Executive Stock Options: Portfolio Effects [†]

Vicky Henderson[‡]

University of Oxford

Jia Sun [§]

China Credit Rating Co. Ltd

A. Elizabeth Whalley [¶]

Warwick Business School

August, 2013

[†]We thank Kris Glover, Egor Matveyev, Kevin Murphy, Rong Wang, Rafal Wojakowski, David Yermack, participants at the EFA 2013, EFMA 2013 and NFA 2012 and seminar participants at Princeton University, Toulouse School of Economics, Lancaster University Management School, LSE, Cass Business School, University of Technology Sydney, Rutgers, King's College, London, Oxford-Man Institute, and Warwick Business School.

[‡]Oxford-Man Institute, University of Oxford, Eagle House, Walton Well Road, Oxford. OX2 6ED. UK. Email: vicky.henderson@oxford-man.ox.ac.uk

[§]China Credit Rating Co. Ltd, Ying Tai Mansion. 28 Jin Rong Street, Beijing, Beijing 100032, China. Email: sunjia@chinaratings.com.cn

[¶]Finance Group, Warwick Business School, University of Warwick, Coventry. CV4 7AL. UK. Email: elizabeth.whalley@wbs.ac.uk

Executive Stock Options: Portfolio Effects

Abstract

More than half of S&P 500 CEOs receive options annually, however extant valuation models have not accounted for portfolio considerations. We show the inability of executives to diversify means portfolio effects matter: exercise thresholds and shareholder costs are lower than for stand-alone options and vary with portfolio size and composition. The model explains which options to exercise first and how exercise can be induced by new grants, matching empirical findings. Ignoring portfolio composition may invalidate estimates of exercise thresholds, bias inferences about executive characteristics (e.g. optimism) drawn from them and overstate costs, particularly for executives with larger portfolios.

Keywords: Executive stock options, employee stock options, executive compensation, option portfolios, risk aversion, exercise thresholds

JEL Classification Numbers: G11, G13, G30, J33

1 Introduction

Stock options have been an important component of executive compensation packages since the early 1990s. Recent studies by Murphy (2012), Conyon, Fernandes, Ferreira, Matos, and Murphy (2011) and Frydman and Saks (2010) document their useage by S&P 500 companies. By 2001, stock options accounted for over half of median compensation for CEOs in S&P 500 firms. Ten years later, stock options still represented 21% of total pay, and 68% of CEOs received stock options. Furthermore, more than half of S&P 500 CEOs have received options annually since 2006, leading to the potential accumulation of a portfolio of exercisable options. Despite the fact that executives often hold multiple option grants, the extant theoretical literature has only considered valuation models for a single grant of options (with the same strike price and time to maturity). Indeed, in a recent study of option exercise and shareholder costs, Carpenter, Stanton, and Wallace (2010) comment, *"It would be useful to understand which options are most attractive to exercise first"* (p 318). In this paper we investigate in a utility-based model the effects of an executive holding several grants of executive stock options (ESOs) on the exercise threshold and cost to the shareholders of each option. By doing this, we remedy the omission to date of portfolio effects from the theory of ESO exercise and costs. We find that simply being part of a portfolio of options makes a difference to individual options' exercise thresholds and costs; it is not possible simply to consider each option on a stand-alone basis and then combine.

Understanding executives' exercise behavior is key to studying their subjective valuation and the cost to shareholders of the options granted. A better knowledge of valuation and costs will assist firms in awarding options more effectively and so is of paramount importance to shareholders and investors. Furthermore, estimates of shareholder costs are required to satisfy accounting regulations in the US and Europe (FASB 123 and IFRS 2). Executives' exercise behavior is also of wider importance in the corporate finance literature. For example, Malmendier and Tate (2005), amongst many others, use option exercise to detect optimism and Bergstresser and Philippon (2006) interpret option exercise as a signal of private information. Our model provides such a *benchmark* for thresholds *excluding* these behavioral or information-based effects, which will allow more accurate and/or more powerful inferences to be drawn. Indeed Van Bekkum and Zhu (2013) stress the lack of models available to identify which options from a portfolio should be exercised first in the absence of insider information. We find that recognition of an executive's whole portfolio of options can make a significant difference to exercise thresholds, lowering them for options exercised early in the exercise order. Furthermore, a new option grant can induce immediate exercise of existing options. These findings are consistent with recent empirical evidence on exercise thresholds, e.g. Klein and Maug (2010), who also remark, *"The explanatory power of the option portfolio effects is very large."*. Thus ignoring portfolio composition invalidates not only estimates of ESO exercise thresholds but also inferences about executive characteristics, such as optimism, drawn from these thresholds.

Our model generates several new "portfolio effects" which have important impacts on ESO thresholds and costs. Firstly, for individual options in portfolios, the option's position in the portfolio's optimal exercise order is a new key factor determining both an option's exercise threshold and its cost within the portfolio. The earlier in the optimal exercise order, the lower the option's threshold and the lower its cost as part of a portfolio. The optimal exercise order depends on the overall composition of the executive's option portfolio and is determined endogenously. All else equal, options are earlier in the exercise order, the higher their moneyness and the shorter their time to maturity; however, where these conflict, but differences in moneyness are not too great, the optimal exercise order can switch from exercising an option with higher moneyness first, to one with shorter maturity (but lower moneyness). Secondly, portfolios of ESOs are always valued at a discount to the sum of the standalone utility-based values of their constituent options. The proportional discount depends on both the size and the overall composition of the portfolio: it is greatest for portfolios where there is a switch in the optimal exercise order and increases when new options are added to a portfolio. Finally, the relevant cost of a new option grant is its *incremental* cost, which is always lower than the option's stand-alone cost, and which differs depending on the size and composition of the portfolio it is being added to, with the effect that the cost of a new option grant varies non-monotonically with its own strike price.

These results are in stark contrast to a Black-Scholes world, where thresholds and costs are independent of other option holdings. The underlying reason all these "portfolio effects" arise is because the inability of the executive to hedge options using the firm's shares exposes her to unhedgeable or non-diversifiable risk whilst she continues to hold them. The basic intuition for single option grants, that lack of diversification leads to lower (subjective) valuation, early exercise, and lower shareholder costs, is well understood in the literature. Portfolio effects arise because unhedgeable risk increases *nonlinearly* with portfolio size and thus has greater impact as more options are added to a portfolio. This nonlinearity affects not only the executive's subjective valuation of her portfolio, but also her exercise strategy, and cost to the shareholders: there are interaction effects between different options in a portfolio which means that thresholds and costs for individual options in a portfolio cannot be calculated on a stand-alone basis, but need to be determined for the portfolio as a whole. Furthermore, thresholds and costs for a particular option will vary depending on the composition of the portfolio it is a part of. Our main findings arise directly from the nonlinearity of unhedgeable risk, and thus do not depend on our exact choice of modeling assumptions.[1]

Early literature (Lambert, Larcker, and Verrecchia (1991), Huddart (1994), Carpenter (1998), Hall and Murphy (2000, 2002) and Cai and Vijh (2005)) used a certainty equivalent or utility-based framework to value a single ESO grant, assuming the executive's non-option wealth is invested in exogenously specified, non-optimized proportions in stock and risk-free bonds. They found the executive's certainty equivalent or subjective valuations are lower and the option is exercised at a lower threshold than under Black-Scholes (which assumes perfect hedging).[2]. Later studies by Detemple and Sundaresan (1999), Henderson (2005), Grasselli and Henderson (2009), Leung and Sircar (2009) and Carpenter, Stanton, and Wallace (2010) allowed the executive to optimize the investment of her outside wealth in risk-free bonds and a market asset but still considered only a single grant of identical options.[3] They showed that hedging using the market can remove the systematic portion of this risk, but the firm-specific risk of the executive's position remains unhedgeable. Under utility-based valuation, the effective cost of this unhedgeable risk reduces the executive's valuation of her option grant whilst she continues to hold it: once it has been exercised, the shares acquired are sold and the proceeds invested (optimally) in tradeable securities. Her early exercise strategy thus takes account of the *removal* of unhedgeable risk on exercise, and so she exercises at a lower level of moneyness

[1]See Section 2 for details of our model, and Section 5 for a discussion of robustness.

[2]Lambert, Larcker, and Verrecchia (1991) also set out basic principles for a portfolio of stock and an option grant.

[3]Related contributions are made by Ingersoll (2006), Jain and Subramanian (2004), Kahl, Liu, and Longstaff (2003) and Rogers and Scheinkman (2007).

than in a world where risks can be perfectly hedged. Early exercise impacts on the cost of an ESO to the firm's (well-diversified) shareholders: the literature shows that although this cost will be higher than the executive's subjective valuation, it is still lower than the corresponding Black-Scholes cost.

Whilst the above intuition and theory of the impact of unhedgeable risk on ESO exercise, valuation and costs is well developed for a single option grant, our framework extends utility-based valuation to treat option portfolios. Importantly, in a portfolio, unhedgeable risk is no longer removed on exercise of one option; instead it is the *reduction* in unhedgeable risk which determines the exercise threshold.

Both the total unhedgeable risk associated with a portfolio and the reduction in unhedgeable risk on the exercise of the first option grant from a portfolio increase non-linearly with portfolio size.[4] So the optimal exercise threshold for an option if it is the first option to be exercised in a portfolio is distinctly lower than the exercise threshold for the same option if it is exercised last (or later in the exercise order). More generally, the optimal exercise threshold for a particular option is lower, the larger the option portfolio remaining on exercise. If the option's position in the optimal exercise order for the portfolio changes, say from penultimate- to last-to-be-exercised, or vice versa, then the unhedgeable risk reduction on exercise also alters instantaneously, so the exercise threshold jumps (up if the new position in the exercise order is later, down if it is earlier).

[4]Unhedgeable risk is proportional to the square of the portfolio's delta. So the unhedgeable risk associated with a holding of two identical options is in the order of four times the unhedgeable risk associated with a single option, and the reduction in unhedgeable risk associated with exercising the first of these two identical options is in the order of three times the unhedgeable risk reduction associated with exercising the last.

Empirical studies have found support for a number of economic and behavioral factors influencing executives' early exercise decisions in practice.[5] Much of this empirical evidence is consistent with utility maximization models which explains their frequent useage in the literature. For example, Huddart and Lang (1996) and Bettis, Bizjak, and Lemmon (2005) document early exercise even in the absence of dividends, and Bettis, Bizjak, and Lemmon (2005) find that options are exercised earlier at higher volatility firms. However, whilst Carpenter, Stanton, and Wallace (2012) recognize that there should be inter-relations between the exercises of different grants to the same executive, to date only Klein and Maug (2010) and Armstrong, Jagolinzer, and Larcker (2007) have tested for aspects of portfolio effects on exercise thresholds explicitly.

There is strong evidence of the explanatory importance of portfolio effects - Klein and Maug (2010) remark, *"These [option portfolio] effects have been neglected in previous studies, but they are empirically of first-order importance"* (p24). Our model is able to generate these empirical findings. Of particular relevance is the finding that the exercise rate increases in the week after an executive has received a grant of new options (Klein and Maug (2010)).[6] We show that the moneyness of the exercise threshold for whichever option is first-to-be-exercised decreases when another option is unexpectedly added to a portfolio. If some of the executive's existing options are in-the-money, these options' thresholds will jump down, potentially from above to below the current stock price, causing an increase in their exercise rate. In contrast, under the Black-Scholes model, a new grant does not alter the thresholds for existing options, and thus

[5] Huddart and Lang (1996), Carpenter (1998), Heath, Huddart, and Lang (1999), Bettis, Bizjak, and Lemmon (2005), Armstrong, Jagolinzer, and Larcker (2007), Klein and Maug (2010), Srivastava (2011), and Carpenter, Stanton, and Wallace (2012).

[6] Ofek and Yermack (2000) and Srivastava (2011) find related evidence.

we would not expect a change in the exercise rate. Similarly, if we were modeling only a single option grant in a utility model, as in the extant literature, we could not make inferences about the increased likelihood of exercise of other options.

In addition to explaining existing empirical findings, we contribute new implications to inform future studies. Existing portfolio holdings need to be considered in empirical studies of option exercise so that differing exercise behavior is not falsely attributed to executive characteristics. To illustrate, consider two executives at the same firm, $\mathcal{A}$ and $\mathcal{B}$, who have the same level of risk aversion and ESO holdings except that $\mathcal{A}$ holds one additional option grant (she has been employed longer at the firm). If it is not optimal for $\mathcal{A}$ to exercise her additional option first, then $\mathcal{A}$ will exercise some, if not all, of the options she has in common with $\mathcal{B}$ at a lower threshold, and hence earlier than $\mathcal{B}$, even though they have the same risk aversion. An examination of exercise strategies at the firm might erroneously suggest that $\mathcal{A}$ is more risk-averse than $\mathcal{B}$. Our work also suggests that measures of optimism based on option exercise should incorporate portfolio heterogeneity.[7] For example, if both $\mathcal{A}$ and $\mathcal{B}$ exercise at the same 'late' level of moneyness, this should provide a stronger indication of optimism for $\mathcal{A}$ than for $\mathcal{B}$ because in the absence of optimism, $\mathcal{A}$ would exercise at a lower threshold.

Shareholders are assumed to be well-diversified with no restrictions on trading and thus value ESOs in a perfect market (Black-Scholes world) taking into account the executive's exercise strategy. Prior literature has shown that, even for a single option, the lower optimal exercise threshold reduces the ESO cost to the shareholders relative to the Black-Scholes value of a

[7] In particular, in constructing measures of optimism, Malmendier and Tate (2005) study "late" exercise behavior and calibrate a *single* threshold for *all* CEOs which does not take into account the heterogeneity of CEOs option portfolios.

similar traded option. However, an executive with multiple options will exercise some of these at even lower moneyness thresholds, increasing the proportional discount relative to Black-Scholes still further for these options and so also increasing the proportional discount for the portfolio as a whole. For example (see Tables 1 and 2) combining two at-the-money options with different times to maturity (5 and 10 years) can reduce the cost of the 5-year option from 53% to 30% of the Black-Scholes value, a proportional discount of 44% relative to its cost as a stand-alone option. The portfolio cost is over 20% lower than the sum of the costs of individual options, and only 40% of the Black-Scholes value.

Importantly, the shareholder cost of an option depends on the composition of the remainder of the executive's portfolio. The primary reason is because the moneyness of the early exercise threshold, and hence the shareholder cost, depends critically on the option's location in the optimal exercise order. Generally, the later the option is in the exercise order, the higher its moneyness threshold and the higher its cost[8] (lower its proportional discount). However, the exercise order depends on the strike prices and times to maturity of the other options in the portfolio. All else equal, the *higher* the strike price and the *longer* the time to maturity, the *later* in the exercise order an option is. So if the strike prices and times-to-maturity are co-monotonic,[9] options are exercised in order of increasing strike price, so options with the lowest strike price (and shortest time to maturity) have the largest proportional cost discounts for shareholders relative to Black-Scholes. However, if the strike prices and times to maturity are not co-monotonic, i.e. a longer-dated option has a lower strike price, the optimal exercise order

[8]If the exercise order changes during the option's life, the longer its position is later in the exercise order, the higher its cost.

