MQiTE WORKING PAPER SERIES

WP # 11-02

Does Stock Return Predictability Affect ESO Fair Value?

Julio Carmona, Ángel León and Antoni Vaello–Sebastià

Universitat d'Alacant

Does Stock Return Predictability Affect ESO Fair Value? *

Julio Carmona

Angel León

Antoni Vaello-Sebastià †

University of Alicante

University of Alicante

University of Illes Balears

January 17, 2012

Abstract

Executive Stock Options (ESOs) are modified American options that cannot be valued using standard methods. With a few exceptions, the literature has discussed the ESO fair value by assuming unpredictable stock returns which are not supported by the available empirical evidence. In this paper we obtain the fair value of American ESOs when stock returns are predictable and, specifically, driven by the trending Ornstein-Uhlenbeck process of Lo and Wang (1995). We solve the executive's portfolio allocation problem for a simple buy-and-hold strategy when his wealth can be distributed between a risk-free asset and a market portfolio. This problem is jointly solved with the executive's optimal exercise policy. We find that executives tend to wait longer the higher the predictability, independently of the composition of executive's asset menu. We have also analyzed the implications under the FAS123R proposals for the ESO fair value and found that, even for low autocorrelations, there is a meaningful mispricing when unpredictable returns are erroneously assumed.

JEL Classification: G11, G13, G17, G35, M52

Keywords: Executive Stock Options, Risk Aversion, Undiversification, Predictability, FAS123R.

* Angel León and Antoni Vaello-Sebastià acknowledge the financial support from the Spanish Ministry for Science and Innovation through the grants ECO2008-02599 and ECO2010-18567 respectively.

† **Corresponding author:** Antoni Vaello-Sebastià, Dpt. Economia de la Empresa, Universitat de les Illes Balears, Palma de Mallorca, 07122, Illes Balears, Spain. Tel: +34 971 172024, fax: +34 971 172 389. E-mail: antoni.vaello@uib.es.

1 Introduction

Executive Stock Options (ESOs hereafter) are typically an important part of executive compensation packages. Hall and Murphy (2002) report that in 1999, 94% of S&P 500 companies granted ESOs to their top executives, and these grants accounted for 47% of total pay of S&P 500 CEOs. ESOs are American-style stock options modified for incentive reasons. Thus, ESOs cannot be sold or transferred, although partial hedge is possible by trading correlated assets. In addition, they can only be exercised after ending the vesting period. Consequently, standard methods for valuing American options are not directly applicable and a growing literature has been searching for a solution to the issue of ESO valuation.

A common assumption in this literature is that the stochastic process driving the dynamics of the underlying stock price can be represented as a geometric Brownian motion (GBM) process. Under this assumption, the stock returns turn out to be independent and identically distributed normal variates. However, there is by now a wide agreement in recognizing that this assumption does not fit well the empirical evidence concerning time series returns. As summarized, for instance, in Taylor (2005), there are three stylized facts characterizing the returns distribution at daily frequencies. First, the distribution of returns is not normal, it has typically fat tails with a high peak (leptokurtosis). Second, the autocorrelation of daily returns (predictability) is extremely low but statistically meaningful and third, there is positive dependence between squared returns. The impact on the ESO valuation of the last issue has been analyzed, among others, in Brown and Szimayer (2008), León and Vaello-Sebastià (2009) and Yiang and Tian (2010). Nonetheless, to the best of our knowledge, modeling return autocorrelation in the stock price dynamics equation has been neglected in the literature of ESO valuation. There are two reasons for that. First, the constant drift term plays no role in the Black-Scholes (BS) formula according to the early result by Grundy (1991),¹ who states that the BS formula still holds even though the underlying asset returns follow an Ornstein-Uhlenbeck (OU) process. Second, although the predictability of stock returns is now widely accepted by financial economists, there has been a long tradition in favor of the "random walk" model for stock prices. Nonetheless, as we will show later, very low levels of autocorrelation generate significant biases in pricing ESOs.

¹See page 1049.

Chapter 2 of Campbell et al. (1997) report autocorrelations for CRSP stock returns for both equally-weighted and value-weighted indexes. They find a statistically significant positive serial correlation at the first lag, which is robust across subsamples. The weekly and monthly return autocorrelations also exhibit a positive and statistically significant value at first lag over the entire sample and for all subsamples. Poterba and Summers (1988) also find negative autocorrelations for monthly CRSP data. Taylor (2005) surveys the evidence of predictability for several return series such as indexes, equities, futures and currencies for different time horizons and finds that these autocorrelations are very small but significant.² For instance, more than 90% of the 600 autocorrelation estimates are between -0.05 and 0.05. These low but statistically significant autocorrelations imply that stock returns are predictable. Campbell and Yogo (2006) show more evidence of a predictable component in stock returns by building an efficient predictability test leading to valid inference regardless of the empirical evidence of high degree of persistence of the predictor variables (such as the dividend-price ratio, earnings-price ratio or measures of the interest rate).

This paper aims to analyze the effects of predictability in a restricted sense, namely, when return predictability comes only from their time series statistical properties. Therefore, we shall focus on univariate processes by following the same line of research as Lo and Wang (1995). They discuss the effect of predictability on the market value of European options by modeling initially a univariate trending Ornstein-Uhlenbeck (TOU) process, which is a continuous time AR(1) process. Under this setting, one can only model negative autocorrelations of returns. To model positive autocorrelations, one has to resort to multivariate TOU processes already introduced by Lo and Wang (1995). Along these lines, in a constant volatility framework, they show that the autoregressive parameter has opposite effects on option prices, with negative (positive) autocorrelations increasing (decreasing) prices. This result is a peculiarity of the TOU process and does not hold in general.³

In any case, what Lo and Wang (1995) definitely show is that predictability of stock returns affects the prices of options written on those stocks. This is remarkable because in their set-up predictability is induced by the drift, which does not enter into the option pricing

²More empirical references about the autocorrelation of returns can be seen in Taylor (2005).

³For instance, Hafner and Herwartz (2001) show that there is no such asymmetry in discrete-time when the data generating probability measure follows an AR(1)-TGARCH(1,1) structure. Thus, the sign of the autoregressive parameter is irrelevant, all that matters is its size.

formula. The rationale is that, unlike the GBM process with a constant drift underlying the BS formula, the sample variance of discretely-sampled returns is not an appropriate estimator of the instantaneous variance when returns are predictable. Now, the estimation for the instantaneous variance is obtained as the sample variance adjusted by a factor that captures predictability depending on the first-order correlation coefficient of a discrete AR(1) process,⁴ which is the exact solution to the TOU process. Finally, observe that the closed-form option valuation they obtain corresponds, in our setting, to the firm's cost of an European ESO for a risk-neutral executive.

Recent examples that can be found in the literature about the impact of stock returns predictability on the valuation of options are, among others, Liao and Chen (2006), Huang et al. (2009) and Paschke and Prokopczusk (2010). In the first paper, the stock return is modeled as a continuous-time MA(1)-type process. They show that the impact of autocorrelated returns on European option prices is significant even when the MA(1) parameter is small. Contrary to the TOU process, the autocorrelated behavior of returns comes from the diffusion term. In the second paper, the authors propose an ARMA process for the stock return and model American options by using the local risk-neutralization principle of Duan (1995) and the Least squares Monte Carlo (LSM) approach by Longstaff and Schwartz (2001). They show that the AR effect is more significant than the MA effect in option pricing. Both papers refer to an earlier paper by Jokivuolle (1998), who values European options on autocorrelated indexes under a discrete-time framework. Finally, the third paper develops a continuous time factor model of commodity prices that allows for higher-order autoregressive and moving average components. In this set-up, they obtain closed-form pricing formulas for futures and options when the price dynamics are driven by a continuous autoregressive moving average (CARMA) process.⁵

Here, we are interested in the behavior of long-term American ESOs under predictability for returns driven by the TOU process and using the algorithm of León and Vaello-Sebastià (2010), which is based on the LSM method, for the ESO valuation. As it is well known, we can distinguish among three possible ESO valuations depending on the restrictions faced by

⁴See equations (11) and (12) below.

⁵See Brockwell (2001) for a description of CARMA(p,q) models. Notice that the TOU process is, in fact, a CARMA(1,0) process.

the holder. The subjective value for the restricted executive, the objective value (the firm's ESO cost) for the issuing firm, which is unrestricted but obliged to follow the executive's exercise policy, and the market (risk-neutral) value for an unrestricted holder. This value is just the ESO objective value for a risk-neutral executive. Although we are only interested on the implications of predictability in stocks returns for the objective valuation of ESOs, we will also analyze the subjective valuation since this will allow us to infer the executive's exercise policy required to obtain the objective value.

