

Raupach, Peter

Working Paper

The Cost of Employee Stock Options

Working Paper Series: Finance & Accounting, No. 123

Provided in Cooperation with:

Faculty of Economics and Business Administration, Goethe University Frankfurt

Suggested Citation: Raupach, Peter (2003) : The Cost of Employee Stock Options, Working Paper Series: Finance & Accounting, No. 123, Johann Wolfgang Goethe-Universität Frankfurt am Main, Fachbereich Wirtschaftswissenschaften, Frankfurt a. M., <https://nbn-resolving.de/urn:nbn:de:hebis:30-16972>

This Version is available at:

<https://hdl.handle.net/10419/76928>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Cost of Employee Stock Options

Peter Raupach

Goethe University Frankfurt am Main

First version: April 11, 2003

This draft: November 13, 2003

Abstract

This paper determines the cost of employee stock options (ESOs) to shareholders. I present a pricing method that seeks to replicate the empirics of exercise and cancellation as good as possible. In a first step, an intensity-based pricing model of El Karoui and Martellini is adapted to the needs of ESOs. In a second step, I calibrate the model with a regression analysis of exercise rates from the empirical work of Heath, Huddart and Lang. The pricing model thus takes account for all effects captured in the regression. Separate regressions enable me to compare options for top executives with those for subordinates. I find no price differences. The model is also applied to test the precision of the fair value accounting method for ESOs, SFAS 123. Using my model as a reference, the SFAS method results in surprisingly accurate prices.

JEL classification: G13; J33; M41; M52

Keywords: Employee stock options; Executive stock options; Exercise Behavior; Fair value accounting; Timing risk

1 Introduction

Employee stock options (ESOs) are a popular instrument to align the interests of employees to those of owners. This paper focuses on the cost of such options to shareholders.

There are two main issues specific to ESOs that are relevant in this context. First, most grantees exercise the options considerably earlier than standard option pricing theory predicts. Second, cancellations before expiry are no less important. Underwater options,

for instance, held by an employee who is leaving the firm, are forfeited shortly after the end of the labor contract. In this way, staff and management turnover have substantial impact on the cost of ESOs.

No matter what the reasons are, all valuation models have to incorporate early exercise and cancellation. There are two main types of valuation models. Type One tries to explain why option holders follow a certain exercise pattern. I call such a model *rational*. Type Two, which I call *heuristic*, attempts to describe the stochastics of exercise and cancellation in a correct way. A proper description is fully adequate for the sole purpose of option pricing from a shareholder's perspective.¹

This study follows the heuristic approach. I present a pricing method that seeks to replicate the empirics of early exercise and cancellation as good as possible.

1.1 Why Focus on Empirics?

Of course, from a theoretical point of view it is more appealing to analyze the reasons behind early exercise and cancellations. But rational models widely rely on hard-to-determine utility concepts. The utility function and factors like borrowing constraints, size and frequency of liquidity shocks, stock ownership or initial wealth, as well as the time to retirement should play an important role in the modelling of exercise decisions – and they will be important in practice. Nevertheless, it is very difficult to get reliable information on the interaction of all of these factors.

The nature of such measurement problems is not just academic but also relevant for accounting. Suppose that the accounting standard for ESOs implements a pricing methodology that relies on a utility function. If the function is not specified in particular, there is large discretion on the reported option value left to the accountant. From that point of view, a precise definition is desirable. But a rational model also determines – explicitly or implicitly – the option's value to the employee. Thus, a rational model in the rank of an official accounting rule affects the interests of more people than a heuristic model. There is more danger of political and legal discussions on the definitions that typically result in an extreme position in favor of the party with higher bargaining power. As an example how the result could look like, remember that ESOs, granted in the context

¹Provided that the probability law of stock price movements accounts for the incentives provided by stock options.

of an IPO, may be valued at zero volatility since a historical volatility is not available – the political outcome of a much less-serious measurement problem. Such problems in mind, the Financial Accounting Standards Board (FASB) is well advised to focus on descriptive aspects of exercise and cancellation, which are much easier to be objectified.

Besides shareholders, investment banks might also be interested in prices and ESO hedging strategies that focus on empirical aspects. It is not uncommon to out-source provision and settlement of ESOs to investment banks. So the bank writes the option, being less interested in economic rationales for the option holder’s exercise decisions but a good hedge for risky obligations.

1.2 A Heuristic Pricing Model

The heuristic approach of the present study is new as it does not presume a certain exercise strategy in advance.

In a first step, a general intensity-based pricing model similar to El Karoui and Martellini [KM01] is developed. The framework is general enough to incorporate a large variety of derivatives, provided that the option payoff is well-defined at the time of exercise. It is based on the absence of arbitrage and the assumption that employees have no private information on future stock prices. Furthermore, it is assumed that the remaining unhedgeable risk is idiosyncratic to employees and not priced by shareholders.

In a second step, I show that it is possible to calibrate the model with a regression analysis of exercise rates from the empirical work of Heath, Huddart and Lang. The model generates a nearly ideal fit with all information that is captured in the regression. In this sense, the model determines ESO prices at a new degree of precision. So it may be useful as a reference for other pricing models.

Separate regressions for different employee levels enable me to investigate if options held by top executives are possibly more (or less) costly than those held by subordinates. There is no evidence of essential differences.

In a further analysis I check whether all of the path-dependent regressors involved are really necessary to make the valuation precise. By varying the sensitivity to certain exercise drivers, I look at changes in the option value. Only about the half of regressors are relevant, yet a reduced model, with all insensitive regressors removed, still includes path-dependent components.

1.3 How Accurate Is SFAS 123?

SFAS 123, the relevant standard for the accounting of ESOs, suggests a simple heuristic model that reflects early exercise and cancellation as follows: First, the dividend-adjusted Black/Scholes price is calculated with a maturity equal to the expected lifetime of the option, given that it vests. In order to correct for premature forfeiture of options, the resulting B/S price is multiplied by the probability that the option vests. I will refer to this procedure as the *SFAS method*. The FASB obviously attached importance to keeping things simple. That is desirable – the simpler the procedure, the less discretion is left to the accountant – but it raises the question whether a plain model possibly blinds out important value drivers.

For the lack of market prices, the SFAS method must be validated with reference models, in the hope of getting closer to the truth with the latter. Due to the large number of factors my model accounts for, I hope to provide good reference prices. Earlier studies, working with reference models as well, found little evidence that the SFAS method performs dramatically wrong, supposed that the reference model is true and input parameters are reliable.² So does this study. Computing SFAS prices, based on inputs that are “observed” in the world of my model, I find that the SFAS method is a robust proxy with a small downward bias.

1.4 Previous Research

As stated above, I classify approaches to modelling exercise behavior into rational and heuristic. By adjusting Black-Scholes, the – heuristic – SFAS method implicitly picks a certain exercise policy: Ignoring the (weak) concavity-in-time of the Black-Scholes price, the SFAS price is correct if options are terminated (cancelled or exercised) at some independent random time – regardless of moneyness and vesting.³ However, as Rubinstein [Rub95] argues, independency of stock price path and termination time is rather implausible for several reasons. Furthermore, it is easy to generate exercise policies that keep the SFAS inputs constant but generate a quite different payoff structure and thus different option values.

The heuristic approach of Jennergren and Näslund ([JN93] and [JN95]) is closely

²See Carpenter [Car98] and Raupach [Rau03].

³Provided that the remaining risk arising from imperfect hedging is not priced.

related to the concept of independent termination. The authors introduce an independent exponentially distributed stopping time as a proxy for option holders who leave the firm. The option, if stopped, is liquidated at its current intrinsic value. If not stopped, the option considered in [JN95] pays off only at expiry (like a European option), which allows for a nearly closed pricing formula. The model can be seen as the prototype of independent termination. The American counterpart is discussed in [JN93]: Given that the option is not stopped, the risk-neutral holder freely decides on exercise.

Rubinstein [Rub95] notes that it is difficult to get reliable estimates of relevant input factors. The option value as suggested by Rubinstein gives a rather radical lower bound of prices but is based on few (and reliable) factors. Such simple estimates are easier to be compared between firms. Yet, the question whether a stock options program has positive value, if seen as an investment in incentives, is then even harder to be answered.