[9]Strikes and times to maturity are ordered so that options with higher strikes also have longer times to maturity, see Henderson, Sun, and Whalley (2013).

needs to be calculated as part of the solution. We find options are generally exercised in order of increasing strike price, but options which are close to maturity may be exercised earlier than their strike price would indicate.

These interaction effects lead to new features of ESO costs. Firstly, the shareholder cost of a particular ESO with fixed strike price and time to maturity varies non-monotonically with the strike prices and times to maturity of the other options in the executive's portfolio. Secondly, holding the remainder of an executive's portfolio constant, the shareholder cost of a particular ESO may no longer be monotonically decreasing in its own strike price; it can have local minima. Both findings are in stark contrast to the standard results for exchange-traded options valued under absence of arbitrage, where an option's value is independent of the portfolio in which it is held.

Firms are concerned about the overall cost of options granted to executives, but specifically the cost of the current option grant. A key implication of the nonlinearity of the shareholder cost of ESOs is that this should be measured as the *marginal* or *incremental* cost of the grant. In a Black-Scholes world, this would equal its stand-alone value; however with portfolio effects this is no longer true: the incremental cost is always lower than the stand-alone cost, and varies with the executive's existing exposure to unhedgeable firm-specific risk. So firms should recognize that the cost of a given option grant (and indeed the value of that grant to the executive) will vary depending on the executive's existing portfolio of options on the grant date. Returning to our earlier scenario, the cost of $\mathcal{A}$'s ESO portfolio is greater than the cost of $\mathcal{B}$'s, but the proportional discount relative to stand-alone values will be greater for $\mathcal{A}$, and the incremental cost to shareholders of $\mathcal{A}$'s additional option is likely to be substantially lower than the cost of that option as a stand-alone. Moreover, firms should recognize that the incremental cost to the shareholders of identical unexpected new option grants to the two executives will be lower for $\mathcal{A}$ than $\mathcal{B}$.

Our findings also suggest that further work is needed in developing suitable approximations of shareholder costs for useage in practice. Modifications of Black-Scholes valuation using fixed adjustments to option maturity[10] do not account for portfolio heterogeneity. In our example, any of the options in common that $\mathcal{A}$ exercises earlier than $\mathcal{B}$ will have lower expected time to exercise for $\mathcal{A}$ than $\mathcal{B}$. Cost approximations thus need to take portfolio size, and ideally, composition, into account.

Finally, as highlighted in Murphy (2012) (p6, and see also Lambert, Larcker, and Verrecchia (1991)), exposure to unhedgeable risk has wider implications for valuing executive pay. Empirical studies have typically added different components of compensation. However, Conyon, Core and Guay (2011) and Fernandes, Ferreira, Matos and Murphy (2013) have recognized the need to consider the executive's *total exposure* to equity risk (rather than the exposure associated with a single option grant) when estimating risk premia associated with executive compensation. Indeed, incorporating these risk adjustments enabled them to make the important conclusion that US CEOs are not paid significantly more than their foreign counterparts, which is in stark contrast to previous findings. Whilst we do not report them here, many of the principles of portfolio effects we document also hold for subjective values: for example the use of incremental values for new option grants, and the need to value portfolios as a whole, since risk effects mean this value is lower than the sum of each option valued on a stand-alone basis.

The remainder of the paper is organized as follows. Section 2 presents the model. Sections 3

[10]For example, the typical FASB approximation uses the Black-Scholes formula with maturity equal to the expected time to exercise.

and 4 investigate the implications for the exercise thresholds and shareholder costs. Robustness issues including vesting, employment termination, restricted stock and alternative assumptions about preferences and stock price evolution are discussed in Section 5 and Section 6 concludes.

2 Model

Consider a risk-averse executive who is granted n (American-style) finite-lived non-transferable call options on the stock of her company. Her investment opportunities consist of a riskless bond with constant riskless rate r and the market portfolio with price M_t. The executive is restricted from trading the underlying stock of her company, with price S_t.[11]

The set of n options have maturities $T_W \leq T_X, ..., T_Y \leq T_Z$ and corresponding strikes $K_W, K_X, ..., K_Y, K_Z$ and we refer to them as option W, X, ..., Y, Z. The notation $X(K_X, T_X)$ denotes option X with strike K_X and maturity T_X. In practice, executive stock options typically have ten year maturities at grant date, but on a given date, the executive may hold a portfolio of options from several grants each with varying time to maturity remaining and different strikes.

[11]Insiders cannot short sell stock as they are prohibited by Section 16-c of the Securities and Exchange Act. For this reason, we use the term *executive* for the option holder. Our model could also apply to lower-ranked *employees* - whilst they are not subject to Section 16-c short sale restrictions, costs of short selling may limit their trading. Indeed, evidence of early exercise across all ranks of employees suggests all employees are constrained as to the hedging they can carry out (Bettis, Bizjak, and Lemmon (2005); Carpenter, Stanton, and Wallace (2012)). Bettis, Bizjak, and Kalpathy (2011) provide evidence of the use of derivatives by insiders and find a significant proportion of the executive's exposure remains unhedged. They also find evidence that derivatives are primarily used for information based (insider information) reasons and no evidence that hedging via derivatives is used as diversification against changes in firm risk. Such deals must be reported to the SEC and generally only involve high-ranking executives.

The prices follow[12]

$$\frac{dS}{S} = (\nu - q)dt + \eta dB \tag{1}$$
$$\frac{dM}{M} = \mu dt + \sigma dZ \tag{2}$$

where standard Brownian motions B and Z are defined on a probability space $(\Omega, \mathcal{F}, \mathcal{F}_u, \mathbb{P})$ where $\mathcal{F}_u$ is the augmented σ-algebra generated by $\{B_u, Z_u; 0 \le u \le t\}$ and their instantaneous correlation is $\rho \in (-1, 1)$. The volatility of stock returns η, expected return on the stock ν, proportional dividend yield $q > 0$, and expected return μ and volatility of the market returns σ are all constants. The mean stock return ν is equal to the CAPM return for the stock, given its correlation with the market, $\nu = r + \beta(\mu - r)$; $\beta = \rho\eta/\sigma$. In common with the utility-based ESO literature, we abstract from leverage and dilution considerations, which should not affect our main results.[13]

The restricted executive faces some unhedgeable risk through her option position since $\rho \in (-1, 1)$. Allowing the executive to trade in the market asset enables her to partially hedge the risk she faces from her option portfolio. She holds a cash amount θ_s in M at time s and invests the remainder of her wealth at the riskless rate r. The executive's trading or outside wealth $W_u; u \ge t$ follows

$$dW_u = (rW_u + \theta_u(\mu - r))du + \theta_u \sigma dZ_u;\ W_t = w \tag{3}$$

[12]We assume the firm's stock price is unaffected by the executive. See Section 5 for further discussion.

[13]Denis and Rendleman (2008) consider multiple warrant issues and potential dilution effect of each issuance on the remaining warrants. See also the classic works of Constantinides (1984) and Spatt and Sterbenz (1988) on sequential exercise. Their focus is the impact of dilution on costs, whereas our paper explores the interdependencies of different American style options arising from the impact of unhedgeable risks and the effect of these dependencies on shareholder cost.

given initial wealth w. The remaining unhedgeable risk that cannot be hedged away is represented by $\eta_e^2 = (1-\rho^2)\eta^2$ and thus is greater, the lower the (absolute value of) correlation between the stock price and the market.[14]

An executive with an option portfolio as we have described will maximize her expected utility of terminal wealth at $\tilde{T}$ (where we take $\tilde{T} \geq T_Z$) over the choice of exercise times $\tau^W, \tau^X, ..., \tau^Y, \tau^Z$ of the options and the choice of outside investment in the market θ_u (satisfying integrability condition $\mathbb{E}\int_t^{\tilde{T}} \theta_u^2 du < \infty$) and bond. We assume the executive has constant absolute risk aversion (CARA) denoted by $U(x) = -e^{-\gamma x}$; $\gamma > 0$. Each time an exercise occurs, it is optimal for the executive to sell all shares acquired immediately.[15] The cash proceeds are added to the executive's outside wealth position or trading portfolio and will continue to be optimally invested in the market and bond until $\tilde{T}$. Denote the ordered exercise times by $\tau^{(n)} \leq ... \leq \tau^{(1)}$ where $\tau^{(i)}$ is the exercise time at which there are i options in the portfolio before exercise.[16]

Denote by Π the portfolio of unexercised options and by $|\Pi|$ the size of this portfolio or the number of remaining options. For example, if options X and Z remain unexercised then $\Pi = \{X, Z\}$ and $|\Pi| = 2$. We also define the shortest maturity left in the portfolio by $T_{min} = min\{T_\pi : \pi \in \Pi\}$. The value to the executive $V^\Pi(u, W_u, S_u)$ with remaining options Π, current

[14]Given our primary interest is in the effect of unhedgeable risk, we take $\eta_e = \eta$ in our numerical examples. The general effect of additional hedgeable risk in utility-based models is to increase exercise thresholds and shareholder costs (see Henderson (2005)). However, the key new portfolio effects we describe will be unchanged.

[15]The value of continuing to hold the stock is lower than the stock's market value due to unhedgeable risk. Thus in the absence of minimum holding constraints immediate sale on exercise is optimal. Empirically, many executives sell shares immediately upon exercise (Ofek and Yermack (2000); Huddart and Lang (1996)).

[16]Hence $\tau^{(n)} = \tau^W \wedge \tau^X \wedge ... \wedge \tau^Y \wedge \tau^Z$ is the earliest time at which an option is exercised, and $\tau^{(1)} = \tau^W \vee \tau^X \vee ... \vee \tau^Y \vee \tau^Z$ the last time an exercise occurs.

wealth W_u and current stock price S_u solves the variational inequalities:

$$V^{\Pi}(u, W_u, S_u) \geq \max_{\pi \in \Pi} \{V^{\Pi \setminus \{\pi\}}(u, W_u + (S_u - K_\pi)^+, S_u)\} \tag{4}$$

$$\frac{\partial V^{\Pi}}{\partial t} + \sup_{\theta} \{\mathcal{L} V^{\Pi}\} \leq 0 \tag{5}$$

where the differential operator $\mathcal{L}$ is defined by

$$\mathcal{L} = \frac{\eta^2 s^2}{2} \frac{\partial^2}{\partial s^2} + (\nu - q)s\frac{\partial}{\partial s} + \rho\theta\sigma\eta s \frac{\partial^2}{\partial w \partial s} + \frac{\theta^2 \sigma^2}{2} \frac{\partial^2}{\partial w^2} + [\theta(\mu - r) + rw]\frac{\partial}{\partial w}. \tag{6}$$

If the stock is worthless then options are out-of-the-money and the executive optimally invests her current wealth into the market and riskless bond until the terminal date $\tilde{T}$, giving boundary condition $V^{\Pi}(u, W_u, 0) = \mathcal{M}(u, W_u, \tilde{T})$ where

$$\mathcal{M}(t, w, \tilde{T}) = \sup_{\{\theta_s\}_{t \leq s \leq \tilde{T}}} \mathbb{E} U(W_{\tilde{T}} | W_t = w) = -e^{-\gamma w e^{r(\tilde{T}-t)}} e^{-\frac{(\mu - r)^2}{\sigma^2}(\tilde{T} - t)} \tag{7}$$

is the indirect utility from the Merton (1971) optimal portfolio choice problem. A second set of boundary conditions is obtained at each option maturity,

$$\forall \{\pi\} \in \Pi; \quad V^{\Pi}(T_\pi, W_{T_\pi}, S_{T_\pi}) = V^{\Pi \setminus \{\pi\}}(T_\pi, W_{T_\pi} + (S_{T_\pi} - K_\pi)^+, S_{T_\pi}). \tag{8}$$

The optimal exercise times $\tau^{(n)} \leq ... \leq \tau^{(1)}$ are characterized by[17]

$$\tau^{(|\Pi|)} = \inf\{t \leq u \leq T_{min} : V^{\Pi}(u, W_u, S_u) = \max_{\pi \in \Pi}\{V^{\Pi \setminus \{\pi\}}(u, W_u + (S_u - K_\pi)^+, S_u)\}\}. \tag{9}$$

This says an option from existing portfolio Π is exercised when the value from continuing to hold it is sufficiently low that it equals the value from exercising a particular option π (and

[17]When only one option remains, $|\Pi| = 1$ and the optimal exercise time is given by

$$\tau^{(1)} = \inf\{t \leq u \leq T_\pi : V^{\Pi}(u, W_u, S_u) = V^{\emptyset}(u, W_u + (S_u - K_\pi)^+, S_u)\}$$

where $V^{\emptyset}(u, W_u, S_u) = \mathcal{M}(u, W_u, \tilde{T})$.

investing the payoff $S_u - K_\pi$ in the bond and market). The option that is exercised from the portfolio is the one that gives the maximum continuation value for the remaining portfolio, including reinvestment of the option proceeds. In the Appendix we give details of the separation of variables, transformations and numerical methods to solve the above free boundary problem.

The solution to the problem involves finding the exercise threshold $S^{(|\Pi|)}(t)$ for the next option to be exercised when the portfolio contains the remaining options Π. These are given in (17) (see Appendix). As with standard American options, we can then describe the optimal exercise time for the next option when $|\Pi|$ remain, as the first time the stock price reaches this exercise threshold,

$$\tau^{(|\Pi|)} = \inf\{t \leq u \leq T_{min} : S_u = S^{(|\Pi|)}(u)\}. \tag{10}$$

Once we have solved for each threshold $S^{(|\Pi|)}(t)$, we can infer the optimal exercise boundary for each *option* $\pi \in \Pi$ calculated *as part of the portfolio*, denoted $S^\pi(t); t \leq T_\pi$. The corresponding optimal exercise time for each *option* $\pi \in \Pi$ is

$$\tau^\pi(t) = \inf\{t \leq u \leq T_\pi : S_u = S^\pi(u)\}. \tag{11}$$

For comparison, we also calculate the exercise thresholds $S^\pi_S(t)$ of each option as a *stand-alone*, i.e. if the portfolio consisted of a single option. These are computed via (17) with $|\Pi| = 1$.

Finally, following Henderson (2005), Carpenter, Stanton, and Wallace (2010), Leung and Sircar (2009) amongst others, we define the subjective value $p^\Pi(t, s)$ of the portfolio of remaining options Π to the executive as the amount of cash which invested optimally would give the same expected utility as the options[18]

$$V^{\Pi}(u,w,s) = \mathcal{M}(u, w + p^{\Pi}(u,s), \tilde{T}). \tag{12}$$

Again, the subjective value of a particular option as a stand-alone is computed with $|\Pi| = 1$.

Since our focus in this paper is on the impact of different option strikes and maturities, we take a portfolio with $n = 2$ options. This is our base portfolio for much of the remainder of the paper. In Section 5 we extend to consider portfolios with further options and we comment there on the potential inclusion of restricted stock.

The parameters in our numerical examples are chosen as follows. We consider values for absolute risk aversion γ of 0.1 and 0.2, which are used by Leung and Sircar (2009) to value a single stock option to a risk averse executive. Furthermore, CARA is widely used in the incomplete-markets real options literature and our values are consistent with those used by Miao and Wang (2007) and Chen, Miao and Wang (2010). Our levels of absolute risk aversion are also comparable to levels of relative risk aversion and outside wealth used in the literature on executive stock options. Consider an executive with a grant of one million options with \$10 strike, on stock worth \$10. If she has relative risk aversion of 4 and an outside wealth of \$20 million, her absolute risk aversion is $\gamma = 0.2$. If she instead has an outside wealth of \$40 million (still with relative risk aversion 4) then her absolute risk aversion is 0.1. If her relative risk aversion is greater, say 6, then for her absolute risk aversion to remain at 0.2, she must have outside wealth of \$30 million. Our range of parameters for the riskfree rate, dividends, and the market asset are broadly consistent with those used in the recent theoretical models of a single executive stock option, see Carpenter, Stanton and Wallace (2010).

