and canceled options. One limitation of calibrating the model to annual cancellation rates is that it involves adding cancellations across overlapping option grants, which requires an assumption about the rate at which firms issue options. My calibration makes the assumption that firms grant equal numbers of options each year.

A variable that does not involve adding cancellations across overlapping option grants is the fraction of options from a particular grant that get canceled sometime in their lives. The mean value of this random variable is the probability that the option is ever canceled. That probability is straightforward to compute from the model, even conditional on realized stock price performance.

Unfortunately, data on the fraction of options that are canceled from a given grant are not publicly available. Nevertheless, most firms possess data on the outcomes of all option grants whose expiration date has already passed. For each of the past option grants, an accountant or consultant provided with the outcomes of the options from that grant could measure the size and timing of the payoffs of exercised options, and the fraction of the options canceled. He could also compute the mean values of these variables according to the model, conditioning on the path the stock price followed over the ten years after the date of that old grant. The researcher could then choose the stopping rate that minimized the average prediction error across the past grants.

If a firm did not have data on the outcomes of past options of its own from which to estimate stopping rates, it could use similar data from other firms in its industry or even economy-wide data. Just as firms use actuarial data on human mortality rates to model the cost of future pension liabilities, insurance companies use data on past insurance claims to model the cost of future claims, and mortgage companies use data on past prepayment rates to model the cost of future prepayments, firms should also draw on all available data on past option outcomes to model the cost of newly granted options.

6. Conclusion

While a great deal of study of executive stock options concentrates on the incentive effects of this form of compensation, only a nascent literature considers the valuation of these options. This paper focuses on the cost of these options to the shareholder who can freely trade or hedge his position. This cost, or market value of the options, is important not only for its implications about the optimal contract between the manager and the firm, but also for anyone trying to value a company, such as stock analyst, merger specialist, or potential shareholder. Although the FASB has backed away from its proposal to require firms to recognize option compensation cost in earnings, investors and economists still face the problem of how to assess the value of this claim on the equity of the firm.

The market value of the option depends on its payoff, and this payoff is controlled by the executive who decides when to exercise. Therefore, a theoretical understanding of option value points to the need for an empirical study of option payoffs, or, equivalently, of option exercise patterns.

Existing models of the optimal exercise policy for an executive who cannot sell or hedge his option demonstrate that with sufficiently high risk aversion and

low wealth, the executive will exercise the option almost as soon as it gets in the money, making its value arbitrarily small. I show that such extreme behavior is not consistent with exercise patterns observed in the data. Executives hold options long enough and deep enough into the money before exercising to capture a significant amount of their potential value.

The main contribution of this paper is its evidence that a simple model can describe actual option exercises just as well as a complex preference-based model. I compare two models of the exercise policy of an executive who holds a nontransferable option on his firm's stock: a utility-maximizing model in which the executive cannot hedge the option, and a naive extension of standard option exercise theory that introduces random, exogenous exercise and forfeiture. Under calibrations of the two models that best match observed exercise patterns, the exercise policies are remarkably similar.

I explore a variety of calibrations of the models and demonstrate that the utility-maximizing model shows no improvement over the more parsimonious extension of the standard American model in explaining the cross-sectional variations in the times and payoffs of exercises in a sample of options from 40 firms. This demonstrates that once we extend standard option pricing theory with an exogenous stopping rate, we gain little more by incorporating a preference-based decision process.

One reason for this finding could be that executives have much greater ability to hedge the option position than the utility-maximizing model allows. For example, an executive can in reality sell short stocks that are highly correlated with his company's stock, or sell stock that he holds and would otherwise retain, and he can take short futures positions on a stock index to eliminate market risk. In addition, tax advantages from delaying exercise may offset the benefits of diversification. Finally, the executive may be more willing to hold the option, despite hedging restrictions, if he knows he has some control of the underlying asset process and inside information about the firm's prospects.

Appendix A.

Solving for the executive's optimal exercise policy in either of the two models is a dynamic programming problem. It involves construction of a so-called value function which gives the value of the optimized objective function in each state and time. In the extended American option model, the value function gives the option value, just as in the ordinary American option model. By contrast, in the utility-maximizing model, the value function gives the executive's optimized expected utility. Note that once the exercise policy in the utility-maximizing model is generated, the utility value function has no further relevance for determining the value of the option.

Formally, the exercise policy is a decision function $D(i, j, k)$, which indicates the exercise decision at each possible state. Let T be the option expiration date, in years, and let (i, j, k) represent the state in which i periods have elapsed, $i = 1, 2, \dots, nT$, the stock price has made j moves up, $j = 0, 1, \dots, i$, and the stopping state indicator is k . In a stopping state, $k = 1$, and otherwise $k = 0$. Let $D = 1$ if the executive leaves the option alive and 0 if his action extinguishes the option.