The rest of the paper is organized as follows. Section 2 presents the setting for ESO valuation. We distinguish between the case in which the ESO cost is driven by only the stock price, or the one state-variable (S1) model, which corresponds to the framework put forward initially by Hall and Murphy (2002), and the case in which the ESO cost is also affected by the availability of a market portfolio, or the two state-variable (S2) model, introduced by Cai and Vijn (2005). Section 3 presents our results concerning the S1 model. In this section we perform an extensive sensitivity analysis examining the size and sign of the bias that is incurred when the objective ESO value is obtained by assuming a GBM process for the stock price. In Section 4, we include a market portfolio as an additional investment for the executive's wealth and show the differences in the results because of its introduction. In Section 5, we apply our framework to evaluate the bias induced by the computation of the firm's ESO cost using the FAS123R method. Section 6 concludes. We have also included three appendices corresponding to Section 2. In Appendix A, we show the exact solution of the TOU process and its properties. Appendix B, we provide a detailed description of the algorithm mentioned above for valuing both, objective and subjective ESOs. Finally, Appendix C summarizes the optimization and numerical details to obtain the ESO subjective valuation.

2 Model description

To analyze how predictability affects the firm's ESO cost, we need to obtain previously the executive's exercise policy. This exercise policy is characterized, as usual, by a threshold price and it comes from solving the ESO subjective valuation problem. Then, by using the risk-neutral measure and the threshold price, the ESO objective value follows. See, for instance,

Kulatilaka and Marcus (1994), Hall and Murphy (2002) or Ingersoll (2006).

2.1 The set-up

We assume an environment in which the executive has an initial wealth, W_0 , of which a proportion α is held in restricted stocks of the company that cannot be sold until the ESO maturity. The larger α , the more related is the executive's wealth to the firm's market value and so, the higher the idiosyncratic risk supported. The remainder of the initial wealth, $(1 - \alpha)W_0$, is distributed between the market portfolio and the risk-free asset in proportions η and $1 - \eta$ respectively. This composition is decided at time 0, the grant date, and held unchanged until the end of the planning period, given by the ESO maturity, at time T . In other words, we assume the executive follows a buy-and-hold strategy for both market asset and risk-free holdings. This is in line with Barberis (2000), who also incorporates the effects of predictability. Of course, a more realistic situation would allow η to vary across the investment horizon period. But, since we also need to find at the same time the executive's exercise policy, such a strategy becomes hard to compute and the implementation of both situations are beyond the scope of our study.

Next, the executive is granted with an amount of N at-the-money ESOs with an exercise price of K . Thus, the executive wealth at maturity is conditioned on the ESO exercise date. Let us denote by $W_{T|t}$ the executive's wealth at maturity conditional on the ESO exercise at time $t \leq T$. Then, if the ESOs are not exercised until maturity, the executive's terminal wealth is

$$W_{T|T} = \alpha W_0 e^{\bar{q}T} \frac{S_T}{S_0} + (1 - \alpha) W_0 \left(\eta \frac{M_T}{M_0} + (1 - \eta) e^{r_f T} \right) + N(S_T - K)^+, \quad (1)$$

where S_T and M_T denote the firm's stock and market portfolio prices at date T . The yearly stock dividend yield, denoted as $\bar{q}$, is assumed to be reinvested in restricted stocks.⁶ The main point here is that the executive does not receive any payment from restricted stocks until T . Finally, r_f denotes the yearly risk-free rate.

Given our assumed buy-and-hold strategy for the executive's portfolio, if the ESOs are

⁶For simplicity, this model assumes that the investment in the market portfolio is done through an index fund, which does not pay dividends since it reinvests all dividend proceeds.

exercised at time $t < T$, he will invest the ESO proceeds according to the same initial combination of market portfolio and risk-free asset. Thus,

$$\begin{aligned} W_{T|t} = & \alpha W_0 e^{\bar{q}T} \frac{S_T}{S_0} + (1 - \alpha) W_0 \left(\eta \frac{M_T}{M_0} + (1 - \eta) e^{r_f T} \right) + \\ & + N(S_t - K)^+ \left(\eta \frac{M_T}{M_t} + (1 - \eta) e^{r_f(T-t)} \right). \end{aligned} \quad (2)$$

Notice that, despite the ESO exercise, the executive must maintain his restricted stocks until T . We shall assume that M_t follows a GBM process driven by

$$dM_t = \mu_M M_t dt + \sigma_M M_t dZ_{M,t}, \quad (3)$$

while $p_t \equiv \ln S_t$ follows a trending Ornstein-Uhlenbeck (TOU) process:

$$dp_t = (\theta - \kappa(p_t - \theta t)) dt + \sigma dZ_{S,t}, \quad (4)$$

where $\kappa \geq 0$ and $dZ_{S,t} dZ_{M,t} = \rho dt$. Notice that the TOU process includes the GBM process as a particular case when $\kappa = 0$. We can rewrite equation (4) as

$$d(p_t - \theta t) = -\kappa(p_t - \theta t) dt + \sigma dZ_{S,t}. \quad (5)$$

Note that, when p_t deviates from its trend θt , it is pulled back at a rate κ , the speed of mean reversion, proportional to its deviation.⁷

The executive is endowed with a standard power utility function defined over his terminal wealth:

$$U(W_T) = \begin{cases} \frac{W_T^{1-\gamma}}{1-\gamma} & \text{if } \gamma \neq 1 \\ \ln(W_T) & \text{if } \gamma = 1, \end{cases} \quad (6)$$

where $\gamma > 0$ is the coefficient of relative risk-aversion.

The executive maximizes the expected utility of his terminal wealth allocating his outside wealth between the market portfolio and the risk-free asset to their optimal sizes, denoted as η^* and $1 - \eta^*$ respectively, and selecting his optimal ESO exercise date, t^* . Therefore, the

⁷The exact discrete-time solution for equation (5) and its properties are described in Appendix A.

executive's maximum expected utility is given by

$$U^* \equiv \max_{\eta, t} \mathbb{E}_0 [U(W_{T|t})]. \quad (7)$$

In short, we assume an asset allocation framework for a buy-and-hold investor (the executive) regarding both risk-free and market asset holdings.

2.2 Model implementation

The methodology used to solve the program described by equation (7) is a modified version of the LSM algorithm of Longstaff and Schwartz (2001) for pricing path dependent derivatives. It consists on two stages. In the first stage, equation (7) is solved for a grid of values of η . This yields the executive's exercise policy, defined by a unique threshold price, for each η in the grid. Observe that we also obtain the corresponding expected exercise date, τ , for each η in the grid. Appendix B provides a detailed description of how these values are obtained. Then, in a second stage, we select the corresponding values for both the market portfolio share and the threshold price denoted, respectively, as η^* and S^* , yielding the maximum executive's expected utility. Finally, after finding S^* , we can solve for the ESO objective value, V^{OBJ} , using the risk-neutral measure.

We next obtain the executive's subjective valuation using the certainty-equivalence principle already introduced by Lambert et al. (1991). It identifies the subjective ESO value with the amount of cash that, delivered at the grant date and invested until ESO maturity, reports the same expected utility to the executive than holding the ESO. This is obtained as follows. Suppose that, at the grant date, a non-restricted amount of cash CE is delivered in place of each ESO. Then, the total executive's wealth at maturity is

$$W_T^{CE} = \alpha W_0 e^{\bar{q}T} \frac{S_T}{S_0} + [(1 - \alpha)W_0 + N \cdot CE] \left(\eta \frac{M_T}{M_0} + (1 - \eta)e^{r_f T} \right). \quad (8)$$

Thus, to calculate the subjective value we only need to find the amount CE that provides the same expected utility than holding an amount N of ESOs. Therefore, as in Tian (2004)

, CE solves the following equation:

$$\max_{\eta} \mathbb{E}_0 [U(W_T^{CE})] = U^*. \quad (9)$$

where U^* is given in equation (7). Notice that, for this maximization program, the optimal market portfolio share in executive's wealth will be, in general, different from that obtained when the executives chooses optimally the ESOs exercise date (t^*) and the proportion of unrestricted wealth to be allocated into the market portfolio (η^*).⁸

Although the paper is not concerned with the properties of the executive's subjective valuation under predictability of stock returns, we can illustrate some of them with the help of Figure 1. In this figure we have plotted the ratio of the executive's subjective valuation when returns are predictable, CE_{TOU} , with respect to the case of no predictability driven by a GBM process, CE_{GBM} . In both cases, the executive's unrestricted wealth can be allocated only in the risk-free asset. Clearly, as one should expect, this ratio increases with the predictability of returns. Furthermore, the ratio is higher the higher the degree of risk-aversion or the higher the weight of restricted stocks on executive's wealth.