[18] In the Appendix, we characterize the free boundary problem that $p^{\Pi}(t,s)$ must solve, see (19)-(21).

3 Exercise thresholds and moneyness

It is well known that the exercise threshold for a single executive stock option lies below the equivalent threshold in a Black-Scholes setting, because of the effect of unhedgeable risk on the executive's option value (see Carpenter, Stanton, and Wallace (2010) amongst others). Whilst she continues to hold the option, the executive bears additional risk due to her exposure to the firm's stock price, which she is unable to perfectly hedge. On exercise however, the executive sells the shares she receives and thus eliminates her exposure to this unhedgeable risk. The unhedgeable risk exposure reduces the value of continuing to hold the option to the executive and so decreases the threshold at which she optimally exercises. To gain intuition, in Section 3.1 we recap findings from the existing literature by considering the sequential exercise of a single grant of two identical options. Then, in Section 3.2, we turn to the new and more realistic case where the strike price and time to maturity of each grant is different and so the optimal exercise order must also be determined. Implications for empirical studies of exercise thresholds are drawn in Section 3.3.

3.1 A single grant

Figure 1 shows how the exercise thresholds evolve over time for a portfolio of two identical options, both with strike $K_Y = 10$ and maturity $T_Y = 10$. The dashed line is the executive's optimal threshold taking account of unhedgeable risk for a single option as a stand-alone. It is significantly lower than the highest solid line, which represents the Black-Scholes threshold, for all times before maturity,[19] showing how unhedgeable risk decreases the level of the threshold

[19]Note the Black-Scholes and both utility-based thresholds are strictly greater than K instantaneously before maturity (T_Y^-). Kim (1990) has shown the Black-Scholes world threshold $S_{\mathcal{BS}}^Y(T_Y^-) = \frac{r}{q}K > K$ if $r > q$.

throughout the option's life. Of course the Black-Scholes threshold remains the same regardless of whether the executive is exercising one or many identical options, since all risk is hedgeable and there is no need to unwind risk gradually over time.

When determining whether to exercise the option, the executive trades off the benefits of continuing to hold the option (interest on the deferred payment of the strike price), with the costs, which include not only the foregone dividend yield, as in the Black-Scholes world, but also the cost of the unhedgeable risk exposure.[20] The additional cost of ongoing unhedgeable risk exposure reduces the threshold from the Black-Scholes threshold, by more, the greater the executive's risk aversion and the greater the risk to which she is exposed.[21] Furthermore, the effects of unhedgeable risk increase with the length of time the risk is borne. So whilst the convexity of the option payoff still means that the option value to the executive, and hence the optimal exercise threshold, increases rapidly as the time to maturity increases from zero, the counter-acting reduction in option value due to unhedgeable risk means the threshold stops increasing at a lower time to maturity and can even decrease as the time to maturity increases (see also Carpenter, Stanton, and Wallace (2010)).

The dashed line in Figure 1 represents the optimal exercise threshold for this option not only as a stand-alone option but also when it is part of any portfolio in which it will optimally be

---

Sun (2011) has shown $S^Y_{\mathcal{BS}}(T^-_Y) - S^Y(T^-_Y) \geq S^Y_{\mathcal{BS}}(T^-_Y) - S^Y_{\mathcal{S}}(T^-_Y) > 0$ and increases with the executive's risk aversion. At maturity, it is of course optimal to exercise if and only if the option is in-the-money, so $S^Y(T) = S^Y_{\mathcal{S}}(T) = S^Y_{\mathcal{BS}}(T) = K$. For ease of reading, we do not plot $S^Y(T)$ in any graphs.

[20]We compare the partial differential equation for the option value (see Appendix, (19)) with the Black-Scholes pde given in (24). The additional term representing the effects of unhedgeable risk in (19) is like a non-constant increase in the dividend yield which scales with the executive's risk aversion, γ, the unhedgeable variance of the stock price, η_e^2, and her exposure to this, measured by the stock price times the option delta.

[21]For more details see Sun (2011).

the last option to be exercised. Once all the executive's other options have been exercised, the executive takes into account only the unhedgeable risk associated with the remaining, single option, which will be eliminated on exercise, and exercises at the same threshold as for a stand-alone option. More generally, when deciding when to exercise a particular option, the executive takes into account the *reduction* in her exposure to unhedgeable risk which occurs upon exercise and sale of the resulting shares. Since the effects of unhedgeable risk are related to the *variance* of the unhedgeable risk exposure, they increase nonlinearly with the size of the option portfolio. So the reduction in unhedgeable risk exposure on exercise of a given option is generally greater, the larger the overall option portfolio. For example, exercising the first option in a portfolio of two identical options reduces the executive's exposure to unhedgeable risk by considerably more than exercising the second (last) option. The lowest dot-dashed line in Figure 1 shows the optimal exercise threshold for the first-to-be-exercised option in such a portfolio. As we see in Figure 1, the executive prefers to unwind risk gradually over time by exercising identical options at different thresholds, in agreement with Carpenter, Stanton, and Wallace (2010), Grasselli and Henderson (2009), Jain and Subramanian (2004) and Leung and Sircar (2009).

3.2 Which option is exercised first?

We have seen that the exercise order has a significant impact on the level of the exercise threshold. But what determines the exercise order? In a Black-Scholes world, holding all else constant, options with lower strike price (or equivalently, greater moneyness) and shorter times to maturity are exercised earlier. If these two effects work in the same direction, so options in a portfolio have co-monotonic strike prices and times to maturity, then options are exercised in order of increasing strike (Chapter 4 of Cox and Rubinstein (1985)). Henderson, Sun and Whalley (2013) have shown that this co-monotonicity result continues to hold under utility-based pricing in incomplete markets[22] i.e. for an executive with a portfolio of two options, Y and Z, if $K_Y \leq K_Z$ and $T_Y \leq T_Z$ then it is always optimal for the executive to exercise Y before Z. Given ESOs are most often granted at-the-money, if there is a bull market, an executive will typically have a portfolio of options where the more recent options have a higher strike price and a longer time to expiry. In this situation we show it is always optimal to exercise the "older" options first since they also have the lower strikes. However, if $T_Y \leq T_Z$ but K_Z ($< K_Y$) is sufficiently small, then it can be optimal never to exercise Y before Z. More generally, the optimal exercise order can *switch* during the life of the shorter dated option. These situations will arise in a bear market when portfolios will tend to have more recent options with a lower strike price but longer time to expiry.

The left panel of Figure 2 shows the exercise thresholds for a portfolio of two options, Y and Z, which are co-monotonic with $K_Y = 10, T_Y = 5, K_Z = 11, T_Z = 10$. The thresholds for option Y are solid lines, whereas those for option Z are dashed. The top pair of lines represent the Black-Scholes thresholds for options Y and Z. The lower pair of lines represent the thresholds under utility-based valuation. Red lines represent thresholds at which the *first* option and blue lines thresholds at which the *second or last* option in the portfolio is exercised. Since Z is always exercised after Y, the dashed (blue) threshold at which option Z is exercised is equivalent to a stand-alone threshold for Z. However, the solid (red) threshold for Y reflects

[22] In fact American call options with co-monotonic strikes and maturities will be exercised in order of increasing strike under very minimal assumptions on prices and preferences; see Henderson, Sun, and Whalley (2013) for details.

the presence of Z and is not equivalent to Y's stand-alone exercise threshold. In contrast, the Black-Scholes thresholds for each option do not change in the presence of other options in the portfolio and thus are the same as stand-alone thresholds for each individual option.

The right panel of Figure 2 lowers the strike of Z to $K_Z = 4$, so that the option strikes and maturities are not co-monotonic. Now we see that under our model, it is never optimal to exercise Y first. The dashed threshold for option Z is always lower than the solid (blue) threshold for option Y. As in a Black-Scholes setting, the effects of differences in times to maturity are more important when the options are closer to maturity,[23] whereas for long times to maturity, differences in time to maturity become relatively less important, so differences in strike prices determine the exercise order. In this case, K_Z is sufficiently low that the strikes determine the exercise order for the whole life of the shorter dated option, so Z is always exercised first.[24]

For some strike/maturity combinations, the tradeoff between the strike/moneyness and the time to maturity changes such that the exercise order switches from initially exercising the option with a longer time to maturity first, say Z with lower strike, to exercising the shorter dated option (with the higher strike price), Y say, first. We denote the switchover time by $T^* \in [0, T_Y]$.[25] If K_Z is sufficiently small, as above, that it is never optimal to exercise Y before

[23] Exercise thresholds change rapidly with time to maturity when the time to maturity is short.

[24] In the Black-Scholes world, we see the thresholds for Y and Z intersect close to the maturity of Y. Z is exercised first before this intersection, then immediately prior to Y's maturity it is optimal to exercise Y first. Again, however, the Black-Scholes thresholds do not change in the presence of other options, and thus the thresholds are the same as they would be for each option as a stand-alone.

[25] For a portfolio consisting of options Y and Z with $T_Y \leq T_Z$, the switchover time T^* is defined via:

$$T^* = \inf\{u \leq T_Y : e^{-\gamma(1-\rho^2)(S_u - K_Y)^+ e^{r(\bar{T}-u)}} H^Z(u, S_u) \leq e^{-\gamma(1-\rho^2)(S_u - K_Z)^+ e^{r(\bar{T}-u)}} H^Y(u, S_u)\}.$$

Z, then $T^* = T_Y$. Alternatively, if $K_Y \leq K_Z$ so we are in the co-monotonic case, then $T^* = 0$.

Figure 3 shows the exercise thresholds for a portfolio of two options, Y and Z, which are not co-monotonic and where the exercise order switches during the life of the option. Option Y has a higher strike price but a shorter time to maturity than option Z: $K_Y = 10 > 8 = K_Z$, $T_Y = 5 < 10 = T_Z$. As before, the top pair of thresholds are the equivalent Black-Scholes thresholds which cross at $T^*_{BS} = 1.09$. The optimal Black-Scholes strategy is to exercise option Z first if $0 \leq t \leq 1.09$ and option Y first if $1.09 \leq t \leq 5 = T_Y$. The optimal exercise order also switches under utility-based valuation (the lower set of thresholds) from exercising Z first to exercising Y first, though at a later time, $T^* = 4.25$. However, in contrast to the Black-Scholes setting, where the switch in the exercise order has no impact on the location of either exercise threshold, the difference in the level of each option's exercise threshold depends on whether it is exercised first or last (more generally, later) and results in a discontinuous jump in the optimal exercise threshold for each individual option, S^Y and S^Z, over the switching time, T^*. S^Z is lower before T^*, when Z is exercised first, *i.e.* before option Y, than afterwards, and so 'jumps up' over the switching time as Z becomes the last to be exercised option. Correspondingly, Y's threshold, S^Y 'jumps down' at T^* as Y becomes first to be exercised.[26]

Figure 3 also shows a dotted (green) line, which represents the threshold for option Y if option Z is exercised before T^*. Since the portfolio consists of only two options, this is the same as the stand-alone threshold for Y. We consider various scenarios for stock prices in Figure 4. In the left-hand panel, both Path A and Path B cross the lowest exercise threshold (for option Z)

[26]More generally, at any switching time, any option whose order switches will have a discontinuity in its exercise threshold. The threshold will be significantly higher after the switching time if the option becomes later in the exercise order, and will decrease significantly if it becomes earlier in the exercise order.

before T^*, so option Z is exercised first. The executive's portfolio then consists just of option Y, which has exercise threshold as for a stand-alone option, given by the combination of the solid (blue) line for $t < T^* = 4.25$ and the dotted (green) line for $T^* \leq t < T_Y$. Stock path A crosses this line so option Y is exercised, whereas stock path B remains below it and option Y expires unexercised. In contrast, the stock paths in the right-hand panel do not cross the lowest exercise threshold before T^*. Between T^* and T_Y it is optimal to exercise Y first if the stock price crosses the lower solid (red) line. Stock path A does, so option Y is exercised; stock price B does not, so option Y expires unexercised. Once option Y has either expired or been exercised the executive's portfolio consists solely of option Z and is exercised only if (as is the case for both stock price paths) the stock price crosses its stand-alone exercise threshold given by the dashed (blue) line before its maturity T_Z. Thus which exercise threshold is relevant for option Y between T^* and T_Y depends on the stock price behavior before T^* (it is the dotted (green) line if Z has been exercised before T^*, or the lower, solid (red) line if not).

The exercise threshold for either one of the two options in this case exhibits significantly different characteristics from those under Black-Scholes. Not only is the threshold no longer monotonically increasing with time to maturity but strikingly, it exhibits a discontinuity whenever the exercise order changes.[27]

The combinations of strike price and maturity for each form of exercise strategy are il-

[27]For the relatively longer dated option, Z, the exercise threshold whilst it is the first-to-be-exercised option increases as time to maturity decreases and, particularly close to the switchover time, at an increasing rate, so the exercise threshold is convex in time to maturity. After the discontinuity in the threshold when the exercise order switches, once the option has become the last-to-be-exercised, the threshold may still increase for long enough times to maturity. The relatively shorter dated option, Y, can start as the last-to-be-exercised with a threshold which may increase or decrease with time to maturity.

lustrated in stylized form in Figure 5. This illustrates the optimal strategy for a portfolio consisting of an option with strike price K and maturity T together with our base-case option Y, with K_Y and T_Y for different combinations of (K, T). The form of the strategy depends on which region, A - F, the strike price and maturity of the second option in the portfolio fall into relative to option Y. The vertical and horizontal quadrant lines represent the boundaries of the co-monotonic regions. So in the top right quadrant, B, Y is always exercised first and the later maturity option is exercised last, as long as both options are still alive, *i.e.* $T^* = 0$. Similarly in the bottom left quadrant E where $K \leq K_Y$ and $T \leq T_Y$, whilst both options are still alive, the earlier maturity option is exercised first and option Y is exercised last. The curved solid line represents, for each maturity T, the maximum, if $T > T_Y$, (minimum if $T < T_Y$) strike price for which it is always optimal to exercise the option with the lower strike price first (even though this has a longer time to maturity).

In the top left and bottom right quadrants, the effects of strike prices and maturities pull in opposite directions. In areas A and F, the impact of the difference in strike prices dominates, so in area A, option Y is always exercised before the shorter-maturity option due to Y's lower strike price, and in F, option Y has higher strike price and is always exercised last. In these areas, $T^* = \min(T_Y, T)$. In areas C and D the strike price difference dominates initially, but there is a switch in the optimal exercise order at $0 < T^* < T_{min}$ so for times close to T_{min} the differences in maturity times dominate.[28]

[28] Particularly for short times to maturity, the boundary between regions B/D and C/E *i.e.* the strike price at which it is just optimal to exercise the shorter-maturity option first for all remaining $t < T_{min}$ may differ from the horizontal lines shown in the figure, so region B may extend below K_Y and region E may extend above K_Y. Horizontal lines represent boundaries for *all* times.