A.1. The extended American option model

The maximized option value at each state, $V(i, j, k)$, is the value function for this dynamic program. The value function V and the exercise policy D are defined as follows.

$$V(nT, j, k) = (u^{nT}d^{nT-j} - 1)^+, \tag{A.1}$$

$$D(nT, j, k) = 0, \tag{A.2}$$

$$V(i, j, 0) = \max(V_c(i, j), u^i d^{i-j} - 1), \tag{A.3}$$

$$D(i, j, 0) = 1_{\{V_c(i, j) > u^i d^{i-j} - 1\}}, \tag{A.4}$$

$$V(i, j, 1) = (u^i d^{i-j} - 1)^+, \tag{A.5}$$

$$D(i, j, 1) = 0, \tag{A.6}$$

$$\text{for } i = nT - 1, nT - 2, \dots, nt_v, \tag{A.7}$$

$$V(i, j, 0) = V_c(i, j), \tag{A.8}$$

$$D(i, j, 0) = 1, \tag{A.9}$$

$$V(i, j, 1) = 0, \tag{A.10}$$

$$D(i, j, 1) = 0, \tag{A.11}$$

for $i = nt_v - 1, nt_v - 2, \dots, 0$, where

$$V_c(i, j) = e^{-rh}((1 - q)(\tilde{p}V(i + 1, j + 1, 0) + (1 - \tilde{p})V(i + 1, j, 0)) + q(\tilde{p}V(i + 1, j + 1, 1) + (1 - \tilde{p})V(i + 1, j, 1))). \tag{A.12}$$

A.2. The utility-maximizing model

At each time i and state j , the value of the Merton portfolio is

$$W(i, j) = \frac{x e^{rhi}}{z^i} \left(\left(\frac{\tilde{p}}{p} \right)^j \left(\frac{1 - \tilde{p}}{1 - p} \right)^{i-j} \right)^{-1/A}, \tag{A.13}$$

where

$$z = \tilde{p} \left(\frac{\tilde{p}}{p} \right)^{-1/A} + (1 - \tilde{p}) \left(\frac{1 - \tilde{p}}{1 - p} \right)^{-1/A}. \tag{A.14}$$

The function W represents the wealth process that would be optimal for the executive in the absence of the option and the random payoff for leaving. It involves holding the stock and bond in a constant proportion determined by the mean excess return and volatility of the stock, and the risk aversion of the executive. If, at any date i , the executive has exercised or forfeited the option, left the firm, and has total wealth w to be invested in this fund until the expiration date, his expected utility is

$$g(w, i) = \frac{(w e^{r(nT-i)})^{1-A}}{1-A} z^{A(nT-i)}. \tag{A.15}$$

The exercise policy of the executive, $D(i, j, k)$, and the utility value function $v(i, j, k)$ are defined recursively as follows:

$$v(nT, j, k) = (W(nT, j) + (u^{nT} d^{nT-j} - 1)^+)^{1-A} / (1 - A), \tag{A.16}$$

$$D(nT, j, k) = 0, \tag{A.17}$$

$$v(i, j, k) = \max(v_c(i, j, k), v_e(i, j, k)), \tag{A.18}$$

$$D(i, j, k) = 1_{\{v_c(i, j, k) > v_e(i, j, k)\}}, \tag{A.19}$$

$$\text{for } i = nT - 1, nT - 2, \dots, nt_v, \tag{A.20}$$

$$v(i, j, k) = \max(v_c(i, j, k), v_f(i, j, k)), \tag{A.21}$$

$$D(i, j, k) = 1_{\{v_c(i, j, k) > v_f(i, j, k)\}}, \tag{A.22}$$

$$\text{for } i = nt_v - 1, nt_v - 2, \dots, 0,$$

where

$$\begin{aligned} v_c(i, j, 0) &= \sum_{m=1}^{nT-i} q(1-q)^{m-1} \sum_{j'=0}^m \binom{m}{j'} p^{j'} (1-p)^{m-j'} g(W(i+m, j+j')) \\ &\quad \times (1 + (u^j d^{i-j} - 1)/W(i, j)) + y, i+m \\ &\quad + (1-q)^{nT-i} g(W(i, j) + (u^j d^{i-j} - 1), i), \end{aligned} \tag{A.23}$$

$$v_e(i, j, 1) = g(W(i, j) + (u^j d^{i-j} - 1) + y, i), \tag{A.24}$$

$$v_c(i, j, k) = (1 - q)(pv(i + 1, j + 1, 0) + (1 - p)v(i + 1, j, 0)) + q(pv(i + 1, j + 1, 1) + (1 - p)v(i + 1, j, 1)), \tag{A.25}$$

and

$$v_f(i, j, k) = g(W(i, j) + ky, i). \tag{A.26}$$

The functions v_e , v_f , and v_c represent the utility value of exercising the option, forfeiting the option, and continuing with the option, respectively. The utility value to the executive of exercising in states with no departure payoff, $v_e(i, j, 0)$, incorporates the possibility of receiving that payoff in the future.