2.3 Discrete-time return for the TOU process

To implement our model we use the exact discrete-time solution of equation (5). The discrete-time return at time t for a time period of length h is defined as

$$r_{t,h} \equiv \ln(S_t) - \ln(S_{t-h}). \quad (10)$$

Then, the mean of $r_{t,h}$ equals θh . When h is measured as a fraction of a year, θ measures the annual mean of log-return. Similarly, the variance, σ_r^2 , and the autocorrelation at lag one, $\rho_r(1)$, of the discrete-time returns are related to the speed of mean reversion according to the following expressions:

⁸We are indebted to an anonymous referee for pointing this to us. Further details about the optimization are explained in Appendix C.

Figure 1: Behavior of the subjective ESO value as a function of predictability, κ .

This figure plots the ratio between the subjective value when the stock price is driven by a TOU process, CE_{TOU} , and the corresponding one when the stock price is driven by a GBM process, CE_{TOU} , as a function of the speed of adjustment parameter, κ . Other parameters used for the TOU process are $\theta = 0.10$ and $\sigma = 0.30$. The remaining parameters describing the executive environment are those detailed at the beginning of Section 3.

$$\sigma_r^2 = \frac{\sigma^2}{\kappa} (1 - e^{-\kappa h}), \quad (11)$$

$$\rho_r(1) = -\frac{1}{2} (1 - e^{-\kappa h}), \quad (12)$$

where to simplify notation, the dependence of the former definitions on h has been omitted. Now, as equations (11) and (12) make clear, when κ increases, σ_r^2 decreases but $\rho_r(1)$ increases in absolute value. Finally, observe that the value $\kappa = 0$ nests the case of a GBM process for the stock price, for which autocorrelations of log-returns are zero at all lags.

3 The one state-variable framework

In this section we focus on V^{OBJ} , that is obtained by using the risk-neutral measure to value an American call option for which the threshold price is defined by the executive's exercise policy. We proceed in two stages. At first, we assume executives will have no availability to a market portfolio to allocate their unrestricted wealth, so that, the parameter η is restricted

to be equal to zero. We have denoted this framework as the *S1* model, since the stock price is the only state-variable to obtain V^{OBJ} . We consider this simplified version of the model because the computational cost is much lower (we do not have to find η^*) and the main results are also valid in the more general case. Nonetheless, in Section 4 we extend our framework to include the availability of the market portfolio, in addition to the risk-free asset and the firm's stock, giving place to the *S2* model to get V^{OBJ} .

Following the benchmark case by Hall and Murphy (2002), we consider a representative executive that owns five million dollars in initial wealth and is granted with 150,000 stock options issued at-the-money with $K = \$1$. We maintain as the true data generating process (DGP) for the stock price the TOU process, as described in equation (5). We explore how V^{OBJ} changes as a result of changes in each of the three parameters characterizing the above process. Our benchmark (yearly) values for these parameters are: $\kappa = 0.50$ for the speed of adjustment, $\theta = 0.10$ for the long-run drift coefficient, and $\sigma = 0.30$ for the diffusion coefficient. Other parameter values are $r_f = 0.06$ and $\bar{q} = 0$.

According to equation (12) and for our benchmark value of $\kappa = 0.50$, $\rho_r(1)$ is equal to -0.020 in monthly terms ($h = 1/12$), -0.005 in weekly terms ($h = 1/52$) and -0.001 in daily terms ($h = 1/360$). This low value of $\rho_r(1)$ can be found in the empirical evidence, see Table 4.8 in Taylor (2005). Although it may seem low, it will help us to highlight, in a later section, the mispricing that can be incurred in the computation of V^{OBJ} if a standard GBM process, that exhibits independent returns, is postulated erroneously as the true DGP for the stock price.

Now, we are mainly concerned with the consequences of the evidence of stock return predictability on V^{OBJ} . A greater predictability is, in the present setting, a higher size of $\rho_r(1)$. Since the autocorrelations for stock returns found in empirical studies are generally so low, the GBM process might be a good approximation for the DGP of stock prices. However, under the true DGP these small autocorrelations come from a mean reversion process such as the TOU process in equation (5).

All this suggests that we are interested in analyzing how changes in $\rho_r(1)$, and hence changes in κ , affect V^{OBJ} . Therefore, to control for the change in the volatility of the discrete-time returns, σ_r , resulting from a change in κ , we have also adjusted the diffusion

coefficient, σ , in order to keep σ_r unchanged. To be sure that nothing is lost by imposing this restriction, we have also examined the case in which σ is not adjusted, so that changes in κ also modify σ_r . It has turned out that the differences between both cases, in terms of objective valuation, average exercise dates and other features of interest, have not been significant. Thus, we focus on the case in which σ is adjusted so as to keep σ_r constant.

Considering the above adjustment, a higher κ implies only an increase in the size of $\rho_r(1)$. The range of values for κ , between 0 and 1, implies a range of first-order autocorrelations between 0 and 0.04 in monthly terms, which are representative of the values found in the empirical literature cited previously. For the diffusion coefficient, σ , the values range from $\sigma = 0.25$ to $\sigma = 0.45$, and for the long-run drift coefficient, θ , the values range from $\theta = 0.065$ to $\theta = 0.15$.

3.1 Discussion of results

We study now the impact of changes in κ , σ and θ in the ESO valuation. Figure 2 summarizes our main findings concerning the behavior of V^{OBJ} as a function of these parameters. In all simulations we have considered the case of a 10-year ESO with a vesting period of 3 years, which is common in both practice and literature. The no vesting case is qualitatively similar and it is omitted to save space.

Several features can be observed in this figure. First, as κ increases, which means a higher predictability, V^{OBJ} seems to converge independently of the size for γ and α (panel A). Second, an increase in either σ (panel B) or θ (panel C) supposes a higher value of V^{OBJ} . Thus, as one would expect, the objective valuation increases with the diffusion coefficient, σ , and this is independent of the values considered for γ and α . Clearly, a higher σ implies a higher σ_r , which leads to higher stock prices in relatively shorter periods of time.

However, the positive relationship between θ and V^{OBJ} may be somewhat surprising. According to the Black-Scholes model, the expected returns are irrelevant for option pricing. Nonetheless, in the present framework θ affects the executive's exercise behavior, and hence the ESO objective valuation. With greater expected returns, a higher θ , it will be more profitable for the executive to hold the ESOs for a longer period. Then, higher expected returns could mitigate the executive's suboptimal early exercise.

Figure 2: Behavior of the objective ESO value in the S1 framework.

We show the behavior of the objective ESO's value for a 10-year ESO with a vesting period of 3 years, V^{OBJ} , is represented as a function of the speed of adjustment, κ , in panel A; the diffusion coefficient, σ , in panel B and the long-run drift, θ , in panel C. In each panel we have considered two values for the degree of relative risk-aversion, $\gamma = \{2, 4\}$, and for the degree of undiversification, $\alpha = \{1/2, 2/3\}$. The other values for the parameters of the TOU process in equation (5) are: $\theta = 0.10$ and $\sigma = 0.30$ for panel A; $\kappa = 0.50$ and $\theta = 0.10$ for panel B, and $\kappa = 0.50$ and $\sigma = 0.30$ for panel C. In all cases, we have set $r_f = 0.06$ and $\bar{q} = 0$.

Table 1: Average exercise dates (in years) and predictability.

γ	α	κ					
		0.0	0.05	0.10	0.25	0.50	1.00
2	1/2	7.41	7.42	7.78	8.14	8.31	8.50
	2/3	6.33	6.79	7.27	7.80	8.18	8.43
4	1/2	5.01	5.71	6.30	7.19	7.75	8.23
	2/3	4.58	5.22	5.73	6.94	7.58	8.13

This table collects the results concerning how the average exercise dates, τ^* , of our 10-year ESO are affected by different values of mean reversion, κ , which drives the autocorrelation of the TOU process in equation (5), respecting different levels of γ and α . The column for $\kappa = 0$ presents the values of τ^* for the corresponding equivalent GBM process as described in Appendix A. For all cases, the risk-free rate, the dividend yield, the ESO maturity and the length of the vesting period have been set to their benchmark values, namely, $r_f = 0.06$, $\bar{q} = 0$, $T = 10$ and $\nu = 3$.