Several authors have modeled the rationales behind early exercise and cancellation by a utility-maximizing behavior of restricted option holders. For instance, Kulatilaka and Marcus [KM94], Huddart [Hud94], Rubinstein [Rub95], or Hall and Murphy [HM02] assume that a representative risk-averse individual continuously decides on holding the option or exercising it and investing the proceeds in the riskless asset. Carpenter [Car98] generalizes the setting with regard to the portfolio where the proceeds of exercises are invested. She introduces additional randomness by a headhunter, occasionally turning up at the employee's and offering a new job, changing in this way the current basis for decision. Furthermore, Carpenter compares the heuristic model of Jennergren and Näslund [JN93] with her rational model. She finds that the three-parameter rational model neither fits with a sample of exercises better than the one-parameter model with independent stopping, nor has it a higher predictive power. The heuristic model gives prices strikingly similar to that of the SFAS approach, thus supporting the appropriateness of SFAS 123. Raupach [Rau03] also generalizes the Jennergren and Näslund model [JN93], focusing on a good fit with empirics. He supposes the option to be exercised at an exponentially growing or constant barrier if it has not been stopped exogenously. Like in the present paper, the model is used as a reference for the SFAS method. Resulting prices are similar. Hull and White [HW03] suggest a pricing model that is basically a particular case of [Rau03] in a binomial framework. Carr and Linetsky [CL00] generalize Jennergren and Näslund's concept of a constant hazard rate of stopping to rates depending on time and current stock price. Particular cases allow for solving parts of the evaluation analytically.

El Karoui and Martellini [KM01] develop a more general theoretical framework for the pricing of assets with uncertain time-horizon, which is based on continuous-time hazard rates as well. The concept of conditional independence, which is central in the technical framework of my paper, is closely related to the ideas of El Karoui and Martellini.

The paper is organized as follows. Section 2 develops the general pricing framework. Section 3 summarizes some empirical results of the work of Huddart and Lang. In Section 4, the model is applied to a typical ESO design and calibrated with regressions. Section 5 presents prices and tests the SFAS method. Furthermore, I determine what individual exercise drivers in the regressions are relevant for prices. Section 6 concludes. Some evidence on management turnover and a number of proofs are relegated to the appendix.

2 A General Pricing Model

This section develops a general framework which allows to derive a unique price of ESOs from arbitrage and diversification arguments. I will present a hedging strategy for a large class of derivatives that minimizes the variance of the hedging error. Provided that the remaining risk is not priced, the price of an option is then the value of the hedge. The hedging strategy also makes explicit what information is essential in order to price ESOs correctly. The model is based on ideas similar to those of El Karoui and Martellini [KM01].

2.1 Assumptions

Let a vector price process X of traded securities be given with paths in $C := C([0, T], \mathbb{R}^n)$, and a filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ on a complete probability space $[\Omega, \mathcal{F}, \mathbf{P}]$. There is a money market account paying out a constant yield rate r . Given a fixed time t , I assume every integrable, \mathcal{F}_t -measurable contingent claim to be perfectly replicable by continuous trading in X and the money market account. For instance, this holds if \mathbb{F} is augmented and (X, \mathbb{F}) is a continuous semimartingale, following a stochastic differential equation with smooth coefficients. In the sequel, I assume all random variables to be square integrable.

To specify some terms, I mean by *termination* the end of the option contract for any reason. Terminations at a positive payoff are called *exercises*, and *cancellations* otherwise. I use *forfeiture* as a synonym of cancellation.

The option payoff is defined by a mapping $f : [0, T] \times C \rightarrow \mathbb{R}^+$. If an option is terminated at time t , the holder receives cash in the amount of $f(t, X) \geq 0$. The payoff shall be uniquely determined by the path of X up to t .⁴ The set of possible times of termination is restricted to a final set⁵ $\{t_1, \dots, t_K\}$. Restricting the determinants of the payoff to t and X precludes that in-the-money options are cancelled, possibly due to explicit disciplinary clauses in the option contract.

The definition of payoff is flexible enough to cover features like vesting periods or non-exercise windows. As well, outperformance options or hurdles fit into the framework the same as path-dependent derivatives like Asian options.

Most of the assumptions could be weakened, yet I will forego generality in favor of compactness.

Following the methodology of heuristic models, I do not explicitly specify how an option holder arrives at an exercise decision. I assume that there is a random time τ with values in $\{t_1, \dots, t_K\}$, at which the option pays out $f(\tau, X)$ (the total of cancellations simply appears to be $\{f(\tau, X) = 0\}$). The joint law of τ and X is assumed to be common knowledge.

The following assumption is key to the possibility of hedging: At every time, the *current* decision on termination and the *future* development of the price process are independent. Formally, I assume

$$\mathbf{P}(\tau > t, X \in B | \mathcal{F}_t) = \mathbf{P}(\tau > t | \mathcal{F}_t) \mathbf{P}(X \in B | \mathcal{F}_t)$$

for $t \in \{t_1, \dots, t_K\}$ and Borel sets $B \in \mathcal{B}(C([0, T]))$. I call the property *conditional independence*. It is equivalent⁶ to the *K-assumption*

$$\mathbf{P}(\tau > t | \mathcal{F}_t) = \mathbf{P}(\tau > t | \mathcal{F}_T), \quad (1)$$

to be found in [KM01] or [MS79]. Conditional independence is *not* total independence of exercise decisions and X . Quite the contrary, the intuition behind is an option holder

⁴Formally, $f(t, \cdot)$ shall be measurable with $\mathcal{B}([0, t])$, the sub- σ -algebra of Borel sets in C that is generated by the natural projection $C \rightarrow C([0, t], \mathbb{R}^n)$.

⁵El Karoui and Martellini [KM01] as well as Carr and Linetsky [CL00] use continuous intensities of exercise, which seems at least partially inappropriate for ESO since the distribution of exercise time will jump at maturity and vesting dates. See Huddart and Lang [HL96, fig. 1] for empirical support. Admitting continuous-time but degenerate distributions is possible but makes the model cumbersome.

⁶A proof is found in the appendix, Section 7.2.

who might experience idiosyncratic impulses, to be intractable by the option writer, but takes all past market information into consideration when deciding whether to terminate options. Idiosyncratic impulses could be sudden liquidity need, an alternative job opportunity, serious illness and things like that. The only restriction imposed by conditional independence is that a termination in t has nothing to do with the further development of X after t . In other words, option holders do not condition exercise on private information about the future development of stock prices.⁷ Note that the path of X on $[0, t]$ may even enforce termination in t . For instance, the optimal, deterministic exercise strategy for a traded American option is covered as well. In this case, $\mathbf{P}(\tau > t \mid \mathcal{F}_t) \in \{0, 1\}$, and conditional independence is trivially given.

2.2 Hedging Strategy

The writer of an ESO – shareholders or an investment bank, servicing the claims arising from ESO exercises – has to pay $f(\tau, X)$ in τ . I assume that she finances this payoff by borrowing it from the money market at the riskless rate of interest just in τ until maturity. Doing so does not narrow her action space. Following this strategy, the option writer has to pay back $H := e^{r(T-\tau)} f(\tau, X)$ to the money market in T . To hedge this liability, she implements a replicating strategy that matures in T as well. It is quite easy to determine the variance-optimal hedge at an abstract level.

Lemma 1 *Among all payoffs that can be replicated by trading in the money market account and X , the contingent claim*

$$H^* := \mathbf{E}^{\mathbf{P}} \left(e^{r(T-\tau)} f(\tau, X) \mid \mathcal{F}_T \right) \quad (2)$$

approximates H best in $L^2(\mathbf{P})$, i.e., it minimizes the variance of the hedging error. The error has a mean of zero.

Proof. Due to \mathcal{F}_T -measurability, H^* can be replicated by trading in X and the money market. Minimal variance and zero expectation for $H^* - H$ are elementary properties of conditional expectations. ■

⁷There is empirical evidence that ESO holders process private information when deciding on exercise. For instance, Huddart and Lang [HL03] report abnormal stock returns following high rates of ESO exercise. I have to ignore this to keep things simple.

Despite its formal elegance, representation (2) of H^* is rather unmanageable. In order to derive an appropriate hedging rule, I disaggregate the random payoff across time: Set for some fixed t

$$\Pi_t := I_{\{\tau=t\}} f(t, X).$$

The sum of all Π_t returns $f(\tau, X)$. Putting this into (2) and applying conditional independence (1) yields

$$\begin{aligned} H^* &= \mathbf{E}^{\mathbf{P}} \left(e^{r(T-\tau)} \sum_t \Pi_t \middle| \mathcal{F}_T \right) = \sum_t e^{r(T-t)} f(t, X) \mathbf{E}^{\mathbf{P}} (I_{\{\tau=t\}} | \mathcal{F}_T) \\ &= \sum_t e^{r(T-t)} f(t, X) \mathbf{P}(\tau = t | \mathcal{F}_t). \end{aligned} \quad (3)$$

Hedging H by H^* can now be reinterpreted as follows.

- For every $t \in \{t_1, \dots, t_K\}$, implement a bundle of replicating strategies, each paying $f(t, X) \mathbf{P}(\tau = t | \mathcal{F}_t)$ in t .
- Aggregate the differences $f(t, X) (\mathbf{P}(\tau = t | \mathcal{F}_t) - I_{\{\tau=t\}})$ in the money market account until T .