The overall effect is that in the darker shaded regions A and B, the optimal strategy whilst both options are still alive is to exercise Y first, whereas in the lighter shaded regions E and F where there is a choice, it is optimal to exercise Y last. In regions C and D, which option it is optimal to exercise first depends on when the stock price first crosses the threshold of whichever is the first-to-be-exercised option. When the exercise order changes at T^*, there is a jump in the thresholds for each individual option. Only in the co-monotonic regions, B and E, are there no jumps in either threshold.

Within each region, although the form of the exercise strategy remains the same, the threshold for any option except the last-to-be-exercised varies with the composition of the remainder of the executive's portfolio. For example, Figure 6 shows the effect of different strike prices for a longer dated option, Z, with $T_Z = 10$, on the exercise threshold for a 5-year option Y with $K_Y = 10$ when held together in a portfolio. The top, solid line is the threshold if Y is always exercised last, which equals the stand-alone threshold. This is Y's threshold if K_Z is sufficiently low (any point in region F in Figure 5) and it is unaffected by changes in K_Z. As K_Z increases, it becomes optimal to exercise Y first when it is sufficiently close to maturity (region D). For example, for $K_Z = 8.5$, the switchover time for these parameter values is $T^* = 3.19$, so before then, Y is still exercised last and the threshold coincides with the top, stand-alone threshold; however for $T^* < t < T_Y$, it is optimal to exercise Y first at a much lower threshold. As K_Z increases further to $K_Z = 9$, there are two effects. Firstly, Y is optimally exercised first over a longer period, so the switchover time, T^*, decreases (to 1.39). Y's threshold is the lower, first-to-be-exercised threshold for longer. Secondly, the first-to-be-exercised threshold itself is slightly higher than when $K_Z = 8.5$. This reflects the lower *marginal* reduction in unhedgeable risk on Y's exercise due to Z's lower moneyness. When $K_Z = 10$, the options are co-monotonic so it is always optimal to exercise Z first (Region B with $T^* = 0$). However the effect of Z's moneyness on Y's first-to-be-exercised threshold continues to hold. So Y's first-to-be-exercised threshold for $K_Z = 10$ is higher than that for $K_Z = 9$ and that for $K_Z = 11$ is higher still.

Thus far we have illustrated the forms of the exercise strategy and the resulting exercise thresholds for particular parameter values. The potential forms of the exercise strategy, the shape of the regions of relative strike prices and maturity dates, within which each strategy applies, and the general shape of the exercise thresholds all remain unchanged when parameter values are varied. Many of the new features in the shape and (lack of) continuity of the exercise thresholds *vs* time are due to the effect of the executive's exposure to unhedgeable risk. So changes in parameter values which increase the magnitude of the effect of this exposure to unhedgeable risk increase the effects on the exercise thresholds. If the magnitude of the effect increases, all exercise thresholds decrease because of the increased cost; however, due to diminishing returns to scale the exercise thresholds of the last-to-be-exercised options decrease more than those of the second-to-last (first) to be exercised options, decreasing the magnitude of the jump in the thresholds when the exercise order changes. More importantly, since the effect of unhedgeable risk exposure increases with the length of time it must be borne, the exercise thresholds for longer-maturity options decrease more then those of shorter maturity options, both when last and when first-to-be-exercised. Hence the switchover time, T^*, increases, so longer-maturity options are exercised first for a longer initial period. In unreported simulations, these effects are found to hold consistently for increases in risk aversion, γ, and in the risk-free rate, r, which increases the impact of the unhedgeable risk exposure over time. Increases in the dividend yield, q, have similar effects, since exercise thresholds decrease due to the increased foregone dividend yield, though less markedly since the increase in this effect over time is less pronounced.

3.3 Implications

Exercise thresholds are important because they are observable: the exercise of an ESO reveals one point of the empirical exercise threshold of the first option to be exercised from the executive's ESO portfolio. This is usually described in terms of the relative moneyness on exercise or S^{π}/K_{π}. What are the implications of our model for the relative moneyness on exercise (hereafter moneyness) of ESOs held in portfolios, specifically the first-to-be-exercised moneyness which is observed in practice?

Most implications follow directly from the equivalent results for exercise thresholds themselves. So for example, for any option in a portfolio (as long as it is not exercised last), the moneyness on exercise is lower than the moneyness on exercise for the same option held as a stand-alone. A particular ESO's, say Y's, moneyness on exercise decreases if the executive is granted additional options which are optimally exercised after Y. It also varies with the terms, (K, T), of other options in the executive's portfolio which are after it in the optimal exercise order.

Like the threshold itself, the moneyness on exercise of the first-to-be-exercised ESO is generally non-monotonic with respect to time to maturity, first increasing, potentially at an increasing rate, and then decreasing as time gets closer to the maturity of each option in the portfolio. It jumps when the identity of the first-to-be-exercised option changes during the portfolio's life: generally jumping up on exercise or expiry of the first-to-be-exercised ESO and jumping down if there is additionally a switch in the optimal exercise order whilst the portfolio composition remains unchanged.[29] For example, the portfolio of Y and Z in Figure 3 with $K_Y = 10, T_Y = 5$ and $K_Z = 8, T_Z = 10$ has switching time, $T^* = 4.25$. Z is exercised first before T^*, Y afterwards. At T^* the exercise threshold of the first-to-be-exercised option, at which the executive is indifferent between exercising Y first or Z first, is 12.49. The relative moneyness on exercise of the first option to be exercised thus drops from $12.49/K_Z = 1.56$ for exercising Z just before the switchover to $12.49/K_Y = 1.25$ for exercising Y just afterwards.

The non-monotonicity and jumps in the relative moneyness on exercise of the first-to-be-exercised ESO in a portfolio makes it difficult to draw empirical predictions about the determinants of observed moneyness on exercise which hold in all circumstances. However, empirical studies of ESO exercise in practice often work with exercise rates or propensity to exercise, and we show how our model can be used to explain empirical exercise patterns around new option grants (Section 3.3.1) and other empirical findings on exercise rates (Section 3.3.2).

3.3.1 Immediate exercise of existing options: the impact of a new grant

Several empirical studies have found a relation between the executive's propensity to exercise existing options and the arrival of a new grant of options. Specifically, Ofek and Yermack (2000) find that executives with larger exposure to firm-specific risk (larger portfolios of stocks and options) were more likely to exercise some existing options after new option grants than executives with lower prior exposure. Similarly, Klein and Maug (2010) also find that new option grants increase executives' propensity to exercise, and Srivastava (2011) finds that existing options' time to maturity on exercise increases with the size of new option grants.

[29] Recall that after the switch, the option exercised first will have a shorter time to maturity and higher strike price.

Under a Black-Scholes framework, a new grant does not alter the thresholds of the existing options, and thus cannot alter the likelihood that any existing options will be exercised. In contrast, the increase in unhedgeable risk brought to the portfolio by the new grant means that it may be optimal to exercise the existing option immediately in our setting. Indeed, as we demonstrate, the decrease in the relative moneyness required to trigger immediate exercise once executives' risk aversion and portfolio effects are taken into consideration is consistent with these results.

Consider an executive with an existing option Y who receives a new option grant Z. If K_Z is sufficiently high that it is optimal to exercise Y first as soon as the new option, Z, is granted, Y's threshold once the new option has been added to the portfolio will be much lower than its threshold in the absence of the new option. New options are often granted at-the-money. So if option Y is sufficiently in-the-money, then an unexpected new at-the-money option grant may induce Y's immediate exercise. The level of moneyness required for exercise once Y is part of the new, larger portfolio will be much lower than that required for Y as a stand-alone option.

For example, consider option Y with $K_Y = 10, T_Y = 5$. For parameters given by $\gamma = 0.1, r = 0.05, q = 0.02, \eta = 0.4$, Y is exercised immediately under Black-Scholes if the stock exceeds 48.8, or equivalently, Y's relative moneyness S/K_Y exceeds 4.88. The equivalent stand-alone threshold for Y in moneyness terms is 1.65. (This takes into account only the unhedgeable risk associated with option Y itself.) Now consider receiving another option Z with $T_Z = 10$. The presence of Z in the portfolio dramatically lowers Y's moneyness threshold such that Y is exercised immediately if $S/K_Y > 1.39$. Hence the probability of immediate exercise of existing options at the date of a new option grant is higher once risk aversion and unhedgeable risk *with portfolio effects* are taken into account. Factors that increase the impact of unhedgeable risk will reduce Y's moneyness threshold and raise the likelihood that the existing option is exercised immediately at the grant date of the new option.

3.3.2 Empirical literature on ESO exercise rates

Recently, several papers have developed empirical models of ESO exercise. Carpenter, Stanton, and Wallace (2012) use a GMM-based methodology to estimate an exercise function and use this to value options. Armstrong, Jagolinzer, and Larcker (2007) and Klein and Maug (2010) have used a hazard rate model to estimate exercise rates. Carpenter, Stanton, and Wallace (2012) recognize the complexity of stock option portfolios in practice in the methodology they use, but do not include specific controls for the composition of the remainder of the portfolio. Only Klein and Maug (2010) and Armstrong, Jagolinzer, and Larcker (2007) test explicitly for portfolio effects and find strong evidence of its explanatory importance. Klein and Maug (2010) find that most of the options exercised have a time value which is lowest or close to lowest of all options in the executive's portfolio[30] and Armstrong, Jagolinzer, and Larcker (2007) find relationships between the exercise rate of a particular option and the intrinsic values of *other* options in an executive's portfolio. Both are consistent with the portfolio effects we describe.[31]

[30]They calculate the time value using Black-Scholes world values, so the optimal exercise ordering may sometimes differ from the ordering of the Black-Scholes world time values on exercise. However in many cases (for e.g. co-monotonic) the two will coincide, so this is broadly consistent with our model.

[31]Armstrong, Jagolinzer, and Larcker (2007) show that the exercise rate of a particular option, say option Y, is decreasing in the overall intrinsic value of the executive's other option grants, if this is positive, and is also decreasing in the absolute value of this intrinsic value if the intrinsic value is negative. Although seemingly inconsistent, both of these results can be explained using our model. To illustrate the main effects, we approximate the remainder of the option portfolio by a holding with common strike price and time to maturity. First, if the intrinsic value of the executive's other options is negative, they are out-of-the-money, and

4 Shareholder costs

Shareholders do not face the same restrictions on trading as executives, so they value ESOs assuming perfect hedging (Black-Scholes), but taking account of the executive's optimal exercise threshold (see Carpenter (1998), Carpenter, Stanton, and Wallace (2010) amongst others). The cost of each option, $C^{\pi}(t,s)$, is calculated using the Black-Scholes equation whilst the stock price remains below the optimal exercise threshold, S^{π}.[32] The portfolio cost, $C^{\Pi}(t,s)$, is given as the sum of the individual option costs, C^{π}.[33]

so are most likely to be exercised after Y. Increasing the strike price of the executive's other options increases the magnitude of the other options' (negative) intrinsic value. This decreases the unhedgeable risk associated with these other options and therefore also decreases the reduction in the unhedgeable risk on exercising Y. This leads to an increase in Y's optimal exercise threshold or equivalently a decrease in Y's exercise rate, as in Armstrong, Jagolinzer, and Larcker (2007). In the other case when the intrinsic value of the executive's other options is positive, Y may be exercised before or after the other options. The primary effect is that a decrease in the strike price of the other options, which increases their intrinsic value, can lead to a change in the optimal exercise order, so Y is exercised first before the change in the strike price but is exercised later (last) afterwards. We show that the optimal exercise threshold for an option jumps up when it moves to a later position in the exercise order, due to the sudden decrease in the change in unhedgeable risk arising from the exercise of this particular option. This is equivalent to a decrease in the option's exercise rate, as found by Armstrong, Jagolinzer, and Larcker (2007).

[32] The pde and boundary conditions for $C^{\pi}(t,s)$ are given in the Appendix.

[33] For example, $C^{Y,Z}$ is the cost of the portfolio of options Y and Z, C^{Y} is cost of Y when considered as part of the portfolio, C^{Z} is the cost of Z as part of the portfolio. The costs can also be calculated for individual options on a stand-alone basis, denoted for example, $C^{Y}_{\mathcal{S}}$, by replacing the threshold S^{Y} with the stand-alone threshold $S^{Y}_{\mathcal{S}}$.

4.1 Portfolio effects: discount relative to Black-Scholes cost

It is well known that the costs of individual ESOs are lower than their equivalent Black-Scholes values due to the executive's optimal early exercise in the presence of unhedgeable risk (Lambert, Larcker, and Verrecchia (1991), Carpenter, Stanton, and Wallace (2010), Leung and Sircar (2009)). Decreases in the executive's optimal exercise thresholds due to the presence of other options in her portfolio automatically decrease the cost of those options to the shareholders still further. So the cost of any option when held in a portfolio is generally lower than its cost as a stand-alone.[34] Furthermore, the overall shareholder cost of any portfolio of options is lower than the sum of the costs of each option evaluated on a stand-alone basis, since all bar at most one of the options making up the portfolio are exercised earlier and thus have lower individual costs.

Thus portfolio effects increase the magnitude of the reduction in shareholder costs relative to Black-Scholes for *any* ESO portfolio and for most individual ESOs: a greater discount is required from the Black-Scholes world cost to obtain the shareholder cost for a given level of executive risk aversion.[35] For example, consider a five year at-the-money ESO Y on a stock with current stock price $S = 10$, dividend yield $q = 5\%$ and unhedgeable volatility $\eta_e = 40\%$, exercisable at any time before its maturity, with risk-free rate $r = 10\%$. The Black-Scholes world value of this grant is 3.48. For a stand-alone ESO grant, taking into account the executive's risk aversion, $\gamma = 0.2$, reduces Y's cost by 47% to 1.85. However, if the executive owning Y also owns a second option, Z, which is also vested and has the same strike price but with 10

[34]The only exception is if the option is optimally exercised last throughout its remaining life (when the cost is unaffected by the presence of other options in the portfolio).

[35]Equivalently, a lower level of risk aversion is required to achieve the same discount.

years to maturity, then the cost of Y decreases further to 1.04, a 70% reduction relative to the original Black-Scholes world value. The cost of the ESO portfolio, 3.12, (the sum of the portfolio costs), is also less than the sum of the costs of each ESO as stand-alone options, 3.93, again increasing the proportional discount relative to Black-Scholes world values (from 49% to 60%).[36]

These proportional discounts can be significant in magnitude. Table 1 shows proportional discounts increase, the greater the portfolio effect on the exercise threshold, i.e. the greater the unhedgeable risk (η_e), the greater the cost of that risk to the executive (measured by the risk aversion, γ) and the greater the potential reduction in threshold (the higher the original threshold), so the lower the dividend yield, q (the higher $r - q$). Overall, portfolio effects increase discounts for the costs of ESOs relative to Black-Scholes values, even for low levels of risk aversion (for $\gamma = 0.01$, the discount for individual options almost quadruples from 3% to 11%, and the discount for the portfolio almost doubles from 4% to 7%). The presence of one single other option in a portfolio can decrease the cost of an individual ESO to the shareholders to less than a quarter of its Black-Scholes world value, and the cost of a portfolio of two options can be less than one third of its Black-Scholes world value.