By definition, V^{OBJ} depends on the executive's exercise policy. This exercise policy can be described in terms of either the threshold price, S^* , or the expected exercise time, τ^* . We next examine the former results in connection with the obtained values for τ^* . The relevant results are presented in Tables 1 and 2.

As Table 1 shows, higher values of κ imply higher values of τ^* for all considered combinations of γ and α . Of course, either a higher γ or α makes the executive to exercise comparatively earlier but, as κ increases, these differences become smaller. Furthermore, the expected exercise time is comparatively lower for the equivalent GBM process,⁹ for which $\kappa = 0$, in all cases.

To understand the role of predictability in the observed behavior for V^{OBJ} and τ^* , we examine Figure 3 in which we have plotted several simulated paths for the firm's stock price under alternative values of κ .¹⁰ We have also plotted the long-run trend of the process by a discontinuous line. As it can be observed, as κ increases from 0.05 to 1.00 (i.e. $\rho_r(1)$ changes

⁹This GBM process for stock prices has been constructed in such a way that the mean and variance of the discrete-time returns are the same than the corresponding ones for the true TOU process. More details are given in Appendix A.

¹⁰To isolate the effects of predictability, the monthly volatility for discrete returns remains constant to 0.0866 in all cases and the four simulated paths in each panel are generated using the same random numbers. Then, the observed differences only depend on κ . The other relevant parameters have been set to $\theta = 0.10$ and $\sigma = 0.30$ in all simulations.

from -0.002 to -0.04), the simulated paths tend to revert faster to that long-run trend, avoiding the paths become very deep in-the-money. Then, as long as the time to maturity is kept fixed, the average values achieved by the TOU process are comparatively lower the higher κ . Hence, the executive tends to postpone the ESO exercise increasing, accordingly, the value of τ^* .

Therefore when prices are weakly mean reverting, so that they resemble closely a GBM process, the executive tends to exercise at relatively high prices. Meanwhile, as the predictability of the process increases, through a higher κ , the executive waits longer on average (see Table 1) but the average value achieved by the firm's stock price turns out to be comparatively lower. This has the effect of reducing V^{OBJ} for a relatively low risk-averse executive and increasing V^{OBJ} for higher risk-averse executives. Of course, as one would expect, when γ and α are relatively low, the executive tends to wait longer so as to get a higher threshold price which raises V^{OBJ} .

Table 2 exhibits the values of τ^* for different values of σ and θ . This table also displays the case in which the stock price is driven by an equivalent GBM process. We can appreciate that τ^* decreases regarding σ but increases according to θ . Once again, the expected exercise times are comparatively lower for the corresponding equivalent GBM processes. There is only one exception that takes place for a combination of high values of the drift parameter ($\theta = 0.15$) and low values of both the degree of relative risk-aversion parameter ($\gamma = 2$) and the proportion of restricted stocks in executive's total wealth ($\alpha = 1/2$). This fact can be explained as a result of the combination of a high value for the drift and absence of mean reversion that gives more upside potential and hence, more value to the call option for the executive. Nevertheless, this exception disappears when the degree of mean reversion is high enough ($\kappa = 1.5$).

As Table 2 shows, higher values of σ lead to lower values for τ^* , for every value of γ and α . This can be explained along the lines used to understand its positive impact on V^{OBJ} . On the other hand, a higher long-run trend θ of stock prices, makes executives to be willing to wait longer in the expectation of higher threshold prices for their ESOs.

Finally, we have also explored how the former results change with a different value for

Figure 3: Alternative paths for different TOU processes.

This figure plots alternative paths for the TOU process in equation (5) as a function of the speed of adjustment parameter, κ . The diffusion coefficient, σ , has been adjusted so as to keep constant the volatility of the discrete returns, σ_r , in all cases. The value for this volatility is 0.30 in yearly terms. The long-run drift is equal to its benchmark value, $\theta = 10\%$. In the four panels the discontinuous line represents the long-run trend as given by $\theta t + p_0 e^{-\kappa t}$. As κ increases, the simulated paths tend to revert faster to this long-run trend, although their volatilities are not altered.

Table 2: Average exercise dates (in years). The effects of σ and θ .

Panel A: Diffusion coefficient, σ											
γ	α	(GBM)					(TOU)				
		0.25	0.30	0.35	0.40	0.45	0.25	0.30	0.35	0.40	0.45
2	1/2	9.31	7.47	6.67	6.17	5.85	8.63	8.32	8.07	7.87	7.71
	2/3	7.72	6.41	5.68	5.34	5.15	8.53	8.20	7.89	7.66	7.50
4	1/2	5.53	5.04	4.80	4.70	4.64	8.20	7.77	7.48	7.26	7.10
	2/3	4.87	4.57	4.49	4.42	4.42	8.03	7.60	7.28	7.02	6.82
Panel B: Drift coefficient, θ											
		(GBM)					(TOU)				
		0.065	0.075	0.10	0.125	0.15	0.065	0.075	0.10	0.125	0.15
2	1/2	6.71	6.90	7.63	8.31	9.43	7.69	7.87	8.40	8.70	9.00
	2/3	5.77	5.95	6.57	6.94	7.58	7.52	7.72	8.28	8.60	8.91
4	1/2	4.95	4.95	5.07	5.10	5.20	7.24	7.39	7.85	8.16	8.48
	2/3	4.58	4.58	4.61	4.76	5.40	7.02	7.18	7.69	8.03	8.39

The right panels collect the results about how the average exercise dates, τ^* , of our 10-year ESO are affected by the volatility and the mean of returns of the TOU process in equation (5). For comparison purposes, the left panels present the values of τ^* for the corresponding equivalent GBM processes as described in appendix A. For all cases, the risk-free rate, the dividend yield, the time to maturity of the ESO and the length of the vesting period have been set to their benchmark values, namely, $\kappa = 0.50$, $r_f = 0.06$, $\bar{q} = 0$, $T = 10$ and $\nu = 3$.

κ .¹¹ In particular, a value of $\kappa = 0.05$ that makes TOU processes to resemble the equivalent GBM processes, produces more variability in both V^{OBJ} and τ^* as a result of changes in γ and α .

3.2 Cumulative probabilities

We have shown previously that as the predictability of stock returns increases, executives find optimal to wait longer for exercising the ESOs. This relationship has been inferred from the behavior of the expected exercise time, τ^* . However, this expectation hides some very interesting pieces of information concerning the probability of the ESO exercise. To complete our understanding of the executive behavior, we have computed the unconditional probability that ESOs will be exercised not later than some given date $t \leq T$. This has been

¹¹The results are available upon request.

done as follows. In our simulations for the sensitivity analysis, we have considered a monthly frequency for the ESO exercise dates. This means that, for a 10-year ESO, there is a total of 120 possible exercise dates. Thus, for each of these exercise dates and using a total number of 200,000 paths (including antithetics), we have computed the ratio

$$\pi_s = \frac{\# \text{ of paths exercised at time step } s}{\text{total } \# \text{ of paths}},$$

which shows the probability the ESO package will be exercised at time s . Notice this is just a conditional probability because it is conditional on not being exercised earlier than at time step s . Then, the unconditional probability that ESOs will be exercised exactly at time t , denoted as Γ_t , is

$$\Gamma_t = \pi_t \prod_{i=1}^{t-1} (1 - \pi_{t-i}).$$

Finally, the cumulative probability that the ESO exercise date will be lower than or equal to some given date t , denoted as $F(t)$, is $F(t) = \sum_{j=1}^t \Gamma_j$. To smooth out the resulting probability distribution, we have replicated the former procedure a total of 50 times using 50 different seeds to generate the random numbers, so that the reported cumulative probabilities are the average of them.

The cumulative probability distributions, $F(t)$, have been obtained for each of the relevant parameters characterizing the TOU process. Finally, to compare those results with those obtained under no predictability, we have also computed the cumulative probabilities of early exercise for each of the corresponding equivalent GBM processes.