2.3 Price

Like other authors⁸, I will assume that the risk of imperfect hedging is not priced. The option price is then immediately derived from the hedging strategy. The following two propositions give examples how the absence of a risk premium for the unhedgeable risk could be reasoned. Proposition 2 uses a CAPM argument, whereas Proposition 3 assumes that the error can be diversified away.

Proposition 2 *Suppose that X covers all assets in the market. The hedging error $H^* - H$ is then uncorrelated with each asset and every \mathcal{F}_T -measurable contingent claim. Hence, a well-diversified option writer will set the value of the hedging error equal to its discounted expected value, which is zero.*

Proof. Let Y be an \mathcal{F}_T -measurable random variable. By definition, $H^* = \mathbf{E}^{\mathbf{P}}(H | \mathcal{F}_T)$, and therefore

$$\begin{aligned} \text{cov}(H^* - H, Y) &= \mathbf{E}^{\mathbf{P}}(H^* - H)Y = \mathbf{E}^{\mathbf{P}} \mathbf{E}^{\mathbf{P}}((H^* - H)Y | \mathcal{F}_T) \\ &= \mathbf{E}^{\mathbf{P}}(H^* - H^*)Y = 0. \end{aligned}$$

⁸See Jennergren and Näslund [JN95], [JN93], or Carpenter [Car98], for example.

■

Proposition 3 *Assume that the option writer has granted ESOs to a large number of employees i . Let the exercise times τ_i follow a distribution common to each, and let furthermore all exercise decisions and X be conditionally independent.⁹ If the number of employees tends to ∞ and if, furthermore, the proportional share of the largest option package vanishes as well, the proportional hedging error vanishes in $L^2(\mathbf{P})$. In particular, let a fixed price path scenario be given from 0 to T , and let the number of options granted be equal to each employee. Then, the cash flows will accrue at an approximate density of $f(t, X) \mathbf{P}(\tau = t | \mathcal{F}_t)$ over the timeline.*

The proof immediately follows from the Strong Law of Large Numbers under $\mathbf{P}(\cdot | \mathcal{F}_t)$, and is therefore omitted.

If no premium is paid for bearing the unhedgeable remaining risk, the cost of an ESO to shareholders equals that of its hedging portfolio, the latter of which allows for application of the standard option pricing theory. Let be \mathbf{Q} the equivalent martingale measure, which is unique by the assumptions made at the beginning. By (3),

$$\text{price} = \mathbf{E}_{\mathbf{Q}} e^{-rT} H^* = \sum_t e^{-rt} \mathbf{E}_{\mathbf{Q}} [f(t, X) \mathbf{P}(\tau = t | \mathcal{F}_t)]. \quad (4)$$

Furthermore, the completeness assumptions on X from Section 2.1 ensure that

$$\mathbf{P}(\tau = t | \mathcal{F}_t) = \mathbf{Q}(\tau = t | \mathcal{F}_t),$$

i.e., the conditional probability of exercise remains unaffected by the change of measure.¹⁰

Then (4) simplifies to

$$\text{price} = \sum_t e^{-rt} \mathbf{E}_{\mathbf{Q}} [f(t, X) \mathbf{Q}(\tau = t | \mathcal{F}_t)] = \mathbf{E}_{\mathbf{Q}} e^{-r\tau} f(\tau, X). \quad (5)$$

In other words, an ESO is priced as if perfectly hedgeable or, equivalently, if τ were an \mathbb{F} -stopping time. This formula has already been used in earlier work.¹¹ In the context of this paper, however, the disaggregated representation on the right side of (4) turns out to be more useful. It provides an opportunity to directly transfer empirical evidence on termination rates into prices.

⁹Formally, the indicators $I_{\{\tau_i > t\}}$ and X are to be independent under $\mathbf{P}(\cdot | \mathcal{F}_t)$ for all t .

¹⁰See appendix, Lemma 6.

¹¹See Jennergren and Näslund [JN93], [JN95] or Carpenter [Car98], for example.

3 Drivers of Exercise Probability

Since the early 90's, when the SEC implemented the disclosure of executive stock option plans, data on stock option grants has grown considerably. Yet, precise data on exercise is still scarce. Even matching EXECUCOMP with insider trades does not enable to uniquely determine when a certain option package was exercised, not to mention terminations. In this context, the dataset collected by Steven Huddart and Mark Lang is unique. It contains detailed information on stock option grants plus exercises of about 58,000 employees from 7 publicly traded firms between 1985 and 1994. The options run over 5 to 10 years, with a majority on 10 years.

The dataset was analyzed in several articles: Huddart and Lang [HL96], [HL03], Huddart [Hud98], [Hud99], and Heath, Huddart and Lang [HHL99]. I will calibrate my option pricing model with results from this work.

In [HL96], the authors choose a *grant month* to be an observation, where a *grant* is the total of all options given to employees at one day in one firm. Every month through the lifetime of an option is a candidate for being an observation, whereas some have been eliminated: “We exclude observations for which the strike price exceeded the mid-month market price, observations after a grant was fully exercised, and observations before the first vesting date since little or no exercise would occur in those cases.” The aggregation ends up with a number of 5,060 observations. The authors do tobit and weighted OLS regressions of the option exercise rate on independent variables such as characteristics of the time series of past stock prices or factors that relate to an option's life stage such as dummies for options being recently vested or those expiring soon. Table 1 summarizes definitions of regressors and coefficients utilized in this paper. The dependent variable is called *fraction exercised*, defined as “the ratio of options exercised in a month to total options in the grant”. Note that *fraction exercised* is not a hazard rate since it does not refer to the number of options remaining from earlier terminations but the total of options granted. *Fraction exercised* corresponds to the probability of an option to be exercised in a certain month and to have been unexercised so far (instead of “given that...”). Regressions are conducted, first, for all firms in the sample¹²; second, for each firm separately; and third, for different classes of the employee level, sorted by the number of options each person is granted.

¹²The employee-owned firm “H” is left out.

With similar OLS regressions on the same database, the paper [HHL99] of Heath, Huddart, and Lang focuses on “psychological factors (...) above and beyond the rational factors considered in standard models of exercise”. Main differences between [HHL99] and [HL96] are *grant weeks* as observations instead of months, slightly different criteria for cutting out a grant week, and some different regressors. One of the findings is that whether or not the recent stock price is larger than all prices through the last year has major impact on the exercise rate. As well, several short-term returns have strong impact. Essential regression results and definitions of regressors are summarized in Table 2.

4 Calibrating the Model

This section attempts to reconcile the general pricing model with the empirical findings of Heath, Huddart and Lang. The following is key to the link between the model and empirics.

Assumption 4 *At every point through the lifetime of an option, the forecasts of the regression models from [HHL99] and [HL96] provide a sufficient statistics for the current exercise probability.*

In other words, all systematic drivers of the exercise probability are assumed to be correctly captured by the regression model. What remains – the noise of the regression model – is equal to the remaining unhedgeable exercise risk.

4.1 Representative Setting

According to the options investigated in [HHL99] and [HL96], I develop a model for American call options on one underlying stock. The empirical results have high explanatory power with respect to the response of individuals to stock price movements. Yet, a variety of 7 firms is not enough to reliably investigate the influence of firm-specific factors (like dividends, industry, firm size) or that of option characteristics (like option term, vesting rules, hurdles). For instance, a shortening of the option term will probably increase the mean exercise rate, simply because there is less time to exercise. The regression models do not directly account for the option term, however. I seek to avoid misspecification due to unconsidered firm- or option-specific factors by choosing a representative setting that is as close as possible to parameters from the sample. Hence, the model is not guaranteed to provide plausible results in other settings as well. It is thus a good signal rather than

Variable	Name Mean	Type	All	Level 0	Level 1	Level 2	Level 3
Returns	$ret_{90,45}$	OLS	-0.0111	-0.0104	-0.0123	-0.0151	-0.0172
-90 days to -45	-0.039	Tobit	-0.0107	-0.0128	-0.012	-0.0158	-0.0197
Returns 0	$ret_{45,30}$	OLS	0.0274	0.0292	0.0319	0.034	0.0347
-45 days to -30	0.0088	Tobit	0.0275	0.0325	0.036	0.0367	0.0382
Returns	$ret_{30,0}$	OLS	0.0231	0.0213	0.0245	0.0313	0.0328
-30 days to 0	0.013	Tobit	0.0257	0.0273	0.0314	0.0391	0.0416
Market-to-strike	mts	OLS	0.0104	0.0125	0.0106	0.0078	0.0061
ratio (cut at 5)	2.222	Tobit	0.0087	0.0137	0.0122	0.0103	0.008
Square of \sim	mts^2	OLS	-0.0015	-0.0022	-0.0017	-0.0011	-0.0009
	5.385	Tobit	-0.0006	-0.0017	-0.0014	-0.0009	-0.0009
Volatility	$vola$	OLS	0.004	-0.0001	0.0082	0.0112	0.0121
	0.393	Tobit	0.0039	-0.0037	0.0082	0.0118	0.0120
Fraction recently	$vest$	OLS	0.0077	0.0029	0.0082	0.0116	0.0086
vested	0.035	Tobit	0.0099	0.0033	0.0095	0.013	0.0108
Fraction of	$avail$	OLS	0.0063	0.0039	0.0057	0.0019	0.0085
grant available	0.344	Tobit	0.0076	0.0065	0.0082	0.0078	0.0175
Life left	t_{left}	OLS	-0.0002	-0.0001	-0.0004	-0.0008	-0.0009
	5.692	Tobit	-0.0004	0.0002	-0.0001	-0.0008	-0.0009
Fraction to	$canc$	OLS	0.1663	0.2632	0.2525	0.3262	0.238
be canceled	0.012	Tobit	0.1624	0.2824	0.2178	0.2571	0.1148