Approximations for shareholder costs based on Black-Scholes methodology are permitted by FASB and are in common useage (Armstrong, Jagolinzer, and Larcker (2007), Carpenter, Stanton, and Wallace (2012)). The typical Modified Black-Scholes method (FAS 123R) uses the Black-Scholes formula with option expiry set equal to the expected time to exercise.[37] For

[36]Since the options are co-monotonic, the earlier maturity option, Y is always exercised first, so the cost (in the portfolio) of the ten year option is unaffected in this case.

[37]Alternatively, approximations based on a fixed maturity adjustment such as an average between vesting and maturity dates can be used by companies with limited exercise data (SAB 110).

individual options, Carpenter, Stanton, and Wallace (2010) find this FASB approximate cost can under- or over-state utility-based option costs, sometimes by large amounts, depending on firm characteristics. Furthermore, they found that varying firm and executive characteristics changed both the utility-based cost and the FASB approximation (through the expected time to exercise), but at different rates, and not always in the same direction. We find portfolio effects lower exercise thresholds, particularly for options earlier in the exercise order. This will reduce the expected time to exercise for these options, and thus their FASB approximations, as well as reducing the utility-based option cost. However, the expected time to exercise for a particular option, and hence its FASB approximation, will vary depending on the composition of the remainder of the executive's portfolio. Estimates of expected time to exercise for current options are often in practice based on the average actual time to exercise of prior options. However portfolio effects mean these estimates will need to be adjusted to take account of both the prior and the current ESO portfolio composition. Moreover, we find that the proportional reduction in ESO costs due to portfolio effects is greater for firms with higher levels of unhedgeable risk and lower dividend yields. Carpenter, Stanton, and Wallace (2010)'s results suggest the FASB approximation tends to overstate the utility-based ESO cost for a single option by more for these types of firm. It is thus likely that portfolio effects will increase the overstatement of the FASB approximate cost for options, even if an accurate portfolio based estimate of the expected time to exercise is used. More generally, our results imply any approximate method for estimating the costs of ESOs which does not adjust for portfolio effects, even if it gives reasonable values for a stand-alone option, can lose significant accuracy if used to value the same option held in a portfolio.

4.2 Portfolio composition: moneyness effects

In the above example, the two options were co-monotonic, so portfolio effects reduced the cost of the shorter dated option, Y, leaving the cost of the longer-dated option unchanged. However, we saw in Section 3 that more generally the exercise order depends on the relative strike prices and times to maturity, so can change during the portfolio's life. This means the cost to the shareholders of an individual option grant to a particular executive depends on the composition of the remainder of that executive's portfolio.

Consider a portfolio of our five-year at-the-money ESO Y combined with a ten-year option, Z, as before, but now allow Z's strike price to vary. We saw in Section 3 (Figure 6) that if K_Z is sufficiently small, Y's exercise threshold is unaffected, since Z is exercised first and Y last, but that as K_Z increases, it becomes optimal for the executive to exercise Y first, at a lower threshold than the stand-alone one, when Y is sufficiently close to maturity. As K_Z increases further both the length of time when Y is exercised first and also Y's exercise threshold whilst it is exercised first increase. This translates directly into an effect on Y's cost, which is shown as a function of K_Z in the top left graph in Figure 7. Other graphs in Figure 7 show how the optimal switchover time T^* (top right-hand graph) and cost of option Z (bottom left graph) vary with K_Z.

The horizontal dotted line in the top left graph represents Y's cost as a stand-alone option. This stand-alone cost is an upper bound on the cost of Y when it is part of a portfolio with Z, which is shown by the lower, solid line. The two costs coincide for low K_Z, when Y is exercised last (so T^* in the top right graph is T_Y). Once K_Z is sufficiently high that it becomes optimal to switch to exercising Y first close to Y's maturity, increasing K_Z decreases the switchover time T^*. It also decreases the cost of Y because the decrease in Y's cost due to the longer period that Y is exercised at its much lower, first-to-be-exercised threshold dominates the effect of the increase in the first-to-be-exercised threshold itself. For higher K_Z,[38] once Y is always exercised first, only the second effect, due to the decrease in unhedgeable risk eliminated on exercise, remains, so Y's cost increases gradually with K_Z.[39] So the cost of one option in the portfolio, Y, varies non-monotonically with the strike price of the other option in the executive's portfolio. This is in stark contrast to a Black-Scholes world, where option values are independent of the portfolio in which they are held.

Portfolio effects also change how the cost of an option varies with its own strike price. In a Black-Scholes world, options are monotonically decreasing in their own strike price, $\frac{\partial C^{BS}}{\partial K} < 0$, and this continues to hold for stand-alone options in a utility-based setting as shown by the top dashed line in the bottom left graph in Figure 7. However, when Z is held as part of a portfolio this is no longer necessarily true: its cost as a function of its own strike price is given by the solid line, which has a local minimum close to the strike price of the other option in the portfolio.[40]

[38] Strictly, it is always optimal to exercise the shorter-dated option first only when the options are co-monotonic, i.e. for $K_Z \geq K_Y$ if $T_Z \geq T_Y$. However the co-monotonicity is a necessary condition only for this to be true for all possible times to maturity. In practice when the shortest time to maturity is small, if K_Z is slightly below K_Y the switchover time occurs long before T_Y and has thus effectively already occurred, so $T^* = 0$ for $K_Z \geq K_Z^*$ where $K_Z^* < K_Y$ varies with T_Y. In the example, $K_Z^* = 9.23 < 10 = K_Y$

[39] In the limit as $K_Z \to \infty$, there is no additional risk and $C^Y \to C_S^Y$.

[40] The graph can be split into four regions. Firstly, for K_Z sufficiently lower than the current asset price $S = 10$, it is optimal for the executive to exercise Z immediately, so the option cost equals its payoff. This is the case whether it is held alone or as part of a portfolio. However, immediate exercise is optimal for higher values of K_Z (up to $K_Z = 6.95$ in this case) when the option is held as part of a portfolio, due to the executive's lower valuation of the portfolio before exercise because of the larger associated unhedgeable risk. Secondly, for higher

Finally, since the cost of the whole portfolio of ESOs granted to an individual executive is the sum of the cost of each option, evaluated as part of the portfolio: $C^{Y,Z} = C^Y + C^Z$, the portfolio cost also varies non-monotonically with the strike prices of its constituent options.

4.3 Discounts arising from portfolio composition

The presence of other options in the executive's portfolio reduces the cost of each individual option grant in the portfolio (as long as it does not remain the last to be exercised over the whole of its life). It is thus valued at a discount relative to its stand-alone cost, simply because it is part of a portfolio. Moreover the discount varies with the strike prices and times to maturity of the other options in the executive's portfolio.[41] The "portfolio proportional discount" or proportional reduction in cost relative to the stand-alone cost due to portfolio effects can be measured as $1-\mathcal{R}_\Pi$ where $\mathcal{R}_\Pi = C^\Pi/(\sum_i C_{\mathcal{S}}^{\pi_i})$ is the cost evaluated as a portfolio as a proportion of the (sum of) stand-alone cost(s). For a portfolio of two or more unexercised options this is

K_Z, it may in principle still be optimal to exercise Z before Y, though the optimal threshold is greater than the current stock price $S^Z > S = 10$. The cost of Z thus reflects its lower threshold as the first-to-be-exercised option; this decreases as K_Z increases due to standard decreasing moneyness considerations. Thirdly, as K_Z increases further, at some point it becomes optimal to switch to exercising Y first close to its maturity. This generates an additional effect when K_Z increases: the decrease in switchover time T^* increases the length of time Z is exercised last, at a higher theshold, which increases Z's cost. Eventually this dominates so the cost of Z has a local minimum before rising to equal its cost as a stand-alone (the fourth region when $T^* = 0$, for $K_Z \geq 9.23$).

[41]Figure 7 shows the effects of varying the strike price of a longer-dated option. Similar effects are found varying the strike price of a shorter-dated option. Details are available from the authors on request.

always non-zero, *i.e.* $\mathcal{R}_{\Pi} < 1$.[42]

The bottom right graph in Figure 7 plots $\mathcal{R}_{Y,Z}$ against K_Z for our two option portfolio of Y with longer dated Z. The discount for the portfolio as a whole arises from option Z when K_Z is sufficiently low, from option Y when K_Z is high (e.g. $K_Z > K_Y$), and from both options when K_Z is slightly lower than K_Y (in the switching region D from Figure 5). The proportional discount for the portfolio as a whole reaches a maximum value (of over 19%) for some K_Z in the switching region, D, but remains above 10% for a wide range of K_Z.[43,44]

[42]At any time all but one of the options in a portfolio are exercised at lower thresholds and thus have lower costs than if held as stand-alone options.

[43]$\mathcal{R}_{\Pi} = 1$ only for low K_Z where Z is so far in-the-money (recall $S = 10$) that it is optimal to exercise Z immediately.

[44]For low K_Z it is optimal for the executive to exercise Z immediately, so the cost of Z as part of the portfolio equals its payoff on exercise. If Z's optimal exercise threshold on a stand-alone basis is higher than the current stock price, S, then the portfolio discount is entirely due to the time value of Z's cost as a stand-alone option. This increases as K_Z, and hence the difference between Z's stand-alone threshold and S increases. Alternatively, if K_Z is sufficiently high, the portfolio discount is entirely due to Y's lower threshold as the first-to-be-exercised option. Increases in K_Z reduce Z's moneyness and hence reduce the marginal unhedgeable risk on exercise of Y, reducing the discount. (This dominates over the additional effect of the increased weight of Y in the portfolio due to the reduction in Z's moneyness and hence cost). For K_Z in the intermediate region where it is initially optimal to exercise Z first, but not immediately, and the optimal exercise order may change during the options' lives, a number of effects arise when K_Z increases. The increase in the marginal unhedgeable risk reduction on exercise of the first-to-be-exercised option continues to reduce the portfolio discount, but the increase in K_Z also changes both the length of time for which Y is exercised first and the relative weights of Y and Z in the portfolio. Which effect dominates depends on the parameter values. Overall, the proportional discount for the portfolio as a whole is maximized within this region where the options are non co-monotonic and initially the longer dated option is exercised first. Moreover, the further away (in (K, T) space) from the switching region, (e.g. higher $K_Z > K_Y$ or lower $K_Z \ll K_Y$), the lower the proportional discount.

In Table 2 we report, for each option, Y and Z, and for the portfolio as a whole, $Y+Z$, the portfolio cost as a proportion of the (sum of) stand-alone cost(s), $\mathcal{R}_Y \equiv C^Y/C^Y_{\mathcal{S}}$ *etc.*, for different strike prices and maturities K_Z and T_Z. Panel A considers the same parameter values as in the earlier example; panels B - E show the effect of varying each key parameter.[45]

The Table shows that portfolio effects can reduce shareholder costs for the portfolio as a whole by more than 20% relative to the costs of the options as stand-alones, and by over 40% for individual options. Moreover, the proportional discount remains significant for a wide range of portfolios (K's and T's) and parameter values. The portfolio proportional discount is larger for firms with more risk-averse executives, with high idiosyncratic volatility and low dividend yield, when the risk-free rate is high, and for executives with portfolios where earlier-maturity options have higher strike prices.[46]

Firms should thus recognize that evaluating the cost of each ESO individually overstates the total cost to the shareholders of a portfolio of options granted to the same executive. The inherent non-linearity induced by portfolio considerations means the cost of an ESO portfolio needs to be calculated as a whole. Moreover, the costs of individual ESOs depend on the portfolio they are a part of, and so may change if the composition of the executive's portfolio changes. This raises the question of how a firm should evaluate the cost of a new option grant.

[45] For each combination of (K_Z, T_Z), either it is always optimal to exercise the shorter-dated option, Y, first, so $\mathcal{R}_Y < 1$ and $\mathcal{R}_Z = 1$, or always optimal to exercise the longer dated option, Z, first so $\mathcal{R}_Y = 1$, $\mathcal{R}_Z < 1$, or the optimal exercise order changes, so both $\mathcal{R}_Y, \mathcal{R}_Z < 1$. In the Table, since $T_Z > T_Y$, these possibilities correspond to regions B, F and D in Figure 5.

[46] So the optimal exercise order changes and both options will optimally be exercised first in different time periods.

4.4 Incremental cost of new options

If the executive receives an additional option grant, portfolio size and total unhedgeable risk increase. The overall portfolio cost increases, but the discount relative to the sum of the stand-alone costs also increases. Consider the effect of an unanticipated new grant of an option, Z, to an executive with an existing portfolio of options (consisting of option Y). Before option Z is granted, the cost to the shareholders of the executive's portfolio would be given by $C_{\mathcal{S}}^{Y}$, the cost of option Y as a stand-alone. After option Z is granted, the cost of the portfolio would be $C^Y + C^Z$, where both costs take into account the presence of the other option in the portfolio. Thus the incremental cost of the new grant to the shareholders is $C_{Z|Y}^{marginal} \equiv C^Y + C^Z - C_{\mathcal{S}}^{Y}$. This is plotted in Figure 8 for different strike prices for the newly granted option, K_Z, assuming option Y is currently at-the-money. Since C^Y and C^Z vary non-monotonically with respect to K_Z, as shown in Figure 7, the incremental cost of option Z also depends on K_Z in a more complex way than would be expected in a Black-Scholes world. In particular, it is no longer monotonically decreasing with respect to K_Z: there is a local minimum incremental cost at a strike price slightly smaller than K_Y. Equivalently, the local minimum incremental cost for the new option occurs if the option is granted slightly in-the-money, and the incremental cost of the new option is less sensitive to changes in its strike price when it is granted close to the money.

The left-hand graph in Figure 9 shows $\mathcal{R}_\Pi$ or (1 - "portfolio proportional discount") for the incremental cost of Z *vs* K_Z, assuming option Y is at-the-money i.e. $S = 10$. Note that as long as it is optimal to hold both options, the proportional discount is non-zero. It has a similar shape to the graph for the portfolio as a whole (bottom right hand graph in Figure 7), but at a lower level.[47] The maximum proportional discount occurs if the new option is granted slightly in-the-money and remains above 20% for most choices of strike price and time to maturity.

It is common for firms to grant all ESOs at-the-money with the same time to maturity (10 years)(Carpenter, Stanton, and Wallace (2012)). The right-hand graph in Figure 9 shows how $\mathcal{R}_{\Pi}$, and hence the portfolio proportional discount, varies with the stock price for an incremental option Z granted at-the-money with 10 years to maturity to an executive who already holds an ESO with strike price 10 and five years to maturity. This proportional discount is also non-monotonic with respect to $S = K_Z$. It is non-zero as long as it is not optimal to exercise the existing option, Y, immediately, and over 10% for all stock prices apart from those for which the existing option Y is far from the money (more than 30% in- and out-of-the-money for these parameters).

So firms also need to recognize that the cost of the same option grant to otherwise identical executives, with in particular the same level of risk aversion, will not be the same if the portfolio of other options they hold differs. The magnitude of the portfolio proportional discounts suggest any approximate method for evaluating the cost of ESOs needs to be flexible enough to take these portfolio considerations into account. Increasing the size of the existing portfolio (adding options) will decrease the cost of making a new option grant to the executive. So the cost to the shareholders of an option grant to an executive is highest when the executive has no outstanding unexercised ESOs.