Panel A of Figure 4 represents $F(t)$ as a function of κ . Clearly, as κ increases, this probability is lower for any date before the expiration date. Thus, in comparison with the GBM process, the executive always waits longer for exercising the ESO. Panel B depicts the case for different values of σ when the stock price is driven by either a TOU or its equivalent GBM process. Generally, a higher volatility leads to earlier exercise (lower τ^*) and specifically when the level of predictability is low, as in a GBM process. A larger value of σ means that i) the executive's wealth becomes more volatile due to the ESOs but also because of the restricted stocks, and ii) the option may be deep in-the-money with a higher probability. Both facts suppose an earlier exercise of the ESO, according to panel B findings. Finally,

Figure 4: Cumulative probabilities of early exercise.

We plot the cumulative probabilities that the exercise time is not greater than some given date t , $F(t)$, for some values of κ , σ and θ . Panel A shows the curves corresponding to $\kappa = 0$, $\kappa = 0.10$, $\kappa = 0.50$ and $\kappa = 1.0$. The remaining parameters have been set to $\theta = 0.10$ and $\sigma = 0.30$. In panel B, the curves represent the cases $\sigma = 0.30$ and $\sigma = 0.40$, where the other parameters have been set to $\kappa = 0.50$ and $\theta = 0.10$. Finally, in panel C the curves represent the cases $\theta = 0.10$ and $\theta = 0.15$ with the other parameters set to $\kappa = 0.50$ and $\sigma = 0.30$. For comparison purposes, we have also plotted the same probability for the corresponding equivalent GBM process described in Appendix A. The general finding is that, executives will tend to wait longer to exercise ESOs as the predictability of the process increases, that is, a higher value of $|\rho_r(1)|$. In all cases, we have focused on an executive with $\gamma = 4$ and $\alpha = 2/3$.

panel C shows the same situation for several values of θ and it does not deserve further comments.

3.3 The objective bias

In this subsection, we analyze the sign and size of the biases that can occur because of a misspecification of the underlying process for the stock price. As we have mentioned before, the literature has typically postulated a GBM process for the stock price. This assumption implies that the discrete-time returns will follow a white noise with drift. But, if the log-price were to follow a TOU process, as we have assumed in our computations, the discrete-time returns would follow an ARMA(1,1) process with trend. Thus, the firm's cost V^{OBJ} would be biased as a result of an erroneous choice for the stochastic process of the price dynamics. We designate this bias as¹²

$$Bias_1 = \frac{V_F^{OBJ} - V^{OBJ}}{V^{OBJ}} \times 100, \quad (13)$$

where V_F^{OBJ} denotes the objective value computed under the false hypothesis that prices are driven by a GBM process. We maintain the notation V^{OBJ} to denote the objective value computed under the true process. We have calculated this bias for different values of κ and the results are displayed in panel A of Figure 5. For $\gamma = 2$, the objective bias is generally positive and increasing with predictability. Note that the size of the bias is not high in this case, meanwhile for $\gamma = 4$ the sign of the bias turns out to be negative and its size becomes quite substantial even for moderate degrees of undiversification.

The results of performing the same analysis but for different values of σ are depicted in panel B of Figure 5. We have found the same pattern described before, namely, when the executive has a low degree of relative risk-aversion and is well diversified, V_F^{OBJ} tends to overstate the true cost but not much. As the degree of relative risk-aversion increases and the executive becomes worse diversified, V_F^{OBJ} understates the true cost in such a way that the size of the bias turns out to be above 10% for most values of σ .

We have also computed the objective bias for several alternative values of θ . The results

¹²Notice that the objective value can be computed either using a conventional binomial model or, as we do, by using the LSM method. We illustrate the size of bias using this latter method, but the results should be similar with other methods of computing V^{OBJ} .

Figure 5: Objective ESO valuation bias under the S1 framework (S1), $Bias_1$

We plot the percentage bias in equation (13), which is incurred when the firm's cost of a 10-year ESO with a vesting period of 3 years is evaluated using a GBM process for the stock price when the true one is a TOU process. Panels A to C depict this bias as a function of, respectively, κ , σ and θ , for four combinations of the relative degree of risk-aversion and the executive's degree of undiversification. In all cases, we have set $r_f = 0.06$ and $\bar{q} = 0$.

are shown in panel C of Figure 5. Along the lines of the previous results, the sign of the bias is positive when the degree of risk-aversion is low and the executive is relatively well diversified. However, as the degree of risk-aversion increases and the executive becomes less well diversified, the sign of the bias becomes negative and its size is also well above 10% for almost all cases considered.

We can conclude that for high degrees of relative risk-aversion, the objective value assuming a GBM process can understate substantially the true firm's ESO cost. Finally, the size of the bias does not appear to be significantly affected by the length of the vesting period and it has not been reported here.

4 The two state-variable framework

Here we explore how the executive's exercise policy and hence, the ESO objective value is modified by including a market portfolio. This extended set-up is assumed to provide the true objective value, V^{OBJ} . We keep the benchmark values of the previous section for describing the environment of the representative executive as well as those describing the TOU process. The parameters characterizing the dynamics of the market portfolio are $\mu_M = 0.10$ and $\sigma_M = 0.20$. The initial value of the market portfolio will be set to one and the correlation between the innovations in the stock and the market portfolio, ρ , will be set to 0.50. Basically, to obtain V_{S2}^{OBJ} we solve the executive's problem for a grid of values of η in order to obtain the value η^* which maximizes the executive's utility.¹³ This makes *S2* model more demanding in computational terms than the *S1* model, which doesn't need finding η^* .

The results are summarized in a ratio that measures the difference between the objective value in the one state-variable model, V_{S1}^{OBJ} , and the objective value in the two state-variable model, V_{S2}^{OBJ} , as a percentage of the latter. Hence,

$$Bias_2 = \frac{V_{S1}^{OBJ} - V_{S2}^{OBJ}}{V_{S2}^{OBJ}} \times 100 . \quad (14)$$

Recall that V_{S2}^{OBJ} measures the true cost of ESOs for firms. Cai and Vijh (2005) have shown that this ratio is positive when the stock price is driven by a GBM process. The reason is that,

¹³Appendix B provides more details on the numerical algorithm.

Table 3: Predictability and ESO cost: S1 versus S2 models.

ν	α	κ				
		0.05	0.10	0.25	0.50	1.00
0	1/2	13.34	8.16	1.41	-0.24	-0.62
	2/3	19.93	13.29	3.11	0.46	-0.34
3	1/2	8.99	6.80	1.01	-0.07	-0.53
	2/3	10.02	8.14	2.62	0.44	-0.31

We present here for $\gamma = 4$ the percentage bias defined in equation (14), concerning the firm's ESO cost when executives have available a market portfolio to allocate their unrestricted wealth (S2 model) against the case in which they have not (S1 model). As it is shown, the differences decrease as κ increases to become negligible for relatively high values of this parameter. For the rest of the parameters see Table 2 and Section 4.

under model S1, executives can allocate their unrestricted wealth only in the risk-free asset whereas under model S2, they can also invest in the market portfolio. Since this alternative becomes more attractive, the average exercise times will be generally lower in the S2 model.

Our results are presented in Table 3. We have restricted our attention on the impact of changes in predictability, so that we only report the results concerning changes in κ . In this regard, we have maintained our choice of adjusting σ such that the volatility of discrete-time stock returns is not affected.

For the case of $\gamma = 4$, V^{OBJ} increases with κ in both S1 and S2 models.¹⁴ For $\kappa \geq 0.25$, which implies a first order autocorrelation higher than 0.0103 in absolute value, the firm's ESO cost for both models becomes indistinguishable. To get some intuition for this result, recall that the market portfolio is driven by a GBM process. Then, when executives observe a higher predictability in the firm's stock returns, they do not see any advantage in exercising earlier to place the proceeds into the market portfolio. As a result, the average exercise time and the threshold price are essentially the same in both the S1 and the S2 models. This result becomes quite relevant to the extent that V_{S1}^{OBJ} can be obtained at a lower computational costs than V_{S2}^{OBJ} .

Finally, we have also studied how the share of the market portfolio is affected by changing

¹⁴ESO cost for firms behaves differently when $\gamma = 2$. In the S1 model, V^{OBJ} decreases with κ . In the S2 model, V^{OBJ} increases with κ . However the behavior of the percentage bias is completely analogous, namely, as predictability increases the bias decreases until being negative.

Table 4: Optimal share of market portfolio (η^*).