Table 1: *Regression estimates from Huddart and Lang [HL96, table 5, 6] “...[by employee level] of options exercised on stock price variables, options recently vested, options available, life left, and options to be canceled.” The end of an event month serves as reference time for returns. Accordingly, $ret_{90,45}$ is the log stock price return between 60 and 15 days before beginning of the event month. Further definitions from [HL96, table 4]: “The unit of observation is a grant month. Statistics are for all grants with more than ten grantees and all grant months with market-to-strike ratios in excess of one. There are 5.060 such grant months. (...) Market-to-strike ratio is the lesser of five and the ratio of the market price of the stock to the strike price of the option at the end of the exercise month. Volatility is the standard deviation of log daily stock price returns over the year prior to the grant month. Fraction recently vested is the number of options that vested in the three months prior to the exercise month expressed as a fraction of options granted for months in which the market-to-strike ratio exceeds 1.15, and zero otherwise. Fraction of grant available is the ratio of options available to be exercised (i.e., vested and unexercised) to the options granted as of the beginning of the exercise month. Life left is the number of years remaining in the option life prior to expiration. Fraction to be canceled is the number of vested, unexercised options from a grant that will be canceled in the coming three months, expressed as a fraction of the total grant.” Further definitions from [HL96, table 6; indices adapted]: “Level 0 employees were among the top 5% of employees receiving options at their company; level 1, among the next 20%; level 2, among the next 25%; and level 3, among the final 50%.”*

Variable	Mean	Standard deviation	Coefficient	t-Statistic
EXER	0.0020	0.0081	—	—
Intercept	—	—	−0.00219	213
AVAIL	0.3695	0.2224	0.00264	14.6
CANCEL	0.0101	0.0367	0.05466	33.3
VEST	0.0785	0.1130	0.00108	3.9
RATIO	0.7673	0.1834	0.00251	12.9
RETWK1	0.0081	0.0573	0.01055	14.3
RETWK2	—	—	0.01232	17.1
RETWK3	—	—	0.00491	6.9
RETWK4	—	—	0.00032	0.5
RET6MO1	0.1466	0.2619	0.00008	0.4
RET6MO2	0.0954	0.2823	−0.00075	24.8
MAX	0.2632	0.4404	0.00194	20.6
Adjusted R^2	0.2849			
Number of observations	12,145			

Table 2: *Descriptive statistics and regression from Heath, Huddart and Lang [HHL99, table 3, 4]. Definitions: “There are 12,145 weekly observations of options exercised expressed as a fraction of options granted. EXER, AVAIL, CANCEL, and VEST are the fraction of the total number of options awarded from a single grant that, relative to observation week, are as follows: exercised, available for exercise, and to be canceled within six months; and, that have vested in the prior six months, respectively. RATIO is the difference between the market price of the stock on the Monday of the observation week and the strike price, divided by the option’s Barone-Adesi and Whaley [1987] value as of the same date. RETWK1 is the return on the stock in the week prior to exercise. RET6MO1 is the return on the stock over months −7 to −2, inclusive relative to the observation week. RET6MO2 is the return on the stock over months −13 to −8, inclusive. Returns are the logarithm of the ratio of closing stock prices on the days bracketing the relevant period. MAX is a dummy variable that takes the value one if the stock price in the observation week exceeds the maximum of the daily closing stock prices computed over trading days −21 to −260, i.e., the maximum over the prior year excluding the month prior to the observation week.”*

a stringent seal of approval if the SFAS pricing method turns out to be consistent with my model in Section 5.1. Large *inconsistency*, in contrast, could disprove SFAS 123 to be appropriate.

I will determine the price of a particular call option with the following characteristics: The option runs from $t = 0$ to $T = 10$ y. It is not exercisable from grant until $V = 2$ y and fully vested afterwards. There are no further exercise restrictions such as block periods around financial statement disclosures. The option is granted at the money, with a strike price K equal to the normalized stock price $X_0 = 100$. So the payoff has the form

$$f(t, X) = I_{\{t \geq V\}} [X_t - K]^+.$$

I assume that the stock price path follows a stochastic differential equation

$$dX_t = \mu X_t dt + \sigma X_t dW_t$$

with constant coefficients, as in the Black/Scholes model. The drift $\mu = 13.5\%$, volatility $\sigma = 41.6\%$, and continuous dividend yield $\delta = 3.0\%$ are set equal to the mean value over the 7 publicly traded firms in [HL96, table 1]. The risk-free interest rate is set to $r = 7\%$.

4.2 Modelling the Probability of Termination

The pricing formula $price = \sum_t e^{-rt} \mathbf{E}_{\mathbf{Q}} [f(t, X) \mathbf{P}(\tau = t | \mathcal{F}_t)]$ is evaluated by path simulation. To do that, I have to determine the termination probability $\mathbf{P}(\tau = t | \mathcal{F}_t)$ for every drawn path X and every potential exercise time $t \in \{t_1, \dots, t_K\}$. I will specify the probability in two different ways, depending on whether the option is *exercisable* or not: The regression models in Table 1 and 2 refer only to observations where the option is in the money, where some of the options are vested, and where at least some of the options remain to be exercised. Accordingly, I define an option to be exercisable in t iff $f(t, X) > 0$ and $\mathbf{P}(\tau \geq t | \mathcal{F}_t) > 0$. The following subsections define the termination probability recursively, starting from $t = 0$ until T .

4.2.1 Exercisable Options

Given a path X and some t such that the option is exercisable, the regression returns a *crude exercise probability* according to

$$p_{\text{crude}}^*(t, X) := \alpha + \beta_1 x_1(t, X) + \dots + \beta_n x_n(t, X) \quad (6)$$

where $x_i(t, X)$ are the regressors of Table 1 and 2, respectively.¹³ Time t runs from 0 to 10 in 480 steps (denoted by Δt), the quarter of a month each.¹⁴

Huddart and Lang [HL96] do not report the intercept of the regression analysis since they are mainly interested in the identification of the drivers of exercise intensity. Yet, the general level of intensity is essential for pricing – it has strong impact on options being exercised earlier or later, even on being exercised at all. I reconstruct the intercept from mean values of the dependent variable and regressors. By taking the expectation on both sides of (6), α is eliminated, ending up with

$$p_{\text{crude}}^*(t, X) = \mathbf{E}p_{\text{crude}}^*(t, X) + \beta_1(x_1(t, X) - \mathbf{E}x_1(t, X)) + \dots + \beta_n(x_n(t, X) - \mathbf{E}x_n(t, X)). \quad (7)$$

The sample mean values from Table 1 and 2 now specify the equation in full, enabling me to make numerical calculations.

It may happen in some cases that the regression will forecast values outside of $[0, 1]$, which is not meaningful for probabilities. The model only makes sense if such cases are negligible. Keeping this in mind, I simply “cap” and “floor” the values, setting

$$p_{\text{crude}}(t, X) := 1 \wedge [p_{\text{crude}}^*(t, X)]^+,$$

where $a \wedge b = \min(a, b)$.

4.2.2 Unexercisable Options

The regression models have no explanatory power for the case that the option cannot be exercised. According to my definition, this is the case if $\mathbf{P}(\tau \geq t | \mathcal{F}_t) = 0$ or $f(t, X) = 0$. When the cumulative termination probability has reached 1 already, all subsequent probabilities clearly must be zero. If not, i.e., if $\mathbf{P}(\tau \geq t | \mathcal{F}_t) = 0$ and $f(t, X) = 0$, options are assumed to be cancelled independently of \mathcal{F}_T at a constant hazard rate λ . Independent cancellations before maturity account for the empirical fact that employees who leave the firm usually have to exercise their ESOs shortly. If an option is out of the money or unvested at that time, it is forfeited. For that reason, I set the constant cancellation rate to $\lambda = 3\%$, a value that is used by practitioners as a rule-of-thumb for the fluctuation of staff. Section 7.1 in the appendix summarizes some evidence that 3%

¹³Note that all regressors refer to the past and present of stock price paths, besides deterministic factors. Hence, $p_{\text{crude}}^*(t, X)$ is \mathcal{F}_t -measurable.