[47] This is unsurprising since both incorporate the reductions in cost of all options in the portfolio due to the additional option grant in the numerator but the denominator is smaller for the individual option than for the portfolio as a whole.

5 Robustness

5.1 Larger portfolios

As the number of options increases, the complexity of the numerical solution increases significantly because of the number of potential exercise combinations which need to be considered, thus we leave detailed investigation of optimal exercise strategies for larger portfolios for further work. However, many of the results for the two-option case will carry over to larger portfolios, potentially with even greater effects. This is because of the common underlying cause of the results, that the effect of unhedgeable risk increases non-linearly with portfolio size. This implies that thresholds for a given option when it is $(n+1)$st-to-last to be exercised are distinctly lower than the threshold for the same option when one option is removed from the remainder of the portfolio, so it is the nth-to-last to be exercised, because the decrease in unhedgeable risk associated with the option exercise is greater, the larger the remaining portfolio. This means that the threshold at which an executive exercises her first option decreases when further options are added to her portfolio.[48,49]

Exercise thresholds for a particular option in the larger portfolio can exhibit the same characteristics as in the two-option case: the threshold can increase, even at an increasing rate, as time to maturity decreases, and will jump if its position in the optimal exercise order changes.

[48]If the new option is the first-to-be-exercised, its threshold is automatically lower than the lowest threshold of the existing options. If at some time it is not the first-to-be-exercised, thresholds for all options exercised before it, including the first-to-be-exercised, are reduced since the portfolio remaining after their exercise now includes the new option.

[49]We know for a single grant of identical options, as the number of options becomes infinite, the individual exercise thresholds will limit to the option strike K, see Henderson and Hobson (2011).

However the jumps no longer affect all options in the portfolio: the exercise threshold for options not involved in a particular switch in the exercise order remain continuous across the switching date. An example of exercise thresholds for a portfolio of three options is given in Figure 10. Depending on the combinations of strike prices and maturity dates, there can be multiple switches in the exercise order and jumps in exercise thresholds. For portfolios with comonotonic strike prices and maturities, there are still no jumps and the exercise order remains unchanged throughout the portfolio's life.

In general the exercise order depends on the relationship between the strike prices and maturities of options in the portfolio. Holding all else equal, options with higher strike prices and longer times to maturity are exercised later; where these effects conflict, differences in strike price are more likely to dominate when all the times to maturity are long. However, an option with a higher strike price but shorter maturity can be exercised out of strike-price order when its time to maturity is short, since its thresholds decrease rapidly as its time to maturity decreases to zero.

These effects carry through to shareholder costs. At most one option in a portfolio has the same cost as its stand-alone cost (if it is always last-to-be-exercised); costs for all other options are lower because they are no longer exercised at their stand-alone, equivalently last-to-be-exercised thresholds. Moreover, the cost of each option in a portfolio is affected by its relationship with the strike prices and maturities of all the other options in the portfolio. As a function of its strike price, the shareholder cost of a new option can have multiple local minima.[50] Costs of all existing options will also depend non-linearly on the strike price of the

[50]If it has a sufficiently low strike price, the new option will be exercised first, say $(n+1)$st to last. However, as its strike price increases, its position in the exercise order changes to nth-to-last, $(n-1)$st-to-last, *etc.*.

new option, generally with a similar shape to the graph in the two-option case (the top left graph in Figure 7). Hence the incremental cost to the shareholders of a new option grant will also be non-linear with potentially multiple local minima.

When an option portfolio's size is increased by the addition of an option, the cost of the overall portfolio to the shareholders will increase, but by much less than the stand-alone cost of the new option. The proportional portfolio discount for new options, or equivalently the difference between the incremental cost of a new option taking account of portfolio effects and its cost as a stand-alone will increase.[51]

Overall then, portfolio effects on both exercise thresholds and costs increase when options are added to an ESO portfolio, but the nature of the effects remains broadly the same as in the two option case we considered in detail in Sections 3 and 4.

Eventually, for high enough strike prices, it becomes the last to be exercised. The exercise threshold and shareholder cost of an option exercised later in the exercise order increases, because of the lower level of unhedgeable risk associated with the exercise decision. However the cost of an option with a given place in the exercise order decreases with its strike price. Together this may lead to a local minimum for each switch in the exercise order, as in the two-option case.

[51]On the addition of an option to an existing portfolio, the costs of options in the existing portfolio which will now be exercised after the new option are unaffected by its presence; however any options now exercised before the new option at any time will have a lower cost. If the existing portfolio increases in size and, when added, the new option is still exercised last, its threshold and hence cost will be lower. If it is not exercised last in the larger combined portfolio, then at least one option in the portfolio is exercised at a lower threshold than it would be in the combined portfolio before the increase in size, lowering the incremental cost.

5.2 Further Modeling Issues

We finish by discussing several issues concerning the robustness of our results to inclusion of vesting, employment risk, restricted stock and performance-based compensation. We also comment on alternative preferences.

In practice, ESOs commonly have vesting restrictions in order to tie the executive to the company and provide incentives.[52] Very few of the works taking unhedgeble risks into account also consider the impact of time-vesting[53] however, Leung and Sircar (2009) find shareholder costs increase with the vesting period. In our framework, vesting restrictions would not alter the optimal exercise ordering once the options had vested, and it would not change the thresholds at all if the options vested at the same time.[54]

Executives may be forced to exercise or forfeit options if their employment is terminated. Leung and Sircar (2009) and Carpenter, Stanton, and Wallace (2010) consider exogenous job termination risk and find it reduces the exercise threshold and shareholder costs. We do not include job termination risk in our setting as it would not alter the main impacts of portfolio effects on exercise thresholds.[55] Importantly, its effect would be in the same direction as the impact of portfolio effects from unhedgeable risk and thus inclusion of employment risk would

[52]See Kole (1997) for a comprehensive analysis.

[53]Carpenter, Stanton, and Wallace (2010) incorporate vesting into their shareholder costs (but not their analysis of exercise thresholds) but do not compare to costs of an equivalent option without vesting.

[54]We anticipate that if the vesting times are different, there will be a small impact on some of the thresholds of options which have vested when other options remain unvested, but we do not believe this will alter our main findings or have a significant quantitative effect.

[55]Incorporation of exogenous job termination would result in a free boundary problem of reaction-diffusion type. Our transformation to the heat equation (see Appendix) would no longer apply and we would need to use alternative more computationally intensive methods, which would necessitate the development of new code.

reduce shareholder costs further.

Restricted stock is playing an increasingly important role in the overall compensation of executives, see Murphy (2012) and Conyon, Fernandes, Ferreira, Matos, and Murphy (2011) for recent statistics. Including restricted stock (in addition to the option portfolio) in our model would increase exposure to unhedgeable risk and hence further reduce exercise thresholds and shareholder costs of the options. Its inclusion would not alter the main findings of the paper. In single option grant models, Carpenter, Stanton, and Wallace (2010) show that their conclusions are robust to the addition of restricted stock, whilst Whalley (2008) contrasts the reduction in exercise thresholds when stock is optimally sold on exercise with the increased thresholds if stock is retained. Although we have studied portfolios of standard ESOs, many of our conclusions will be relevant also for performance-based compensation where there are additional performance-vesting or payout conditions based on stock price, accounting performance or industry comparisons (see Bettis, Bizjak, Coles, and Kalpathy (2010); Johnson and Tian (2000)).

As we discussed at the end of Section 2, in common with much of the related literature,[56] we assume CARA utility, however most notably Carpenter, Stanton and Wallace (2010) employ constant relative risk aversion (CRRA) to study wealth effects on single option exercise and shareholder costs. They find that the exercise threshold and shareholder cost differ most from their Black-Scholes counterparts when wealth is small. Our study of the effect of other options in the portfolio highlights that the greater the number of options, the larger the difference in threshold and the larger the cost discount relative to Black-Scholes. We do not consider the joint impact of wealth and portfolio effects for several reasons. It will make computations, even

[56] For instance, Leung and Sircar (2009), Chen, Miao and Wang (2010), and Miao and Wang (2007)

in a two-option portfolio, substantially more complex as an additional state variable will be needed. Despite this, the key observations of the paper concerning option portfolios will remain true. Finally, an executive's non-option wealth is difficult to identify whereas their share-based compensation is more readily available.

Finally, in our model the firm's stock price is unaffected by the executive. A number of models have allowed the executive to affect the future distribution of stock prices via effort or volatility choice under various utility assumptions and are able to investigate the optimal form of compensation.[57] Under some sets of assumptions cash and options are found to be a good approximation of the optimal contract (Dittman and Yu (2011); Armstrong, Larcker, and Su (2007), see also He (2012)) whilst others suggest cash and equity (Dittman and Maug (2007)) or a continuously rebalanced 'Dynamic Incentive Account' invested in cash and equity (Edmans, Gabaix, Sadzik, and Sannikov (2012)). Some models consider a single period optimization setting, whereas others involve continuous rebalancing of incentives. In practice firms do adjust executives' compensation packages, so the problem is multi- rather than single-period, but they are generally adjusted only on an annual basis, so incentives are not always optimal. We abstract from optimality considerations and consider a general portfolio of options, and document effects which need to be taken into account in any setting with multiple option grants, optimal or not. Additional modeling features will likely change the magnitude but not the existence of these portfolio effects: the position in the portfolio's *exercise order* will still have an important effect on an individual option's exercise threshold, and portfolios will still need to be valued *as a whole*, which magnifies the impact of risk premia on costs and incentives and implies that new

[57]For example, Dittman and Maug (2007), Armstrong, Larcker, and Su (2007), Dittman and Yu (2011), Dittman, Maug and Spalt (2010), Edmans, Gabaix, Sadzik, and Sannikov (2012) and He (2012).

option grants need to be valued at their *marginal* value.

6 Conclusion

In this paper we have modeled the behavior of a risk-averse executive with a portfolio of ESOs with various strikes and time to maturity. Since unhedgeable risk varies non-linearly with portfolio size and composition, the executive's exercise strategy and thus also the shareholder cost of an option held as part of a portfolio depend on the remainder of the executive's ESO portfolio holdings. These portfolio effects are both new to the literature and important - lowering both the moneyness required for exercise and the shareholder cost of most options in a portfolio. In fact, the proportional reduction in cost relative to stand-alone options, even in the case of only two options, can be over 40% for individual options and 20% for portfolios. In contrast to a risk-neutral setting, both exercise thresholds and costs depend on an option's position in the optimal exercise order: when the exercise order switches, the thresholds of options exercised earlier (later) jump down (up). Given these dependencies, the company should re-evaluate the costs of *all* outstanding ESOs each time they give a new grant of options to their executives, as executives may alter the moneyness and order at which they exercise their *existing* options when they receive a new grant. We use the model to explain several empirical findings in the literature: which options are attractive to exercise first and how exercise can be induced by a new grant. This highlights the importance of including portfolio effects in an empirical study of ESO exercise and accounting for such effects in any measure based on exercise behavior.

Whilst there is scope for further work on ESO portfolios,[58] this paper has set out the key principles of how portfolio effects impact ESO exercise and shareholder costs, which, as discussed in Section 5, we expect to be robust to our specific modeling assumptions. Firstly, for individual options in portfolios, the option's position in the portfolio's optimal exercise order is the new key factor determining both an option's exercise threshold and its cost within the portfolio. Secondly, for portfolios as a whole, what matters is the overall strength of the portfolio interaction effects. This determines the discount for the portfolio as a whole and the moneyness on exercise of the first option to be exercised in the portfolio. It depends on both the size and the overall composition of the portfolio, increasing when an option is added to a portfolio, and reaching a maximum when the strike prices of longer-dated options are slightly less than those of shorter-dated options within the portfolio. Finally, the relevant cost of an option grant is its *incremental* cost, which depends on the composition of the portfolio it is being combined with, and which is always lower than the option's stand-alone cost, since it reflects the *extra* unhedgeable risk associated with the option's *addition* to the executive's existing portfolio.

[58]For example, there are implications for incentives. ESO subjective deltas and vegas have been shown to affect executive's risk-taking incentives (See Armstrong and Vashishtha (2012) and references therein). The executive's subjective delta (and vega) of her overall portfolio of options will differ from the sum of the subjective deltas (vegas) of each individual option evaluated on a stand-alone basis. We leave detailed investigation of this for future work.

7 Appendix

7.1 The model

We first use separation of variables and a power transformation[59] via $V^{\Pi}(u,w,s) = \mathcal{M}(u,w,\tilde{T})H^{\Pi}(u,s)^{1/(1-\rho^2)}$ to restate (4), (5) and (9) as

$$H^{\Pi}(u,S_u)^{1/(1-\rho^2)} \leq \min_{\pi\in\Pi}\{e^{-\gamma(1-\rho^2)(S_u-K_\pi)^+e^{r(\tilde{T}-u)}}H^{\Pi\setminus\{\pi\}}(u,S_u)^{1/(1-\rho^2)}\} \tag{13}$$

$$\frac{\partial H^{\Pi}}{\partial t} + \tilde{\mathcal{L}}H^{\Pi} \geq 0 \tag{14}$$

where the differential operator $\tilde{\mathcal{L}}$ is defined by

$$\tilde{\mathcal{L}} = \frac{\eta^2 s^2}{2}\frac{\partial^2}{\partial s^2} + (r-q)s\frac{\partial}{\partial s} \tag{15}$$

and

$$\tau^{(|\Pi|)} = \inf\{t \leq u \leq T_{min} : H^{\Pi}(u,S_u) = \min_{\pi\in\Pi}\{e^{-\gamma(1-\rho^2)(S_u-K_\pi)^+e^{r(\tilde{T}-u)}}H^{\Pi\setminus\{\pi\}}(u,S_u)\}\}. \tag{16}$$

The boundary conditions are given by $H^{\Pi}(u,0) = 1$ and

$$\forall\{\pi\}\in\Pi;\quad \frac{H^{\Pi}(T_\pi,S_{T_\pi})}{H^{\Pi\setminus\{\pi\}}(T_\pi,S_{T_\pi})} = \frac{M(T_\pi,X_{T_\pi}+(S_{T_\pi}-K_\pi)^{\cdot}\tilde{T})}{M(T_\pi,X_{T_\pi},\tilde{T})} = e^{-\gamma(1-\rho^2)(S_{T_\pi}-K_\pi)^+e^{r(\tilde{T}-T_\pi)}}$$

and we have $H^{\emptyset}(u,S_u) = 1$.