A.3. The cost of the option given an exercise policy

The exercise policy D , generated by either model, determines the cost of the option. This cost, $C(D)$, is the expected discounted value of the option’s random future payoff, using true probabilities to measure the risk of a stopping state, and risk-neutral probabilities to measure stock price risk. When D comes from the extended American option model, $C(D)$ coincides with $V(0, 0, 0)$, the value function for that model. More generally, $C(D) = c(0, 0, 0; D)$ where the function $c(i, j, k; D)$ is defined as follows:

$$c(i, j, k) = \begin{cases} (u^i d^{i-j} - 1)^+ & \text{if } D(i, j, k) = 0 \text{ and } i \geq nt_v, \\ 0 & \text{if } D(i, j, k) = 0 \text{ and } i < nt_v, \\ c_c(i, j, k) & \text{if } D(i, j, k) = 1, \end{cases} \tag{A.27}-(A.29)$$

for $i = nT, nT - 1, \dots, 0$, where

$$c_c(i, j, k) = e^{-r_h}((1 - q)(\tilde{p}c(i + 1, j + 1, 0) + (1 - \tilde{p})c(i + 1, j, 0)) + q(\tilde{p}c(i + 1, j + 1, 1) + (1 - \tilde{p})c(i + 1, j, 1))). \tag{A.30}$$

References

Agrawal, A., Mandelker, G., 1987. Managerial incentives and corporate investment and financing decisions. *Journal of Finance* 42, 823–837.

Antle, R., Smith, A., 1985. An empirical investigation of the relative performance evaluation of corporate executives. *Journal of Accounting Research* 24, 1–39.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 31, 637–654.

Brickley, J., Bhagat, S., Lease, R., 1985. The impact of long-range managerial compensation plans on shareholder wealth. *Journal of Accounting and Economics* 7, 115–129.

Cox, J., Ross, S., 1976. The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 3, 145–166.

- Cuny, C., Jorion, P., 1995. Valuing executive stock options with an endogenous departure decision. *Journal of Accounting and Economics* 20, 193–205.
- DeFusco, R., Johnson, R., Zorn, T., 1990. The effect of executive stock option plans on stockholders and bondholders. *Journal of Finance* 45, 617–627.
- Financial Accounting Standards Board, 1995. Accounting for stock-based compensation. Financial Accounting Series No. 154-C.
- Foster, T., Koogler, P., Vickrey, D., 1991. Valuation of executive stock options and the FASB proposal. *The Accounting Review* 66, 595–610.
- Harrison, J., Kreps, D., 1979. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 20, 381–408.
- Holmstrom, B., Hart, O., 1987. Theory of contracts. In: Bewley, T. (Ed.), *Advances in Economic Theory*. Cambridge University Press, New York and Melbourne, pp. 71–155.
- Huddart, S., 1994. Employee stock options. *Journal of Accounting and Economics* 18, 207–231.
- Huddart, S., Lang, M., 1996. Employee stock option exercises: An empirical analysis. *Journal of Accounting and Economics* 21, 5–43.
- Jennergren, L., Naslund, B., 1993. A comment on “Valuation of executive stock options and the FASB proposal”. *The Accounting Review* 68, 179–183.
- Kim, I., 1990. The analytic valuation of American options. *Review of Financial Studies* 3, 547–572.
- Lambert, R., Larcker, D., Larcker, D., 1989. Executive stock option plans and corporate dividend policy. *Journal of Financial and Quantitative Analysis* 24, 409–424.
- Lambert, R., Larcker, D., Verrechia, R., 1991. Portfolio considerations in valuing executive compensation. *Journal of Accounting Research* 29, 129–149.
- Marcus, A., Kulatilaka, N., 1994. Valuing employee stock options. *Financial Analysts Journal* 50, 46–56.
- Merton, R., 1969. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics* 51, 247–257.
- Merton, R., 1971. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory* 51, 373–413.
- Merton, R., 1973. Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4, 141–183.
- Rubinstein, M., 1995. On the accounting valuation of employee stock options. *The Journal of Derivatives* 3, 8–24.
- Seyhun, H., 1992. The effectiveness of the insider-trading sanctions. *Journal of Law and Economics* 35, 149–182.
- Smith, C., Zimmerman, J., 1976. Valuing employee stock option plans using option pricing models. *Journal of Accounting Research* 14, 64–357.
- Vargus, M., 1994. The implications of executives’ stock option bailout on compensation contracts: an empirical analysis. Unpublished working paper. Wharton School, University of Pennsylvania, Philadelphia.
- Yermack, D., 1995. Do corporations award CEO stock options effectively? *Journal of Financial Economics* 39, 237–269.