ν	α	κ				
		0.05	0.10	0.25	0.50	1.00
0	1/2	0.17	0.18	0.26	0.43	0.55
	2/3	0.09	0.08	0.25	0.58	0.76
3	1/2	0.18	0.18	0.26	0.43	0.55
	2/3	0.09	0.09	0.26	0.58	0.76

The table displays, for $\gamma = 4$, how η^* is affected by a higher predictability of the firm's stock returns. The volatility of stock returns in discrete-time is kept constant.

the predictability of the firm's stock returns. The results are displayed in Table 4, which focuses on the impact of different values of κ and distinguishes between the 3-year vesting and the no vesting cases. It is shown that a greater predictability leads the executive to put more wealth in the market portfolio and less into the risk-free asset. In fact, the executive puts even more weight in the market portfolio for a higher undiversification level, α , when the predictability of the stock returns is above a specific threshold (see the last two columns in Table 4). Note that this pattern is independent of the vesting period length.

We have seen that a greater predictability in stock returns leads executives to hold a higher share of their unrestricted wealth in the market portfolio and also to wait longer for exercising the ESOs. From an intertemporal point of view, this means that executives are substituting their holdings of the risk-free asset by ESOs because of their higher expected return, which is subject to less uncertainty as predictability increases.

5 The FAS 123R bias

As a result of the increasing relevance of ESOs in managers' compensation packages, the International Financial Reporting Standard (IFRS) and the Financial Accounting Standards Board (FASB) have issued several requirements for a fair valuation of those expenses. Both standards require to consider the same factors to obtain the fair value: the market price of the shares, the exercise price of the options, the risk-free interest rate, the expected volatility of stock prices, the expected dividend on stock and the number of years until options expire.

However, there are some differences between both standards. The main accountant difference concerns the classification of the award as a liability or equity. Following FAS 123R, if there exists a possibility to settle the award in cash, it is always classified as a liability award, meanwhile IFRS 2 classification is based on the method of expected settlement (cash or shares), that is, the possibility of cash settlement does not imply a liability award. Other minor differences concern the definition of grant date and the vesting period effects. For instance, the IFRS 2 establishes the grant date when the agreement is reached, but under the FAS 123R the grant date is the earlier between the mutual understanding date and the date when employee begins to provide his services. As regards the vesting period, when accounting for ESO plans with graded vesting under the "European" standard (IFRS 2), we must treat each tranche as a separate award while under the FAS 123R we can use a straight-line method for the entire award.

Notice that the above differences do affect accountant practice but not the concept of fair value. Thus, the fair value is the same under both standards, independently of which one is followed to account for the expense. Furthermore, both provide similar recommendations for computing the fair value and, as a matter of choice, we have selected FASB's recommendations to compare our model with these suggestions.

In its Financial Accounting Standard Board (2004) No 123 revised statement (FAS 123R), the FASB requires firms to disclose the method used to estimate the grant-date fair value of their ESO compensation packages. Among the valuation techniques that the FAS 123R considers acceptable, it appears as candidates both the BS and binomial models. The BS formula is appealing because of its simplicity. However, there are some features of ESOs that are not well captured by using this formula. In particular, the BS model assumes European-style options whereas ESOs are typically American options that executives tend to exercise before maturity, see Bettis et al. (2005). To consider this fact, the FAS 123R (paragraph A26) explicitly requires computing this fair value by replacing the ESO expiration date, T , with its expected exercise time, or expected life, τ^* . Specifically, the ESO price should be calculated as

$$V^{FAS} = BS(\tau^*). \quad (15)$$

This procedure leaves firms with the problem of estimating τ^* . Indeed, we can distinguish

two possibilities in this regard. First, the firm can use historical data about the executives' exercise behavior. Alternatively, the early exercise can be described using some binomial model that, by capturing the executives' exercise policy, determines the ESO's expected life, the required input in the BS model.

Kulatilaka and Marcus (1994), argue that the FAS123 proposal may yield a biased estimation of the firm's ESO costs if historical data are used to estimate τ^* , because these historical data are only one particular realization of stock prices. However, Carpenter (1998) compares the performance of three alternatives models, an utility-maximizing model, an extended American option model, in which she introduces an exogenous stopping rate, and the FAS 123 proposal. She obtains that the three alternatives yield similar values for the ESO cost when the models are calibrated to yield the same executive's exercise policy. Ammann and Seiz (2004) also compare similar alternative models to compute the ESO fair value. They clearly show that the expected life of the ESO is the key variable underlying the models, and that pricing differences become negligible when the models are calibrated to the same expected life of the option.

In any case, and independently of the model used to obtain the executive's exercise policy, our contention is simply that firms typically assume that stock prices are driven by a GBM process. If the true model is a TOU process, the resulting executive's exercise policy is flawed and the implied average exercise date would be biased, leading also to a biased fair value.

Therefore, it is interesting to compare the firm's ESO cost under the TOU process, V^{OBJ} , with the BS value obtained under the erroneous DGP that prices are driven by a GBM process, $BS(\tau_F^*)$, where τ_F^* denotes its corresponding expected life. Specifically, we will analyze the following percentage bias:

$$Bias_3 = \frac{BS(\tau_F^*) - V^{OBJ}}{V^{OBJ}} \times 100. \quad (16)$$

We have focused on the behavior of this bias for different values of the parameter κ . We have taken as τ_F^* the value for the case of $\kappa = 0$, which is the equivalent GBM process for the stock price. Furthermore, the volatility of the discrete-time returns, σ_r , has been kept constant by a suitable adjustment of the diffusion coefficient, σ , as in Section 3. As a result, the appropriate value of the firm's stock return volatility to be plugged into the BS formula

Figure 6: FAS123 bias. The one state-variable framework.

We plot the percentage bias defined in equation (16), which is incurred when the FAS123 procedure is used and the expected life of the ESO is computed using erroneously the GBM process instead of the TOU process for stock returns. Both TOU and GBM have been made comparable by equalizing their means and volatilities. The procedure is described in detail in Section 5. The values of the remaining parameters are: $r_f = 0.06$, $\bar{q} = 0$ and $T = 10$.

is 30% for all values of κ . In our analysis we have assumed a 10-year ESO for the cases of no vesting and a 3-year vesting period. The results are depicted in Figure 6 for the S1 model.

We can see that the bias is generally negative. Hence, the estimation of the ESO expected life using a GBM process for the stock price, when TOU is the true process, understates the true expected cost systematically. This result is related to the behavior of the expected exercise date already discussed in subsection 3.1. We have shown that executives tend to wait longer for exercising ESOs when predictability increases, so that the lowest τ^* is achieved when $\kappa = 0$. This, in turn, implies a low BS value for the ESO in comparison with the firm's ESO cost under the TOU process. We show that the size of this bias increases with the predictability of the process for $\gamma = 2$, but decreases for $\gamma = 4$. This is explained by the different behavior of V^{OBJ} for each value of the relative risk-aversion parameter.

We end this section by reporting briefly the results for the FAS bias when the S2 model is considered. By using a similar procedure to that described for the S1 model, we find that the inclusion of a market index only increases the undervaluation incurred when the FAS123 recommendations are used.

6 Conclusions

We have shown that predictability matters for valuing American ESOs from the firm's perspective. The objective value, or firm's ESO cost, is biased if one assumes erroneously that stock prices are driven by a geometric Brownian motion (GBM) instead of the true process driven by a trending Ornstein-Uhlenbeck (TOU) process. This bias is significant even for relatively low values of the first order autocorrelations. The executive performs his ESO exercise under two alternative asset menu settings. One of them consists only of the risk-free asset and the other one is extended by including also the market portfolio. When predictability increases, the pricing differences between both approaches vanish. Moreover, independently of the executive's wealth composition, he waits longer for the ESO exercise the higher the predictability. Finally, we examine the consequences of predictability for the FAS123 proposals. When the erroneous GBM process is used for prices, it generates an undervaluation of the firm's ESO cost, even for moderate low levels of predictability.

Some important extensions for future research are suggested next. First, it would be interesting to analyze the executive's asset allocation problem, when there is the possibility of reallocating the market portfolio along the planning horizon (dynamic allocation) as in Barberis (2000), jointly with holding American ESOs. The quadratic approximation for the valuation of ESOs can be a good candidate to make easier its implementation. See, for instance, Kimura (2010). Finally, it would be also interesting to introduce the executive forfeiture and the case of perpetual options into our framework. See again Kimura (2010) and also, Jennergren and Näslund (1996) among others.