We see in (16) that associated with each $H^{\Pi}(u,S_u)$ there is a free boundary

$$S^{(|\Pi|)}(u) = \inf\{s \geq 0 : H^{\Pi}(u,s) = \min_{\pi\in\Pi}\{e^{-\gamma(1-\rho^2)(s-K_\pi)^+e^{r(\tilde{T}-u)}}H^{\Pi\setminus\{\pi\}}(u,s)\}; u\in[t,T_{min}]\} \tag{17}$$

[59]The separation of variables is simply the observation that wealth factors out under exponential utility. We follow many authors who employ the power transformation in pricing of European and American options under exponential utility, for example Henderson (2005).

which represents the exercise boundary for the *next option* when the options Π remain. As with standard American options (see (26)), the optimal exercise times can be represented as

$$\tau^{(|\Pi|)} = \inf\{t \leq u \leq T_{min} : S_u = S^{(|\Pi|)}(u)\}. \tag{18}$$

Subjective value

We can use the definition of the subjective value to the executive in (12) to re-derive the free boundary problem as follows: $p^{\Pi}(t,s)$ solves

$$\frac{\partial p^{\Pi}}{\partial t} + \tilde{\mathcal{L}} p^{\Pi} - r p^{\Pi} - \frac{1}{2}\gamma(1-\rho^2)\eta^2 s^2 e^{r(\tilde{T}-t)} \left(\frac{\partial p^{\Pi}}{\partial s}\right)^2 \leq 0 \tag{19}$$

$$p^{\Pi}(t,s) \geq \max_{\pi\in\Pi}\{(s-K_\pi)^+ + p^{\Pi\setminus\{\pi\}}(t,s)\} \tag{20}$$

with boundary conditions $p^{\Pi}(t,0) = 0$ and

$$\forall\{\pi\} \in \Pi; \quad p^{\Pi}(T_\pi, S_{T_\pi}) - p^{\Pi\setminus\{\pi\}}(T_\pi, S_{T_\pi}) = (S_{T_\pi} - K_\pi)^+. \tag{21}$$

Shareholder costs

We can use the optimal exercise threshold for each option $S^\pi(u); \pi \in \Pi$ as input into the shareholder costs. To compute the shareholder cost of a portfolio Π, we compute the cost C^π for each option $\pi \in \Pi$ and sum over all options. Each C^π satisfies for $s \leq S^\pi(t)$

$$\frac{\partial C^\pi}{\partial t} + \tilde{\mathcal{L}} C^\pi = 0 \tag{22}$$

with boundary conditions $\forall\pi \in \Pi$, $C^\pi(u,0) = 0$, $C^\pi(T_\pi, S_{T_\pi}) = (S_{T_\pi} - K_\pi)^+$ and $C^\pi(u, S^\pi(u)) = (S^\pi(u) - K_\pi)^+; u < T_\pi$.

Black-Scholes

Consider a single call option Y with strike K_Y, maturity T_Y and assume we price under the

Black-Scholes model. The value $V^{BS}(u, S_u)$ of holding the American call solves the well known variational inequalities:

$$V^{BS}(u, S_u) \geq (S_u - K_Y)^+ \tag{23}$$

$$\frac{\partial V^{BS}}{\partial t} + \tilde{\mathcal{L}} V^{BS} - r V^{BS} \leq 0. \tag{24}$$

Boundary conditions are given by $V^{BS}(u, 0) = 0$ and $V^{BS}(T_Y, S_{T_Y}) = (S_{T_Y} - K_Y)^+$. The optimal exercise time τ^{BS} is defined by $\tau^{BS} = \inf\{t \leq u \leq T_Y : V^{BS}(u, S_u) = (S_u - K_Y)^+\}$ which defines an exercise boundary

$$S^{BS}(u) = \inf\{s \geq 0 : V^{BS}(u, s) = (s - K_Y)^+; u \in [t, T_Y]\} \tag{25}$$

and thus

$$\tau^{BS} = \inf\{t \leq u \leq T_Y : S_u = S^{BS}(u)\}. \tag{26}$$

7.2 Numerical method

Utility model

We first transform to work with the heat equation which is useful for comparison to Black-Scholes and to simplify the coding. Define new variables $x \in \mathbb{R}$ and $\tau \in [0, 0.5\eta^2 \tilde{T}]$ via $s = \delta e^x$ and $t = \tilde{T} - \tau/0.5\eta^2$, where δ is a constant. Write $H^{\Pi}(\tilde{T} - \tau/0.5\eta^2, \delta e^x) = \delta u^{\Pi}(x, \tau) e^{-\xi x - \xi^2 \tau}$ where $\xi = 0.5((r - q)/0.5\eta^2 - 1)$. Then (14) is given by

$$\partial u^{\Pi}/\partial \tau - \partial^2 u^{\Pi}/\partial x^2 \leq 0 \tag{27}$$

and (13) becomes

$$u^{\Pi}(x, \tau) \leq g^{\Pi}(x, \tau) \tag{28}$$

with

$$g^{\Pi}(x,\tau) = \min_{\pi\in\Pi}\{e^{-\gamma(1-\rho^2)(\delta e^x - K_\pi)^+ e^{r\tau/0.5\eta^2}} u^{\Pi\backslash\{\pi\}}(x,\tau)\}.$$

The boundary conditions are $u^{\Pi}(x_{min},\tau) = \frac{1}{\delta}e^{\delta x_{min}+\delta^2\tau}$ and

$$\forall\{\pi\}\in\Pi;\quad \frac{u^{\Pi}(x,\tau_\pi)}{u^{\Pi\backslash\{\pi\}}(x,\tau_\pi)} = e^{-\gamma(1-\rho^2)(\delta e^x - K_\pi)^+ e^{r\tau_\pi/0.5\eta^2}}$$

where $\tau_\pi = 0.5\eta^2(\tilde{T} - T_\pi)$.

To solve (27)-(28) subject to the boundary conditions, we use a Crank Nicolson finite difference method on a uniform grid. The free boundary constraint (28) is enforced by a projected successive over relaxation algorithm (PSOR), see Wilmott, Howison, and Dewynne (1995) for similar schemes in a Black-Scholes framework. We restrict the domain $\mathbb{R}\times[0, 0.5\eta^2\tilde{T}]$ to a finite domain $\{(x,\tau) : x_{min} \leq x \leq x_{max}, 0 \leq \tau \leq 0.5\eta^2\tilde{T}\}$. We introduce a uniform grid with nodes $\{(x_{min} + j\Delta x, n\Delta\tau) : j = 0, 1, ..., j_{max}, n = 0, 1, ..., n_{max}\}$ with grid spacings $\Delta x = (x_{max} - x_{min})/j_{max}$, $\Delta\tau = 0.5\eta^2\tilde{T}/n_{max}$. We apply discrete approximations $U_j^n \approx u^A(x_{min} + j\Delta x, n\Delta\tau)$ and approximate the derivatives by:

$$\frac{\partial u^{\Pi}(x,\tau)}{\partial\tau} \approx \frac{U_j^{n+1} - U_j^n}{\Delta\tau}$$

$$\frac{\partial^2 u^{\Pi}(x,\tau)}{\partial x^2} \approx \frac{1}{2}\left(\frac{U_{j+1}^n - 2U_j^n + U_{j-1}^n}{(\Delta x)^2} + \frac{U_{j+1}^n - 2U_j^n + U_{j-1}^n}{(\Delta x)^2}\right)$$

giving

$$U_j^{n+1} - \frac{1}{2}a(U_{j-1}^{n+1} - 2U_j^{n+1} + U_{j+1}^{n+1}) = U_j^n + \frac{1}{2}a(U_{j-1}^n - 2U_j^n + U_{j+1}^n) \tag{29}$$

where $a = \Delta\tau/(\Delta x)^2$. We solve (29) (together with the discretized boundary conditions) backward in time using the PSOR algorithm to compute the optimal boundary $s^* = \delta e^{x_{min}+j^*\Delta x}$ at each time step $n\Delta\tau$ by finding the minimum index $j = j^*$ such that $U_{j^*}^n = g^{\Pi}(x_{min} + j^*\Delta x, n\Delta\tau)$.

Shareholder costs

Similar transformations and approximations are used to solve (22) together with the given boundary conditions. (In this case there is no free boundary, so we do not require the PSOR.)

Black-Scholes

Each option in the portfolio can be treated separately under the Black-Scholes model since prices are linear. It's price solves (23) and (24) together with the given boundary conditions. We follow the same transformation to the heat equation and numerical finite difference scheme as outlined earlier.

Robustness

We perform various robustness checks for the utility model algorithm. We experimented with the size of the grid versus speed and use $j_{max} = 2000, n_{max} = 9000, x_{min} = -3$ and $x_{max} = 3$. Using the utility model algorithm gives thresholds and values that converge to the correct Black-Scholes quantities (calculated from (23) and (24)) as $\gamma \to 0$ or $\rho \to 1$. For large maturities and identical strikes, we can recover the time-homogeneous thresholds and values derived explicitly by the perpetual approximation of Grasselli and Henderson (2009).

References

Armstrong, C.S., Jagolinzer, A.D., Larcker, D.F., 2007. Timing of employee stock option exercises and the valuation of stock option expense. Rock Center for Corporate Governance Working Paper No. 34.

Armstrong, C.S., Larcker, D.F., Su, C., 2007. Stock options and Chief Executive Officer compensation. Working Paper, University of Pennsylvania.

Armstrong, C.S., Vashishtha, R., 2012. Executive stock options, differential risk-taking incentives, and firm value. Journal of Financial Economics, 104, 70-88.

Bergstresser, D.B., Philippon, T., 2006. CEO incentives and earnings management. Journal of Financial Economics, 80, 511-529.

Bettis, J.C., Bizjak, J.M., Coles, J., Kalpathy, S., 2010. Stock and option grants with performance-based vesting provisions. Review of Financial Studies, 23, 10, 3849-3888.

Bettis, J.C., Bizjak, J.M., Kalpathy, S., 2011. Why do insiders hedge their ownership? An empirical examination. Working paper.

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, 445-470.

Cai, J., Vijh, A.M., 2005. Executive stock option valuation in a two state variable framework. Journal of Derivatives, Spring, 9-27.

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

Carpenter, J.N., Stanton, R., Wallace, N., 2010. Optimal exercise of executive stock options and implications for firm cost. Journal of Financial Economics, 98, 315-337.

Carpenter, J.N., Stanton, R., Wallace, N., 2012. Estimation of employee stock option exercise rates and firm cost. Working Paper, UC Berkeley.

Chen H., Miao, J., Wang, N., 2010. Entrepreneurial finance and nondiversifiable risk, Review of Financial Studies, 23, 12, 4348-4388.

Constantinides, G., 1984. Warrant exercise and bond conversion in competitive markets. Journal of Financial Economics, 13, 371-397.

Conyon, M.J., Core, J.E., Guay, W.R., 2011. Are U.S. CEOs paid more than U.K. CEOs? Inferences from risk-adjusted pay. Review of Financial Studies, 24, 2, 402-438.

Conyon, M.J., Fernandes, N., Ferreira, M.A., Matos, P., Murphy, K.J., 2011. The executive compensation controversy: a transatlantic analysis. Fondazione Rodolfo De Benedetti Report.

Cox, J.C., Rubinstein, M., 1985. Options Markets. Prentice Hall Inc., Englewood Cliffs, NJ.

Dennis, P.J., Rendleman, R.J., 2008. Valuing Multiple Employee Stock options issued by the same company. Journal of Derivatives, 16(1), 44-69.

Detemple, J., Sundaresan, S., 1999. Nontraded asset valuation with portfolio constraints: a binomial approach. Review of Financial Studies, 12, 835-872.

Dittmann, I., Maug, E., 2007. Lower salaries and no options? On the optimal structure of executive pay. Journal of Finance, 62, 303-343.

Dittmann, I., Maug, E., Spalt, O., 2010. Sticks or carrots? Optimal CEO compensation when managers are loss averse. Journal of Finance, 65, 2015-2050.

Dittmann, I., Yu, K., 2011. How important are risk-taking incentives in executive compensation? Working Paper, Erasmus School of Economics.

Edmans, A., Gabaix, X., Sadzik, T., Sannikov, Y., 2012. Dynamic CEO compensation. Journal of Finance, 67, 1603-1647.

Fernandes, N., Ferreira, M.A., Matos, P., Murphy, K.J., 2013. Are U.S. CEOs paid more? New international evidence. Review of Financial Studies, 26, 2, 323-367.

Frydman, C., Saks, R.E., 2010. Executive compensation: a new view from a long term perspective, 1936-2005. Review of Financial Studies, 23, 5, 2100-2138.

Grasselli, M., Henderson, V., 2009. Risk aversion and block exercise of executive stock options. Journal of Economic Dynamics and Control, 33, 109-127.

Hall, B.J., Murphy, K.J., 2000. Optimal exercise prices for executive stock options. American Economic Review Papers and Proceedings, 90(2), 209-214.

Hall, B.J., Murphy, K.J., 2002. Stock options for undiversified employees. Journal of Accounting and Economics, 33, 3-42.

He, Z., 2012. Dynamic compensation contracts with private savings. Review of Financial Studies, 25, 5, 1494-1549.

Heath, C., Huddart, S., Lang, M., 1999. Psychological factors and stock option exercise. Quarterly Journal of Economics, 114(2), 601-627.

Henderson, V., 2005. The impact of the market portfolio on the valuation, incentives and optimality of executive stock options. Quantitative Finance, 5(1), 35-47.

Henderson, V., Hobson, D., 2011. Optimal liquidation of derivative portfolios. Mathematical Finance, 21, 3, 365-382.

Henderson ,V., Sun, J., Whalley, A.E., 2013. Portfolios of american options under general preferences: results and counterexamples. Mathematical Finance, forthcoming.

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

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

Ingersoll, J.E., 2006. The subjective and objective valuation of incentive stock options. Journal of Business, 79, 453-487.

Jain, A., Subramanian, A., 2004. The intertemporal exercise and valuation of employee stock options. The Accounting Review, 79, 3, 705-743.

Johnson, S.A., Tian, Y.S., 2000. The value and incentive effects of nontraditional executive stock option plans. Journal of Financial Economics, 57, 3-34.

Kahl, M., Liu, J., Longstaff, F.A., 2003. Paper millionaires: how valuable is stock to a stockholder who is restricted from selling it? Journal of Financial Economics, 67, 385-410.

Klein, D., Maug, E., 2010. How do executives exercise their stock options? ECGI Working paper.

Kim, I.J., 1990. The Analytic Valuation of American Options. Review of Financial Studies, 3, 547-572.

Kole, S.R., 1997. The complexity of compensation contracts. Journal of Financial Economics, 43, 79-104.

Lambert, R., Larcker, D.F., Verrechia, R.E., 1991. Portfolio considerations in valuing executive compensation. Journal of Accounting Research, 29, 129-149.

Leung, T., Sircar, R., 2009. Accounting for risk aversion, vesting, job termination risk and multiple exercises in valuation of employee stock options. Mathematical Finance, 19, 99-128.

Malemendier, U., Tate, G., 2005. CEO overconfidence and corporate Investment. Journal of Finance, 60, 2661-2700.

Merton, R.C., 1971. Optimum consumption and portfolio rules in a continuous time model. Journal of Economic Theory, 3, 373-413.

Miao J., N. Wang, N., 2007. Investment, consumption and hedging under incomplete markets. Journal of Financial Economics, 86, 608-642.

Murphy, K.J., 2012. Executive compensation: Where we are and how we got there. In Constantinides, G., Harris, M., Stulz, R. (eds.), Handbook of the Economics of Finance, Elsevier Science North Holland.

Ofek, E., Yermack, D., 2000. Taking stock: equity-based compensation and the evolution of managerial ownership. Journal of Finance, 55, 1367-1384.

Rogers, L.C.G., Scheinkman, J., 2007. Optimal exercise of executive stock options. Finance and Stochastics, 11, 357-372.

Spatt, C.S., Sterbenz, F.P., 1988. Warrant exercise, dividends and reinvestment policy. Journal of Finance, 43, 2, 493-506.

Srivastava, A., 2011. The early exercise of executive stock options and firm's future systematic and idiosyncratic risks. Working paper, Kellogg School of Management, Northwestern University.