¹⁴Since the dependent variable of the regression in Table 1 is a *monthly* rate, the coefficients are divided by 4 in 6.

is a good proxy for the large group of top managers (large if weighted by the number of options granted). Obviously, the calibration of the model could be improved by a regression model that includes cancellations directly. It is likely that fluctuation rates are negatively related with firm performance, so that the option cancellation rate might be higher than 3%.

4.2.3 Recursion

Let $p(t, X) := \mathbf{P}(\tau = t | \mathcal{F}_t)$ be the (marginal) probability of termination at t . Let furthermore $p_{\text{canc}} := 1 - \exp\{-\lambda\Delta t\}$ be the hazard rate of cancellations for one step and denote by $p_{\text{remain}}(t, X) = 1 - \sum_{s \leq t} p(s, X)$ the probability of options not being terminated until t (including t). The definition

$$p(t, X) := \begin{cases} p_{\text{crude}}(t, X) \wedge p_{\text{remain}}(t - \Delta t, X) & : f(t, X) > 0 \\ p_{\text{canc}} p_{\text{remain}}(t - \Delta t, X) & : f(t, X) = 0 \end{cases} \quad (8)$$

fulfills the above requirements (with formally setting $p_{\text{remain}}(-\Delta t, X) := 1$). Capping $p_{\text{crude}}(t, X)$ with $p_{\text{remain}}(t - \Delta t, X)$ ensures that the cumulated termination probability does not exceed one. Note that p_{canc} corresponds to a hazard rate, whereas $p_{\text{crude}}(t, X)$ is an *unconditional* probability, which is believed to account for the decline of exercisable options through time in a correct way without further modifications.

The definition of $p(t, X)$ ensures $p(T, X) \leq 1$ but not $p(T, X) = 1$ in general, which conflicts with the fact that each option must either be exercised or cancelled. If the option expires out of the money, this is no problem for the determination of the price since the lack of probability occurs in cases without payment. The case that the option matures in the money is considered below in detail.

Finally, let me remark that the specification of cancellation frequencies has large impact on prices. Although such events contribute zero to expected option payoffs, more frequent cancellations make p_{remain} run off earlier, setting all subsequent probabilities to zero. Otherwise, subsequent steps could contribute valuable payoffs to expectation with positive probability. Above that, cancellations influence subsequent $p(t, X)$ via some regressors, too.

4.3 Computing

The stock price process X is approximated by a recombining binomial tree of 480 time steps. The simulation of X is conducted by drawing ups and downs over the tree under the

physical measure. Given a realization of X , the probability $p(t, X)$ is determined at each t and added to the sum of probabilities of the corresponding node X_t from earlier simulations that hit the same node. After 20.000 paths and normalization of the probabilities I obtain an approximation of the joint distribution of (τ, X_τ) under the physical measure. The moments are computed directly, whereas the pricing formula (5) is evaluated with probabilities multiplied by the Radon/Nikodym derivative of each node.

4.4 Criteria of Consistency

Of course, I was initially doubtful whether a regression model that was not designed to price an option gives plausible results in this context. Above, I stated that $\{p_{\text{crude}}^*(t, X) \in [0, 1]\}$ should hold in nearly all cases. Furthermore, the distribution of (τ, X_τ) should have moments similar to those of the sample. I introduce two further criteria of consistency, both of which concern the cumulative termination probability. I will check under what conditions the outcome of the cancellation probability is consistent with the hypothesis of a constant $\lambda = 3\%$. It has become clear already that the model is automatically correct if the total termination probability $p_{\text{cum}}(t, X) := 1 - p_{\text{remain}}(t, X)$ equals one for some $t \leq T$. The same is true if the option expires out of the money. If, possibly $p_{\text{cum}}(T, X) < 1$ on such a path, the lack can be addressed to forfeitures at maturity, which are neither recognized by the regressions nor in a causal connection with the fluctuation of employees. It does not matter whether the “lack” is removed by an extra-portion of cancellations at T or not.

Suppose now that $p_{\text{cum}}(T, X) < 1$ and $f(T, X_T) > 0$. If the assumption of a given, constant hazard rate of cancellations were perfectly true, one should observe $p_{\text{cum}}(T, X) = 1$. Since this is not the case, I compute a hypothetical cancellation rate that would fill the gap between $p_{\text{cum}}(T, X)$ and 1. It is defined as follows. Let q be the total probability of exercise for a certain path, excluding termination. As specified above, cancellations must go back to the vesting period and out of-the-money periods. Summing up the total length of periods out of the money plus the vesting period to l , the *implicit cancellation rate*

$$\lambda_{\text{impl}}(X) := -1/l \ln(q)$$

defines the first criterion of consistency. This rate is still constant in time but individual to each path. It is meaningful only on the subsample of options expiring out of the money and should take values close to the pre-specified λ if the model is calibrated well.

The second criterion of consistency has the same idea behind but the scope is narrower: Let $A := \{f(T, X_T) > 0 \text{ for all } t \in [V, T]\}$ be the total of paths that keep the option continuously in the money. If the model is miscalibrated, this should crop up most strikingly on A where cancellations are limited to the vesting period $[0, V]$. The conditional cumulative probability of termination for such paths,

$$p_{\text{cum},A} := \mathbf{E}^{\mathbf{P}}(p_{\text{cum}}(T, X) \mid A),$$

should be equal or close to one. I will refer to $p_{\text{cum},A} = 1$ as the *all-in-the-money condition*.

4.5 Final Calibration

I start with an attempt to choose empirical parameters without modification. Expectations in (7) are set equal to the sample means from Table 1, coefficients are those from Column “all”, which are estimates on the whole sample of employees. $\mathbf{E}p_{\text{crude}}^*(t, X)$ equals 0.007, the mean monthly exercise rate as reported in [HL96]. The simulated mean of $p_{\text{crude}}^*(t, X)$ is biased upward to 0.0080, which is no problem from the outset since simulated expectations of the regressors do just loosely correspond to the real-world counterparts. That higher mean value is simply the forecast of the empirical model in another situation. Many simulations of $p_{\text{crude}}^*(t, X)$ however fail to be within $[0, 1]$. There are more than 10% negative values. Furthermore, I observe $p_{\text{cum},A} = 0.78$ and $\sup p_{\text{cum}}(T, X) = 0.87$, both of them values that should plausibly approximate one. Alternative coefficients from Table 1 and 2 lead to similar results. I conclude that the overall termination probability is clearly underestimated.

As a consequence, I limit the input that derives from empirics to regression coefficients henceforth, returning from (7) to the original regression equation (6). The intercept α is not considered to be given by the regression anymore but calibrated such that the criteria of consistency as specified above are met as good as possible.¹⁵ Figure 1 shows prices for different models that arise from a variation of the intercept. Here, the (resulting) probability $\mathbf{E}p_{\text{crude}}^*(t, X)$ is mapped to the price. The left edge at 0.70% corresponds to the original calibration.

Negative exercise “probabilities” become negligible at a mean exercise probability in excess of 0.9%. They occur in less than 1% of the sample, mostly combined with low

¹⁵Alternatively, one could also recalibrate the cancellation rate λ to fill the gap between $p_{\text{cum},A} = 1$ and the observed value of 0.78. But λ were to be set to 18%, which is much too high to be associated with a rate of staff turnover. Besides, resetting λ could not correct the problem of negative $p_{\text{crude}}^*(t, X)$.




Figure 1: *Option prices, depending on the expectation of the monthly termination probability, determined under the physical measure of the pricing model. By variation of the regression intercept, different such expectations are obtained as well as different prices. The graph maps the expectation (“Mean Exercise Frequency”) to the price. The left edge at 0.70% corresponds to the first specification of section 4.5. “HL96 OLS - all” denotes prices for OLS regressions on the full sample in [HL96] (see Table 1). “HL96 TOBIT - all” is the counterpart with tobit coefficients. Prices of “HHL99 OLS” result from the regression in [HHL99] (see Table 2). Price curves for regressions over subsamples according to the employee level are very similar. “Amer. Call” denotes the unique price an unrestricted investor would pay under assumptions of standard option pricing theory.*




Figure 2: Option prices depending on $p_{\text{cum},A}$, the cumulative termination probability, given that the option is continuously in the money from vesting to maturity. Pairs of $p_{\text{cum},A}$ and price are obtained by variation of the regression intercept; $p_{\text{cum},A}$ is determined under the physical measure of the pricing model. The categories “OLS” and “TOBIT” refer to the corresponding regression coefficients from [HL96]; see Table 1. “HHL99” denotes prices arising from coefficients of Table 2, which stem from [HHL99]. The appendices “All” or “Level i ” refer to subsamples of different employee level, sorted by the number of options a person received. Level 0 are the top 5% of employees; Level 1, among the next 20%; Level 2, among the next 25%; and Level 3, among the final 50%.

payoffs.