References

- Ammann, M., Seiz, R., 2004. Valuing employee stock options. Does the model matter? *Financial Analysts Journal* 60 (5), 21–37.
- Barberis, N. C., 2000. Investing for the long run when returns are predictable. *The Journal of Finance* 55 (1), 225–264.
- Bergstrom, A. R., 1984. *Handbook of Econometrics*. Vol. 20. North Holland, Amsterdam, Ch. Continuous time stochastic models and issues of aggregation over time, pp. 1145–1212.
- Bettis, J. C., Bizjak, J. M., Lemmon, M. L., 2005. Exercise behavior, valuation, and the incentive effects of employee stock options. *Journal of Financial Economics* 76 (2), 445–470.
- Bouchard, B., Warin, X., 2011. Monte-Carlo valorisation of american option: facts and new algorithms to improve existing methods. *Numerical Methods in Finance*, ed. R. Carmona, P. Del Moral, P. Hu and N. Oudjane, *Springer Proceedings in Mathematics*.
- Brockwell, P., 2001. *Handbook of Statistics*. Elsevier Science, Amsterdam, Ch. Continuous-time ARMA processes.
- Brown, P., Szimayer, A., 2008. Valuing executive stock options: Performance hurdles, early exercise and stochastic volatility. *Accounting and Finance* 48, 363–389.
- Cai, J., Vihj, A. M., 2005. Executive stock and option valuation in a two state-variable framework. *Journal of Derivatives* 12 (3), 9–27.
- Campbell, J. Y., Lo, A. W., Mackinlay, A., 1997. *The Econometrics of Financial Markets*. Princeton, Princeton University Press.
- Campbell, J. Y., Yogo, M., 2006. Efficient tests of stock return predictability. *Journal of Financial Economics* 81, 27–60.
- Carpenter, J. N., 1998. The exercise and valuation of executive stock options. *Journal of Financial Economics* 48 (2), 127–158.
- Duan, J.-C., 1995. The GARCH option pricing model. *Mathematical Finance* 5, 13–32.
- Financial Accounting Standard Board, 2004. Statement of financial accounting standards n° 123 (revised 2004): Share-based payments. FASB.
- Grundy, B. D., 1991. Option prices and the underlying asset’s return distribution. *Journal of Finance* 46, 1045–1069.
- Hafner, C. M., Herwartz, H., 2001. Option pricing under linear autoregressive dynamics, heteroskedasticity, and conditional leptokurtosis. *Journal of Empirical Finance* 8, 1–34.

- Hall, B. J., Murphy, K. J., 2002. Stock options for undiversified executives. *Journal of Accounting and Economics* 33 (1), 3–42.
- Huang, Y.-C., Wu, C.-W., Wang, C.-W., 2009. Valuing American options under ARMA processes. *International Research Journal of Finance and Economics* 28, 152–159.
- Ingersoll, J. E., 2006. The subjective and objective evaluation of incentive stock options. *Journal of Business* 79 (2), 453–487.
- Jennergren, L. P., Näslund, B., 1996. A class of options with stochastic lives and an extension of the black-scholes formula. *European Journal of Operational Research* 91, 229–234.
- Jokivuolle, E., 1998. Pricing European options on autocorrelated indexes. *Journal of Derivatives* 6, 39–52.
- Kimura, T., 2010. Valuing executive stock options: A quadratic approximation. *European Journal of Operational Research* 207, 1368–1379.
- Kulatilaka, N., Marcus, A. J., 1994. Valuing employee stock options. *Financial Analysts Journal* 50 (6), 46–56.
- Lambert, R. A., Larcker, D. F., Verrecchia, R. E., 1991. Portfolio considerations in valuing executive compensation. *Journal of Accounting Research* 29 (1), 129–149.
- León, A., Vaello-Sebastià, A., 2009. American GARCH employee stock option valuation. *Journal of Banking and Finance* 33 (6), 1129–1143.
- León, A., Vaello-Sebastià, A., 2010. A simulation-based algorithm for American executive stock option valuation. *Finance Research Letters* 7, 14–23.
- Liao, S.-L., Chen, C.-C., 2006. The valuation of European options when asset returns are autocorrelated. *The Journal of Futures Markets* 26 (1), 85–102.
- Lo, A., Wang, J., 1995. Implementing option pricing models when returns are predictable. *The Journal of Finance* 50 (1), 87–129.
- Longstaff, F., Schwartz, E., 2001. Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies* 14 (1), 113–147.
- Moreno, M., Navas, J. F., 2003. On the robustness of least-squares Monte Carlo (LSM) for pricing American derivatives. *Review of Derivatives Research* 6 (2), 107–128.
- Paschke, R., Prokopczusk, M., 2010. Commodity derivative valuation with autoregressive and moving average components in the price dynamics. *Journal of Banking and Finance* 34, 2742–2752.
- Poterba, J. M., Summers, L. H., 1988. Mean reversion in stock prices. *Journal of Financial Economics* 22, 27–59.

- Stentoft, L., 2004. Assessing the least squares Monte-Carlo approach to American option valuation. *Review of Derivatives Research* 7, 129–168.
- Taylor, S. J., 2005. *Asset Price Dynamics, Volatility and Prediction*. Princeton University Press.
- Tian, Y. S., 2004. Too much of a good incentive? the case of executive stock options. *Journal of Banking and Finance* 28, 1225–1245.
- Yiang, G. J., Tian, Y. S., 2010. Forecasting volatility using long memory and comovements: An application to option valuation under SFAS 123R. *Journal of Financial and Quantitative Analysis* 45 (2), 503–533.

Appendix A. ARMA representation of exact discretization of TOU process.

We rewrite equation (5) in terms of the detrended process of p_t :

$$dq_t = -\kappa q_t dt + \sigma dZ_{S,t}, \quad \kappa > 0$$

where $q_t \equiv p_t - \theta t$ and initial condition $q_0 = p_0 = \ln S_0$. The exact solution to this univariate Ornstein Uhlenbeck (OU) process reverting to an unconditional mean of zero is, according to Bergstrom (1984), the following discrete-time process:

$$q_t = f_h q_{t-h} + \varepsilon_t, \quad \varepsilon_t \sim iid N(0, \sigma_{\varepsilon,h}^2)$$

where $f_h \equiv e^{-\kappa h} < 1$ and $\sigma_{\varepsilon,h}^2 = \frac{\sigma^2}{2\kappa} (1 - e^{-2\kappa h})$. Next, we obtain the equation for $\Delta_h q_t \equiv q_t - q_{t-h}$ and rewrite in terms of the stock return, $r_{t,h}$, given in equation (10). Then,

$$r_{t,h} = \theta h(1 - f_h) + f_h r_{t-h,h} + \eta_t, \quad \eta_t = \varepsilon_t - \varepsilon_{t-h},$$

where η_t follows a MA(1) process verifying that

$$\mathbb{E}[\eta_t] = 0, \quad Var(\eta_t) = 2\sigma_{\varepsilon,h}^2 \quad \text{and} \quad Cov(\eta_t, \eta_{t-jh}) = \begin{cases} -\sigma_{\varepsilon,h}^2 & \text{if } j = 1 \\ 0 & \text{if } j \geq 2 \end{cases}$$

Hence, $r_{t,h}$ is described by a stationary discrete-time ARMA (1,1) process with an unconditional mean of θh , an unconditional variance σ_r^2 shown in equation (11) and a negative first-order autocorrelation $\rho_r(1)$ exhibited in equation (12). Finally, the returns are negatively correlated at all lags. The higher order autocorrelation coefficients are easily obtained as $\rho_r(j) = f_h^{j-1} \rho_r(1)$, $\forall j \geq 2$.

The following table summarizes the moments of the continuously compounded returns, defined as in equation (10) under the two alternative processes.

	TOU	GBM
Mean ($\bar{r}$)	θh	$(\tilde{\mu} - \tilde{\sigma}^2/2)h$
Variance ($\sigma_{r,h}^2$)	$\frac{\sigma^2}{\kappa} (1 - e^{-\kappa h})$	$\tilde{\sigma}^2 h$
Autocorrelation at lag 1	$-\frac{1}{2} (1 - e^{-\kappa h})$	0

$\tilde{\mu}$ and $\tilde{\sigma}$ denote the trend and diffusion coefficients of the GBM process.

By matching both unconditional means and variances for the h -period returns under both processes, we obtain the following constraints in the parameters:

$$\tilde{\mu} = \theta + \frac{\tilde{\sigma}^2}{2} \quad , \quad \tilde{\sigma} = \sigma \sqrt{\frac{1 - e^{-\kappa h}}{\kappa h}}.$$

In short, these parameter restrictions lead to a fair comparison across the paper when we aim to analyze exclusively the effects of predictability on the objective ESO valuation.