By virtue of its rigor, I choose the all-in-the-money condition as the primary benchmark of consistency. As a preliminary analysis for that step, Figure 2 shows option prices as a function of $p_{\text{cum},A}$ instead of the mean exercise frequency as in Figure 1. Note that point “1” in the abscissa refers to the *lowest* α that entails $p_{\text{cum},A} = 1$. For higher α , the probability clearly stays constant at 1, while the price may decrease in α , as is seen in Figure 1.

Switching from the mean exercise rate as common attribute of calibration to $p_{\text{cum},A}$ reduces most of the price differences between the models. Given some $p_{\text{cum},A} \in [0.7, 1]$, the maximal deviation of an individual model’s price from the mean over models is below 3.5%.

In order to determine the final option value, the intercept α is now chosen to be the least value such that the all-in-the-money condition is met. The right edge of Figure 2 thus provides the option value of each model. In adapting α this way, I change the original model as cautious as possible. Yet, higher α would result in consistent models as well, except that the mean exercise frequency were even more distant from the empirical mean.

The implicit cancellation rate $\lambda_{\text{impl}}(X)$ shows mean values between 3.5% and 5.9%, which I consider to be in a reasonable scope.

5 Results

Table 3 summarizes prices and characteristics of the joint distribution of (τ, X_τ) for different model set-ups. Characteristics such as mean values under the physical measure and the correlation of τ and X_τ are rather consistent across models. They roughly correspond with their empirical counterparts.¹⁶

Prices according to different regression models are quite similar in general. Option holders capture about 72% of the value of a corresponding American option.¹⁷ Furthermore, options under the OLS models are – slightly, but systematically – less expensive than under the tobit models. They show surprisingly low differences across employee levels within each OLS / tobit model class. Models of rational option holders such as Hall and Murphy [HM02] or Carpenter’s utility maximizing model [Car98] provide many factors that potentially generate a difference in the cost of an option held by a top executive as opposed to a floor manager. Primary candidates for such factors are the utility function, in combination with initial wealth, stock ownership, the degree of diversification, borrowing constraints, or the stochastic size of liquidity shocks relative to wealth. My model suggests that these factors cancel each other out. Yet, there are effects that could generate a systematic dependency on employee level – under the same regression coefficients, but in another context. For instance, Huddart and Lang [HL96] conclude from a higher sensitivity to historical volatility that lower-level employees are more risk-averse on average than top executives. Simulated volatilities, however, vary much less than in reality, which turns the sensitivity to volatility into a fixed effect. If I was introducing

¹⁶The empirical characteristics are recalculated in part from other figures in the descriptive statistics of [HL96]. Hence, they are unreliable. For details, see the description of Table 3.

¹⁷Calculated according to the procedure of Barone-Adesi and Whaley [BAW87]. A vesting period of two years can be neglected for the given stock price process.

Regression Type	Employee Level	$\mathbf{E}p$	Price	SFAS Error	p_{cum}^{ex}	$\mathbf{E}\tau$	$\mathbf{E}X_\tau$	ρ	$\mathbf{E}\lambda_{imp}(X)$
96-OLS	All	1.04%	37.25	-1.3%	51.7%	6.1	3.07	14.2%	5.60%
96-OLS	Lev. 0	1.43%	36.76	-3.4%	59.9%	5.5	2.59	13.6%	3.50%
96-OLS	Lev. 1	1.16%	37.46	-2.2%	56.0%	6.0	2.89	14.4%	4.10%
96-OLS	Lev. 2	1.05%	38.50	-1.3%	52.8%	6.8	3.28	11.0%	4.50%
96-OLS	Lev. 3	1.09%	37.62	-1.2%	54.6%	6.3	3.05	14.4%	4.50%
96-TOB	All	1.14%	38.61	-5.5%	50.4%	5.9	2.98	-2.0%	5.90%
96-TOB	Lev. 0	1.10%	38.70	-4.7%	51.5%	6.2	3.02	2.5%	5.20%
96-TOB	Lev. 1	1.12%	38.48	-5.0%	51.4%	6.0	2.97	2.6%	5.40%
96-TOB	Lev. 2	1.13%	39.26	-5.4%	51.4%	6.3	3.05	-4.5%	5.10%
96-TOB	Lev. 3	1.12%	37.45	-3.7%	53.5%	5.7	2.85	10.8%	5.30%
HHL99	All	1.16%	38.63	-4.8%	56.6%	6.1	2.82	4.9%	3.80%
directly from the sample			—	—	—	≈ 6.7	2.22	$\approx 20\%$	—
Price _{SFAS} (sample inputs)			37.94						
Price _{American}			51.21						

Table 3: *Prices and selected characteristics of the probability law of termination arising from different regression models. The intercept of all regression equations was taken to be the least value such that the cumulative probability of termination, given the option was completely in the money, just equals one. “96” refers to the coefficients of [HL96], to be found in Table 1, whereas “HHL99” denotes those from [HHL99]. For the meaning of the employee level subsamples, see Figure 2. The lower the figure, the higher the level of the employees. All characteristics are expectations under the physical measure of the corresponding pricing model – except price. $\mathbf{E}p$ is the mean monthly exercise probability. “SFAS Error” denotes the proportional deviation of the option price according to SFAS 123 from the model price, given that the latter is true: SFAS prices are derived from Black/Scholes prices with a maturity equal to expected option lifetime, given that it vests. The result is multiplied by the probability that the option vests, here $\exp\{-2\lambda\}$. p_{cum}^{ex} denotes the probability of exercise over the full option lifetime. $\mathbf{E}\tau$ is the mean exercise time and $\mathbf{E}X_\tau$ the mean stock price performance at exercise, ρ denotes the correlation of τ and X_τ . All means are computed under the condition that the option is exercised. $\mathbf{E}\lambda_{imp}(X)$ is the mean implicit cancellation rate. Row 3 from below reports reference values of the descriptive statistics in [HL96, tables 1, 3, 4]. The mean exercise time is an unweighted average over firms, the correlation is privately reported by Steven Huddart. Price_{SFAS} (sample inputs) gives the SFAS price for a mean exercise time of 6.7 and a probability of vesting at $\exp\{-2\lambda\}$. Price_{American} is the option price according to the procedure of Barone-Adesi and Whaley [BAW87].*

stochastic volatility into the probability law of X , differences in the sensitivity to volatility could become more important. With all due care, one could argue as follows. Given that volatility tends to persist some time, a higher sensitivity leads the option holder to forfeit more option time value¹⁸ in turbulent times since more options are exercised just when the option has a high time value. Altogether, an introduction of stochastic volatility should reduce the option value under high sensitivity to volatility as opposed to low sensitivity.

5.1 Testing the Accounting Standard

This section tests whether the SFAS pricing method accounts for the most important factors. Recall that the SFAS method has two exercise related input parameters: the probability that an option vests and the mean lifetime of an option, given that it vests. Do these parameters capture the major part of factors influencing the “true” value? Of course, there is no such truth, just other models. If one believes, as I do, that it is important to achieve a good fit with empirics, my model is interesting since it nearly exactly replicates the empirical results of the regressions. So I validate a two-factor model with a ten- or eleven-factor model, depending on the number of regressors in use.

The usual procedure would be to estimate the exercise related inputs for both models on the same sample, and to compare the corresponding prices. To compute the SFAS price, I take the average across firm-specific exercise times from [HL96] as a proxy for the mean lifetime, given vesting. For lack of data on options cancelled before vesting, I set this probability equal to $\exp\{-V\lambda\} = \exp\{-0.06\}$, in accordance with the assumption $\lambda = 3\%$ in my model. The resulting SFAS price of 37.94 is located in the middle of the regression model prices, suggesting so far that the SFAS method is a strikingly good proxy. Of course, it must be noticed that the representative setting is quite special. For the lack of data, I cannot carry out further checks.

In a second test, I assume my model to be true. Computing the mean exercise time under my model, I get SFAS prices, the proportional errors of which are listed in Table 3. A systematic but low downward bias and very low variance add further evidence about the impressive accuracy of the SFAS method.

I also test other stock price volatilities because σ is a delicate point in my model. The intercept is recalibrated to meet the all-in-the-money condition for every volatility chosen.

¹⁸From the perspective of the “risk-neutral” option writer.




Figure 3: *Model prices and SFAS prices for different volatilities. Other parameters are kept constant at the values of the representative setting (Section 4.1). The regression model’s intercept α is fitted to fulfill the all-in-the-money condition of consistency. Prices according to SFAS 123 are computed under the assumption that my model were true. For the notation of model specifications, see Figure 1. The coefficients stem from estimates across all employee levels.*

Figure 3 shows that the pricing error is stable.

5.2 What Driver Is Relevant?

This section investigates if really all regressed exercise drivers are important for option valuation. For instance, short-term returns are suspected of being irrelevant since a positive return (increasing the exercise probability if $\beta_i > 0$) is offset by a mirror-inverted negative one (decreasing the probability) in most nodes. Because short-term returns are nearly uncorrelated with other factors, – especially with payoff – the contribution to the expected payoff of such a return is presumably a constant, regardless of its β_i .

The following modification of regression models aims at changing the sensitivity to a single regressor, while keeping the left-hand side of the regression at a constant mean value. In doing so, I seek to give an answer to the following question: “Do I need to know to what degree *single values* of a factor drive *single values* of exercise probability, or is it enough to know the impact of the factor’s *mean* on *mean* exercise probability?” If the sensitivity to a regressor appears to be irrelevant for prices, the influence of the regressor

is fully captured by a constant. In the set-up of Section 4.5, even the constant becomes irrelevant since it is superimposed by calibration of the intercept α .

The sensitivity to regressor i is changed by altering β_i . The expectation of $p_{\text{crude}}^*(t, X)$, however, shall be invariant, which is achieved by a formal rearrangement of (6) to

$$p_{\text{crude}}^*(t, X) = \mu_0 + \beta_1 (x_1(t, X) - \mu_1) + \dots + \beta_n (x_n(t, X) - \mu_n) \quad (9)$$

with $\mu_1, \dots, \mu_n \in \mathbb{R}$. The constant μ_0 is set equal to the expected exercise probability *under the simulation* (denoted by $\mathbf{E}^{\mathbf{P}} \dots$). In other words, I set $\mu_0 := \mathbf{E}^{\mathbf{P}} p_{\text{crude}}^*(t, X)$, given that α has been calibrated to fulfill the all-in-the-money condition. If, furthermore,

$$\mu_i = \mathbf{E}^{\mathbf{P}} x_i(t, X) \quad (10)$$

for $i = 1, \dots, n$, equation (9) is still equivalent to the original regression model (6). Except μ_0 , each term on the right side has expectation zero, implying that $\mathbf{E}^{\mathbf{P}} p_{\text{crude}}^*(t, X)$ is invariant to a change in β_i .

Equation (10) is not easily achieved since the regressors are interdependent: At t , the regression variable *avail*¹⁹ equals $p_{\text{cum}}(t - \Delta t, X)$, which is a function of earlier values $p_{\text{crude}}^*(s, X)$. The variable *canc* $:= \text{avail } I_{\{t \geq T - 0.25\}}$ introduces further dependencies. So all regression equations are linked, albeit not strongly since the corresponding β_i are small. Let some $\mu = (\mu_1, \dots, \mu_n) \in \mathbb{R}^n$ be given. Formally, (10) is fulfilled exactly when a fixed point of the mapping $\mu \mapsto (\mathbf{E}^{\mathbf{P}} x_1(t, X), \dots, \mathbf{E}^{\mathbf{P}} x_n(t, X))$ is found. Because that mapping is a contraction, I can start with a μ equal to the expected values from [HL96] and iterate the process of entering μ into (9), computing expectations $\mathbf{E}^{\mathbf{P}} x_i(t, X)$, re-entering them as new μ and so forth. While μ_0 is kept constant at the goal level, μ converges to a fixed point, whereas $\mathbf{E}^{\mathbf{P}} p_{\text{crude}}^*(t, X)$ tends to μ_0 . In practice, about four iterations are needed to get a sufficiently stable μ .

So far, nothing has been changed except that α from (6) is now split into pieces fulfilling (10). I am now in a position to investigate the price effect of an expectation-neutral change of coefficients β_i in (9). If the price does not react, the corresponding regressor can be left out.

Because of the obvious interaction between the market-to-strike ratio and its square, which is a regressor, too, I rearrange equation (9) once more such that the new coefficients can be interpreted as steepness and convexity of the dependence on the market-to-strike ratio (see appendix, Section 7.4).

¹⁹See Table 1 and 2.

Figure 4 shows how the price reacts to a change in β_i for the models of Table 1. Within each column, the modified coefficient takes values in a geometric sequence from $0.25\beta_i$ (left edge) via β_i (center) to $4\beta_i$ (right edge). A flat curve signals low relevance of sensitivity to the corresponding exercise driver. As expected, short-term returns play a minor role. The low impact of the sensitivity to volatility must be interpreted with care since the variance of historical volatility over one year, based on monthly returns, is too low within the model to have any effect. The irrelevance of *vest*, the fraction of options recently vested, is surprising for the moment since the basic coefficient is in the same order as that of *avail*. But the variable is unequal zero only within a quarter after vesting, so that the small number of relevant observations does not have much power. I conclude that the complexity can be downsized to considering five variables. The most important driver, *mts*, depends just on the current stock price. However, the path dependency that comes into play by *avail* cannot be resolved.

An analysis of the model from Table 2 gives similar results (Figure 5). Again, all short-term returns are negligible, as well as *vest*.

6 Conclusion

In this paper I present a new pricing model for employee stock options. The general version is able to react to a principally unlimited number of factors driving exercises and cancellations of options. I show that the model can be calibrated by regression analyses of the exercise frequency done by Heath, Huddart and Lang²⁰.

The valuation method of the accounting standard SFAS 123 is validated with my model as a reference. Given the model is true, the corresponding SFAS prices – computed with inputs gained from my model – are strikingly similar, suggesting that the SFAS method captures all essential features of exercise behavior well. Of course, the result is limited to plain call options similar to those in the regression sample.

Based on separate regressions for different employee levels, I compute option values assignable to top executives and groups of subordinates. I find no evidence that the differences in exercise behavior have implications on the option value. A further analysis shows that only a part of exercise-driving factors is essential for the determination of prices.

²⁰See [HL96] and [HHL99].




Figure 4: *Expectation-neutral variation of OLS regression coefficients β_i from [HL96]. Within each column, β_i is changed in a geometric sequence from $0.25\beta_i$ (left edge) via β_i (center) to $4\beta_i$ (right edge). For the definition of regressors and the precise meaning of the subsamples “All” ... “Level 3”, which correspond to the employee level, see Table 1. The coefficients for “mts” are in fact γ_1 and γ_2 from section 7.4 in the appendix, which can be interpreted as steepness and convexity of the dependence on the market-to-strike ratio.*




Figure 5: *Expectation-neutral variation of OLS regression coefficients β_i from [HHL99]. Within each column, β_i is changed in a geometric sequence from $0.25\beta_i$ (left edge) via β_i (center) to $4\beta_i$ (right edge). For the definition of regressors, see Table 2.*

This paper leaves room for improvement in many respects. Tailoring the regression analysis to meet the needs of my pricing model would make the results more significant. First, a logit or probit regression is more appropriate than a linear model, since its forecast is a probability from the outset. Second, cancellations are both informative and price-relevant so that the termination rate (exercise plus cancellation) should be the dependent variable. Third, I suspect option-specific factors (such as vesting time) and employee-specific factors (such as time to retirement) of being heavily price-relevant. It would be interesting to take them into consideration, in particular since some of the factors could be observed by outsiders.

Further evidence should be added to the relation between the regression model and the FASB method. It well may be that the price discrepancy between the SFAS price and the model price is larger under other conditions.

In addition to the primary goal of ESO valuation, the model is an appropriate starting point for the analysis of rent-extracting exercises, conducted by managers who possess private information on the future firm value.

7 Appendix

7.1 Some Evidence on Management Turnover

This section summarizes some empirical results on the turnover of managers in several countries, in order to support the choice of $\lambda = 3\%$ for the continuous hazard rate of cancellations. The link between cancellations and turnover relies on the fact that currently unexercisable options are typically forfeited if an option holder leaves the firm. The turnover of top executives, as collected here, might be less representative for all employee levels. Yet, the weight of options received by the top group is very large in most cases. Board members are often the only grantees.

Hadlock and Lumer [HL97] report an annual rate of 3.8% for CEOs from a sample of 259 U.S. firms. Kaplan [Kap94] compares the CEO turnover in large U.S. and Japanese firms, resulting in rates of 2.2% (Japan) and 2.9% (U.S.), provided that CEOs who enter the supervisory board are left out. I assume that they may continue to hold the options. Kang and Shivdasani [KS95] find 3.1% p.a. for Japanese firms when the turnover is corrected for executives remaining on the board. The U.S. sample of Denis, Denis and Sarin

[DDS97] yields a weighted mean rate of 7.5%, yet it is not corrected in the above sense. The same problem holds for the rate of 9.2% from Mikkelsen and Partch [MP97], where CEO turnover in unacquired U.S. firms is measured over ten years. Dahya, McConnell and Travlos [DMT02] report a forced CEO turnover at rates between 2.7% and 5% from a dataset of 470 industrial firms in the U.K.