Appendix B. Subjective and objective ESO pricing algorithm.

The algorithm used to solve the executive's problem in equation (7) is based on the popular LSM algorithm of Longstaff and Schwartz (2001) for valuing American options. The modification to consider a risk-averse option holder was introduced by León and Vaello-Sebastià (2010). We really adapt this approach, summarized in steps 1 to 6 below, by changing the dynamics of the stock price process. Given m simulated paths, the algorithm consists on creating a vector Y of length m with the expected utility for each path obtained from the optimal exercise of the ESOs. For the one state-variable case ($S1$) the algorithm simplifies, since the executive only has to search the optimal exercise time, t^* ($\eta = 0$). Then, we first explain the algorithm for the $S1$ model and later, we will extend it for the $S2$ case.

Step 1: Given m simulated paths of the state-variable (S_t), the optimal exercise rule at maturity (T) is obvious, to exercise all paths in-the-money. Vector Y_T is obtained using equations (1) and (6) on all paths at T .

Step 2: At $T - 1$ the executive has to choose between exercising or continuing with the ESOs for one period. The utility at T in case of exercise at $T - 1$, $U(W_{T|T-1})$, is computed

using equations (2) and (6). Finally, to estimate the expected utility of exercise conditional on the current values of the state-variables, $\mathbb{E}_{T-1} [U(W_{T|T-1}) | S_{T-1}]$, we regress $U(W_{T|T-1})$ over some basis functions of the state-variable (S_{T-1}). The projected values of that regression are $\hat{\mathbb{E}}_{T-1} [U(W_{T|T-1}) | S_{T-1}]$.

Step 3: Compute the expected utility in case of keeping the ESO one period ahead. We only have to regress the one period ahead utility, Y_T , over basis functions of S_{T-1} . Then $\hat{\mathbb{E}}_{T-1} [Y_T | S_{T-1}]$ is the projected value of the previous regression.

Step 4: Compare the expected utility when exercising and holding the ESOs and then, select the maximum. That is, we compare $\hat{\mathbb{E}}_{T-1} [U(W_{T|T-1}) | S_{T-1}]$ (step 2) with $\hat{\mathbb{E}}_{T-1} [Y_T | S_{T-1}]$ (step 3). For those y paths where $\hat{\mathbb{E}}_{T-1} [U(W_{T|T-1}) | S_{T-1}] > \hat{\mathbb{E}}_{T-1} [Y_T | S_{T-1}]$ the ESOs are exercised and the y elements of Y_T are updated with $\hat{\mathbb{E}}_{T-1} [U(W_{T|T-1}) | S_{T-1}]$.

Step 5: Repeat steps 2 to 4 up to t_0 , where t_0 can be either the grant date or the first time step after the vesting period, and compute the expected utility as the mean of Y_T .

Step 6: Find the certainty-equivalence using equations (8) and (17).

To compute the expectations at steps 2 and 3 we must select the basis functions. We have noticed that the algorithm works better when using powers of $\log(S_t)$ instead of powers of S_t , specifically we use $\log(S_t)$, $[\log(S_t)]^2$ and $[\log(S_t)]^3$. Regarding the number of basis functions, adding higher order powers yields less stable results because the matrix of basis functions becomes close to singular. This yields less accurate results when computing the inverse.¹⁵

To obtain the firm's ESO cost V^{OBJ} , we follow the approach of Hall and Murphy (2002), which is based on the exercise threshold. This threshold establishes a frontier where the executive is indifferent between holding or exercising the option. Thus, if the stock price is above (below) the frontier, the executive will exercise (hold) the ESO.

To compute the exercise threshold, at step 4 in the previous algorithm, we record the minimum value of S_t for which the executive decides to exercise the ESO. Once we have ob-

¹⁵Moreno and Navas (2003) and Stentoft (2004) provide numerical details about the robustness of the LSM method. New Monte-Carlo type methods, based on either the Malliavin calculus or the regression based approaches, to improve the efficiency for the pricing of American options can be found in Bouchard and Warin (2011).

tained the exercise threshold, we simulate the price dynamics under the risk-neutral measure. We record the payoffs associated to those paths which are exercised at the given threshold. Then, the payoff is discounted at the risk-free rate.¹⁶

The algorithm for the $S2$ model consists on repeating the previous steps 1 to 6 for a grid of values of η . Then, we take that value η which maximizes the expected utility, η^* . To compute the two expectations into the two state-variable model, we have used up to the second order powers of $\log(S_t)$ and $\log(M_t)$ plus the cross-product, $\log(S_t) \times \log(M_t)$.¹⁷

Appendix C. Optimization and numerical details.

To obtain the subjective ESO valuation, we must find the amount CE satisfying equation (9). It depends upon the model considered. For the $S1$ model, executives cannot allocate their unrestricted wealth into a market portfolio. Therefore, for this model we propose the following distance minimization program:

$$\min_{CE} \left(1 - \frac{\mathbb{E}_0 [U(W_T^{CE})]}{\mathbb{E}_0 [U(W_{T|t^*})]} \right)^2. \quad (17)$$

To justify this approach, notice that for the standard parameters used in this article, the expected utility, $\mathbb{E}_0 [U(W_{T|t^*})]$, can be of order 10^{-21} . Thus, in order to avoid problems in the numerical optimization we have selected the minimization of a quadratic relative distance instead of the simple quadratic distance.

Figure 7 shows the shape of the objective function in equation (17) with respect to CE for different degrees of diversification α . In this case $\kappa = 0.1$, $\gamma = 4$ and the remaining parameters have been settled to their benchmark values. The left-hand graphic shows the joint effects of CE and α in the objective function, while the right-hand one displays level curves. We can appreciate that the objective function is well behaved and exhibits a global minimum. Accordingly, the minimization in equation (17) converges quickly. This result holds for a wide range of values for κ and γ , which are available upon request.

For the $S2$ model, there is a market portfolio available for allocating the executive's unrestricted wealth. Therefore, the previous distance minimization program is no longer

¹⁶Remind that V^{OBJ} is a risk-neutral valuation but considering the executive's exercise rule.

¹⁷Although the inclusion of $\log(S_t) \times \log(M_t)$ is recommended by the relevant literature, its contribution here is small. Further details are available upon request.

applicable. To solve equation (9), we have implemented a simple iterative procedure in which CE is also obtained from the minimization of a quadratic relative distance as that described by equation (17) but in which the numerator is now given by

$$\hat{U}(CE) \equiv \max_{\eta} \mathbb{E}_0 \left[U(W_T^{CE}) \right], \quad (18)$$

and the denominator by U^* , given by equation (7), which is obtained from a grid-search procedure as mentioned in the last paragraph of Appendix B. In each step of the iteration a new value of CE is evaluated until convergence is achieved for some desired level of accuracy. Further details are available from the authors upon request.

Figure 8 displays $\mathbb{E}_0 [U(W_{T|t^*})]$ for a grid of values of η and for different levels of κ and γ . We observe that $\mathbb{E}_0 [U(W_{T|t^*})]$ achieves a maximum for some given value of η , denoted as η^* . To obtain η^* , we have fitted a quadratic polynomial using the pairs $\{\eta, \mathbb{E}_0 [U(W_{T|t^*})]\}$ obtained in the former grid-search. Then, η^* follows from taking the first derivative, namely $\eta^* = -p_1/2p_2$, where p_i is the i -th order coefficient of the fitted polynomial.

Figure 7: Certainty equivalent optimization

The left-hand graphic displays equation (17), the quadratic distance function for model $S1$, as a function of CE for different degrees of diversification α . The right-hand graphic shows three level curves for $\alpha = \{0.4, 0.5, 0.75\}$. The relevant parameters are $\kappa = 0.10$ and $\gamma = 4$. The remaining parameters have been set to the benchmark values enumerated in Section 3.

Figure 8: Expected utility and portfolio choice

These graphics display the maximum expected utility, that can be achieved when the executive chooses the optimal exercise date for the granted ESOs, as a function of the portfolio choice, η . The left-hand graphic holds a relative risk-aversion coefficient of $\gamma = 2$ while it becomes $\gamma = 4$ for the right-hand one. Each line corresponds to a different degree of predictability, κ . The rest of parameters has been settled to the benchmark values enumerated in Section 3.