7.2 Equivalence of Conditional Independence and K-Assumption

Lemma 5 *Conditional independence and the K-assumption are equivalent.*

Proof. Let the K-assumption hold and let B be some Borel set in C . Then, for $t \geq 0$,

$$\begin{aligned}
\mathbf{P}(\tau > t, X \in B | \mathcal{F}_t) &= \mathbf{P}(\mathbf{P}(\tau > t, X \in B | \mathcal{F}_T) | \mathcal{F}_t) \\
&= \mathbf{P}(\mathbf{1}_{\{X \in B\}} \mathbf{P}(\tau > t | \mathcal{F}_T) | \mathcal{F}_t) \\
&= \mathbf{P}(\mathbf{1}_{\{X \in B\}} \mathbf{P}(\tau > t | \mathcal{F}_t) | \mathcal{F}_t) \quad (\text{K-assumption}) \\
&= \mathbf{P}(\tau > t | \mathcal{F}_t) \mathbf{P}(X \in B | \mathcal{F}_t),
\end{aligned}$$

which means that X and τ are conditionally independent. If, conversely, the latter holds, and $\mathbf{P}(X \in B | \mathcal{F}_t) > 0$, then

$$\mathbf{P}(\tau > t | X \in B, \mathcal{F}_t) = \mathbf{P}(\tau > t | \mathcal{F}_t). \quad (11)$$

Since all Borel sets B span \mathcal{F}_T (when augmented), the left side of (11) may be extended to $\mathbf{P}(\tau > t | \mathcal{F}_T)$, which yields the K-assumption. ■

7.3 The Exercise Probability Under the Change of Measure

Lemma 6 *Let the pricing kernel $d\mathbf{Q}/d\mathbf{P}$ be integrable and \mathcal{F}_t -measurable. (This can, for instance, be achieved by sufficiently smooth coefficients and non-degenerate diffusion of the stochastic differential equation driving X , which enables $\ln(d\mathbf{Q}/d\mathbf{P})$ to be represented as a stochastic integral in terms of X .) The K-assumption or, equivalently, conditional independence implies*

$$\mathbf{P}(\tau = t | \mathcal{F}_t) = \mathbf{Q}(\tau = t | \mathcal{F}_t).$$

Proof. Let Y be some bounded, \mathcal{F}_t -measurable random variable. Then

$$\begin{aligned}
\int \mathbf{P}(\tau = t | \mathcal{F}_t) Y \, d\mathbf{Q} &= \int \mathbf{P}(\tau = t | \mathcal{F}_T) Y \, d\mathbf{Q} \quad (\text{K-assumption}) \\
&= \int \mathbf{P}(\tau = t | \mathcal{F}_T) \underbrace{Y \frac{d\mathbf{Q}}{d\mathbf{P}}}_{\mathcal{F}_T\text{-measur.}} \, d\mathbf{P} \\
&= \int I_{\{\tau=t\}} Y \frac{d\mathbf{Q}}{d\mathbf{P}} \, d\mathbf{P} \quad (\text{definition of } \mathbf{P}(\cdot | \mathcal{F}_T)) \\
&= \int I_{\{\tau=t\}} Y \, d\mathbf{Q},
\end{aligned}$$

which means that $\mathbf{P}(\tau = t | \mathcal{F}_t)$ meets the definition of $\mathbf{Q}(\tau = t | \mathcal{F}_t)$. Uniqueness of the latter entails identity. ■

7.4 Separating Steepness and Convexity for mts

Let M be the market-to-strike ratio. It enters the regression model (9) by two instances: directly, as M , and as M^2 . When the importance of the regressors is analyzed in Section 5.2, the obvious interaction of M and M^2 are not to be neglected. I rewrite (9) such that new coefficients give rise for a more intuitive interpretation. Starting with an excerpt of the right side of (9), $\beta_1(M - \mu_1) + \beta_2(M^2 - \mu_2)$, where μ_i are the expected values of M and M^2 , I set

$$f(M) := \beta_1(M - \mu_1) + \beta_2(M^2 - \mu_2) \stackrel{!}{=} \gamma_1(M - \mu_1) + \gamma_2((M - \mu_1)^2 - \bar{\mu}_2). \quad (12)$$

By taking expectation on both sides I get $\bar{\mu}_2 = \mathbf{E}^{\mathbf{P}}(M - \mu_1)^2 = \mu_2 - (\mu_1)^2$. Since (12) must hold for all M , it follows that $\gamma_2 = \beta_2$ and $\gamma_1 = \beta_1 + 2\mu_1\beta_2$. Now, $\frac{\partial}{\partial M}f(\mu_1)$, the steepness at the mean value of M , is independent of γ_2 . Given that the distribution of $M - \mu_1$ is symmetric, even the mean steepness $\mathbf{E}^{\mathbf{P}}\frac{\partial}{\partial M}f(M)$ is untouched of γ_2 , since the derivative is then antisymmetric around μ_1 . Because $\frac{\partial^2}{\partial M^2}f(M)$ is totally independent of γ_1 , it is legitimate to interpret γ_1 and γ_2 as separate coefficients of the model's steepness and convexity with regard to the market-to-strike ratio.

References

- [BAW87] G. Barone-Adesi and R. E. Whaley. Efficient analytic approximation of american option values. *Journal of Finance*, 42:301–320, 1987.

- [Car98] Jennifer N. Carpenter. The exercise and valuation of executive stock options. *Journal of Financial Economics*, 48:127–158, 1998.
- [CL00] Peter Carr and Vadim Linetsky. The valuation of executive stock options in an intensity-based framework. *European Finance Review*, 4:211–230, 2000.
- [DDS97] David J. Denis, Diane K. Denis, and Atulya Sarin. Ownership structure and top executive turnover. *Journal of Financial Economics*, 45:193–221, 1997.
- [DMT02] Jay Dahya, John J. McConnell, and Nickolaos G. Travlos. The cadbury committee, corporate performance, and top management turnover. *Journal of Finance*, 57(1):461–483, 2002.
- [HHL99] Chip Heath, Steven Huddart, and Mark Lang. Psychological factors and stock option exercise. *Quarterly Journal of Economics*, 114(2):601–627, 1999.
- [HL96] Steven Huddart and Mark Lang. Employee stock option exercises. an empirical analysis. *Journal of Accounting and Economics*, 21:5–43, 1996.
- [HL97] Charles J. Hadlock and Gerald B. Lumer. Compensation, turnover, and top management incentives: Historical evidence. *Journal of Business*, 70(2):153–186, 1997.
- [HL03] Steven Huddart and Mark Lang. Information distribution within firms: Evidence from stock option exercises. *Journal of Accounting and Economics*, 34(1-3):3–31, 2003.
- [HM02] Brian J. Hall and Kevin J. Murphy. Stock options for undiversied executives. *Journal of Accounting and Economics*, 33:3–42, 2002.
- [Hud94] Steven Huddart. Employee stock options. *Journal of Accounting and Economics*, 18:207–231, 1994.
- [Hud98] Steven Huddart. Tax planning and the exercise of employee stock options. *Contemporary Accounting Research*, 15(2):203–216, 1998.
- [Hud99] Steven Huddart. Patterns of stock option exercise in the united states. In J. Carpenter and D. Yermack, editors, *Executive Compensation and Shareholder Value*, chapter 8, pages 115–142. Kluwer Academic Publishers, 1999.

- [HW03] John Hull and Alan White. Accounting for employee stock options. Working Paper, 2003.
- [JN93] L. Peter Jennergren and Bertil Näslund. A comment on "valuation of executive stock options and the FASB proposal". *Accounting review*, 68(1):179–183, 1993.
- [JN95] L. Peter Jennergren and Bertil Näslund. A class of option with stochastic lives and an extension of the black-scholes formula. *European Journal of Operational Research*, 91:229–234, 1995.
- [Kap94] Steven N. Kaplan. Top executive rewards and firm performance: A comparison of japan and the united states. *Journal of Political Economy*, 102(3):510–546, 1994.
- [KM94] Nalin Kulatilaka and Alan J. Marcus. Valuing employee stock options. *Financial Analyst Journal*, 50:46–56, 1994.
- [KM01] Nicole El Karoui and Lionel Martellini. Dynamic asset pricing theory with uncertain time-horizon. Working paper, 2001.
- [KS95] Jun-Koo Kang and Anil Shivdasani. Firm performance, corporate governance, and top executive turnover in japan. *Journal of Financial Economics*, 38(1):29–58, 1995.
- [MP97] Wayne H. Mikkelson and M. Megan Partch. The decline of takeovers and disciplinary managerial turnover. *Journal of Financial Economics*, 44:205–228, 1997.
- [MS79] G. Maziotto and J. Szpirglas. Modele général de filtrage non linéaire et équations différentielles stochastiques associées. *Annales de l'Institut Henri Poincaré*, 15:147–173, 1979.
- [Rau03] Peter Raupach. The valuation of employee stock options - how good is the standard? Working paper, 2003.
- [Rub95] M. Rubinstein. On the accounting valuation of employee stock options. *Journal of Derivatives*, 3(1):8–24, 1995.