Sun, J., 2011. Models of executive stock options, PhD Thesis, Warwick Business School.

Van Bekkum, S. and Zhu, G., 2013. Insider Information or CEO Talent? An Analysis of CEO Exercises. Working paper.

Whalley, A.E., 2008. Effect of executive share ownership and private hedging on executive stock option exercise and values. Working paper, Warwick Business School.

Wilmott, P., Howison, S.D., and Dewynne J.N., 1995. The Mathematics of Financial Derivatives. Cambridge University Press.

		$C^Y_{\mathcal{S}}/C^Y_{BS}$	C^Y/C^Y_{BS}	C^Z/C^Z_{BS}	$(C^Y_{\mathcal{S}}+C^Z_{\mathcal{S}})/(C^Y_{BS}+C^Z_{BS})$	$C^{Y,Z}/(C^Y_{BS}+C^Z_{BS})$
γ	0	1	1	1	1	1
	0.01	0.97	0.89	0.96	0.96	0.93
	0.1	0.68	0.45	0.65	0.66	0.56
	0.2	0.53	0.30	0.49	0.51	0.40
η_e	0.2	0.68	0.43	0.66	0.67	0.56
	0.4	0.53	0.30	0.49	0.51	0.40
	0.6	0.43	0.23	0.40	0.41	0.32
$r-q$	0	0.61	0.36	0.58	0.59	0.48
	0.03	0.56	0.32	0.53	0.55	0.44
	0.05	0.53	0.30	0.49	0.51	0.40

Table 1: The ratio of option cost as a proportion of the corresponding Black-Scholes world cost for (i) Y as a stand-alone, (ii) Y in the portfolio with Z, (iii) Z in a portfolio with Y ($= Z$ as stand-alone), (iv) the sum of Y and Z on a stand-alone basis and (v) the portfolio of Y and Z. Option Y has $K_Y = 10$, $T_Y = 5$, option Z has $K_Z = 10$ and $T_Z = 10$. Parameters unless otherwise stated are: $S = 10$, $r = 0.10$, $r - q = 0.05$, $\eta = 0.4$, $\rho = 0$, $\gamma = 0.2$.

Panel A: $\gamma = 0.2$, $\eta = 0.4$, $\rho = 0$, $r = 0.1$, $r - q = 0.05$

T_Z	$K_Z = 8$			$K_Z = 9$			$K_Z = 10$			$K_Z = 11$		
	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$
6	1.000	*0.862*	**0.923**	1.000	*0.631*	**0.804**	0.557	*1.000*	**0.782**	0.620	*1.000*	**0.810**
10	1.000	*0.822*	**0.899**	0.997	*0.632*	**0.797**	0.565	*1.000*	**0.795**	0.620	*1.000*	**0.820**
15	1.000	*0.798*	**0.884**	0.994	*0.636*	**0.795**	0.571	*1.000*	**0.803**	0.630	*1.000*	**0.820**

Panel B: $\gamma = 0.1$, $\eta = 0.4$, $\rho = 0$, $r = 0.1$, $r - q = 0.05$

T_Z	$K_Z = 8$			$K_Z = 9$			$K_Z = 10$			$K_Z = 11$		
	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$
6	1.000	*0.733*	**0.851**	0.999	*0.689*	**0.884**	0.640	*1.000*	**0.830**	0.690	*1.000*	**0.840**
10	0.999	*0.732*	**0.847**	0.987	*0.695*	**0.826**	0.650	*1.000*	**0.840**	0.690	*1.000*	**0.850**
15	0.997	*0.733*	**0.845**	0.976	*0.704*	**0.823**	0.660	*1.000*	**0.850**	0.700	*1.000*	**0.860**

Panel C: $\gamma = 0.2$, $\eta = 0.6$, $\rho = 0$, $r = 0.1$, $r - q = 0.05$

T_Z	$K_Z = 8$			$K_Z = 9$			$K_Z = 10$			$K_Z = 11$		
	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$
6	1.000	*0.834*	**0.910**	1.000	*0.602*	**0.792**	0.527	*1.000*	**0.766**	0.593	*1.000*	**0.791**
10	1.000	*0.809*	**0.895**	0.999	*0.605*	**0.789**	0.535	*1.000*	**0.776**	0.596	*1.000*	**0.799**
15	1.000	*0.799*	**0.889**	0.998	*0.607*	**0.788**	0.536	*1.000*	**0.779**	0.596	*1.000*	**0.802**

Panel D: $\gamma = 0.2$, $\eta = 0.4$, $\rho = 0$, $r = 0.05$, $r - q = 0.05$

T_Z	$K_Z = 8$			$K_Z = 9$			$K_Z = 10$			$K_Z = 11$		
	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$
6	1.000	*0.710*	**0.840**	0.998	*0.669*	**0.823**	0.630	*1.000*	**0.810**	0.666	*1.000*	**0.831**
10	0.998	*0.711*	**0.833**	0.980	*0.678*	**0.811**	0.640	*1.000*	**0.830**	0.671	*1.000*	**0.845**
15	0.995	*0.711*	**0.829**	0.948	*0.696*	**0.803**	0.640	*1.000*	**0.840**	0.675	*1.000*	**0.855**

Panel E: $\gamma = 0.2$, $\eta = 0.4$, $\rho = 0$, $r = 0.1$, $r - q = 0.03$

T_Z	$K_Z = 8$			$K_Z = 9$			$K_Z = 10$			$K_Z = 11$		
	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$	$\mathcal{R}_Y$	$\mathcal{R}_Z$	$\mathcal{R}_{Y,Z}$
6	1.000	*0.900*	**0.944**	1.000	*0.640*	**0.809**	0.570	*1.000*	**0.790**	0.640	*1.000*	**0.810**
10	1.000	*0.866*	**0.923**	0.998	*0.648*	**0.806**	0.580	*1.000*	**0.800**	0.640	*1.000*	**0.820**
15	1.000	*0.848*	**0.912**	0.997	*0.654*	**0.806**	0.580	*1.000*	**0.810**	0.640	*1.000*	**0.830**

Table 2: The ratio of the option cost calculated on a portfolio basis to the (sum of the) stand-alone cost(s): $\mathcal{R}_Y = C^Y/C^Y_{\mathcal{S}}$, $\mathcal{R}_Z = C^Z/C^Z_{\mathcal{S}}$ in italics and $\mathcal{R}_{Y,Z} = C^{Y,Z}/(C^Y_{\mathcal{S}} + C^Z_{\mathcal{S}})$ in bold. Option Y has $K_Y = 10$, $T_Y = 5$. The Table varies K_Z and T_Z. Parameters unless otherwise stated are: $S = 10$, $r = 0.10$, $r - q = 0.05$, $\eta = 0.4$, $\rho = 0$, $\gamma = 0.2$.

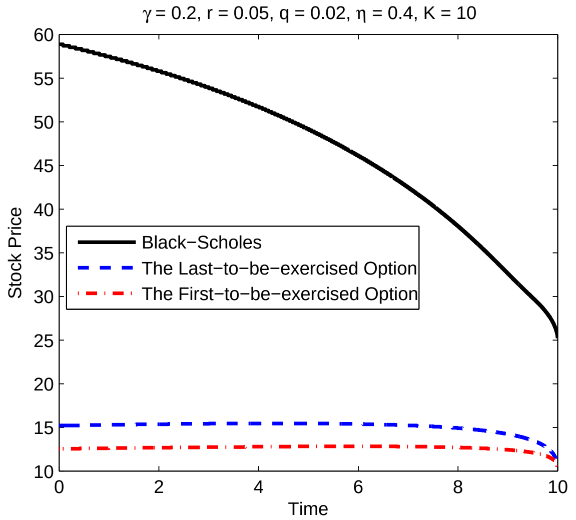

Figure 1: *Exercise thresholds vs time for portfolio of two identical options with* $K_Y = 10, T_Y = 10$. *Parameter values:* $\gamma = 0.2$, $r = 0.05$, $q = 0.02$, $\eta = 0.4$, $\rho = 0$.

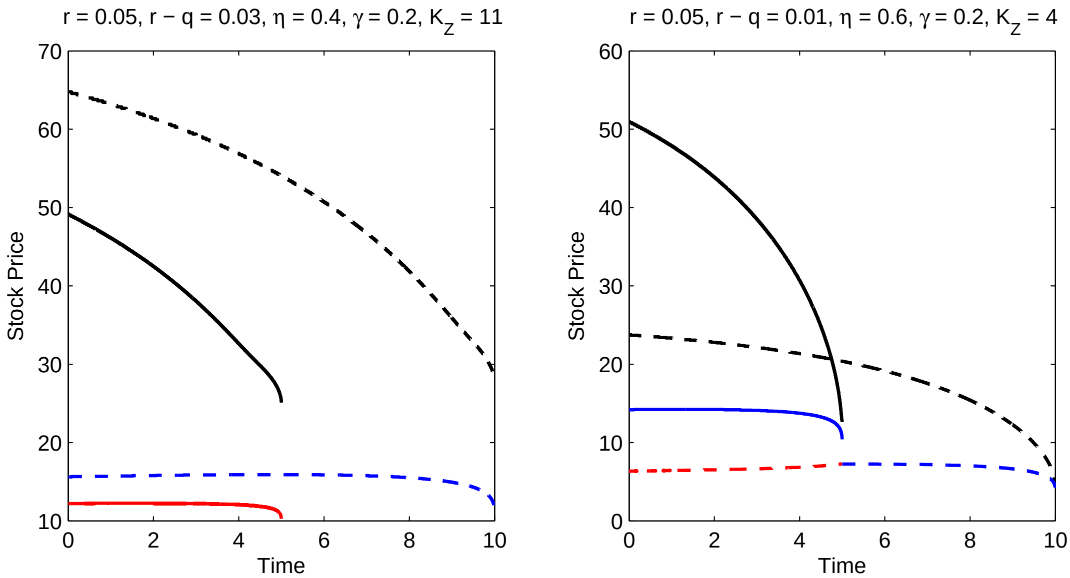

Figure 2: *Black-Scholes and optimal exercise thresholds vs time. Option* Y *has* $K_Y = 10, T_Y = 5$*; option* Z *has* $T_Z = 10$. *In the left panel,* $K_Z = 11$*; in the right panel* $K_Z = 4$. *Parameter values where not stated:* $\rho = 0$.

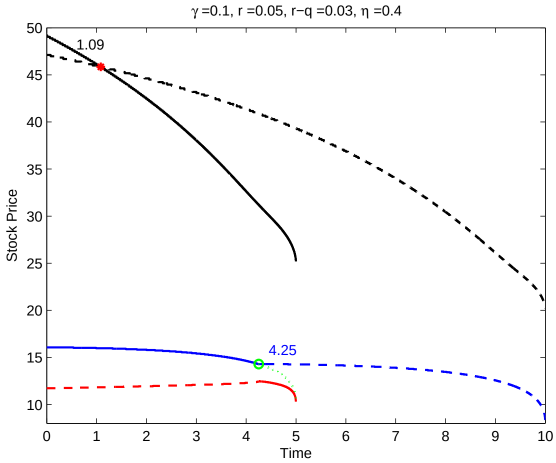

Figure 3: *Black-Scholes and optimal exercise thresholds vs time. Option Y has* $K_Y = 10, T_Y = 5$*; option Z has* $K_Z = 8, T_Z = 10$*. Parameter values:* $\gamma = 0.1$*,* $r = 0.05$*,* $q = 0.02$*,* $\eta = 0.4$*,* $\rho = 0$

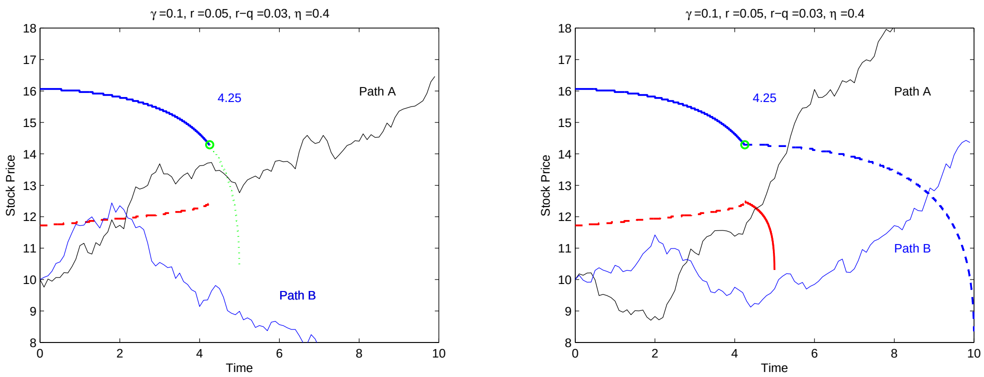

(a) (b)

Figure 4: *Stock price scenarios and exercise thresholds vs time. Option Y has* $K_Y = 10, T_Y = 5$*; option Z has* $K_Z = 8, T_Y = 10$*. Parameter values:* $\gamma = 0.1$*,* $r = 0.05$*,* $q = 0.02$*,* $\eta = 0.4$*,* $\rho = 0$*.*

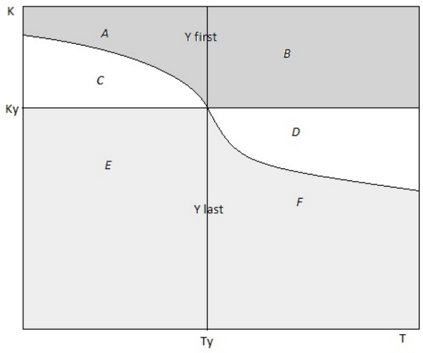

Figure 5: *Stylized representation of strike/maturity combinations giving various exercise ordering strategies for base option* Y *with strike* K_Y*, maturity* T_Y *and a second option with a varying strike* K *and maturity* T.

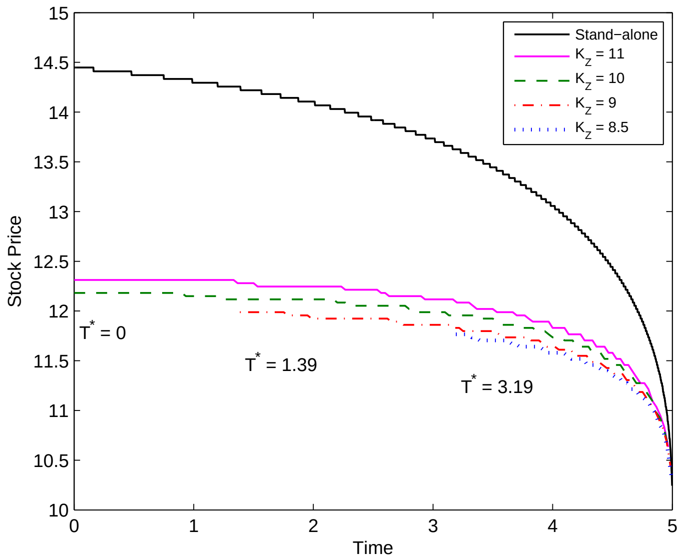

Figure 6: *Impact of* K_Z *on thresholds of option* Y $(K_Y = 10, T_Y = 5,\ T_Z = 10)$*. Parameter values:* $\gamma = 0.2,\ r = 0.05,\ q = 0,\ \eta = 0.2,\ \rho = 0.$

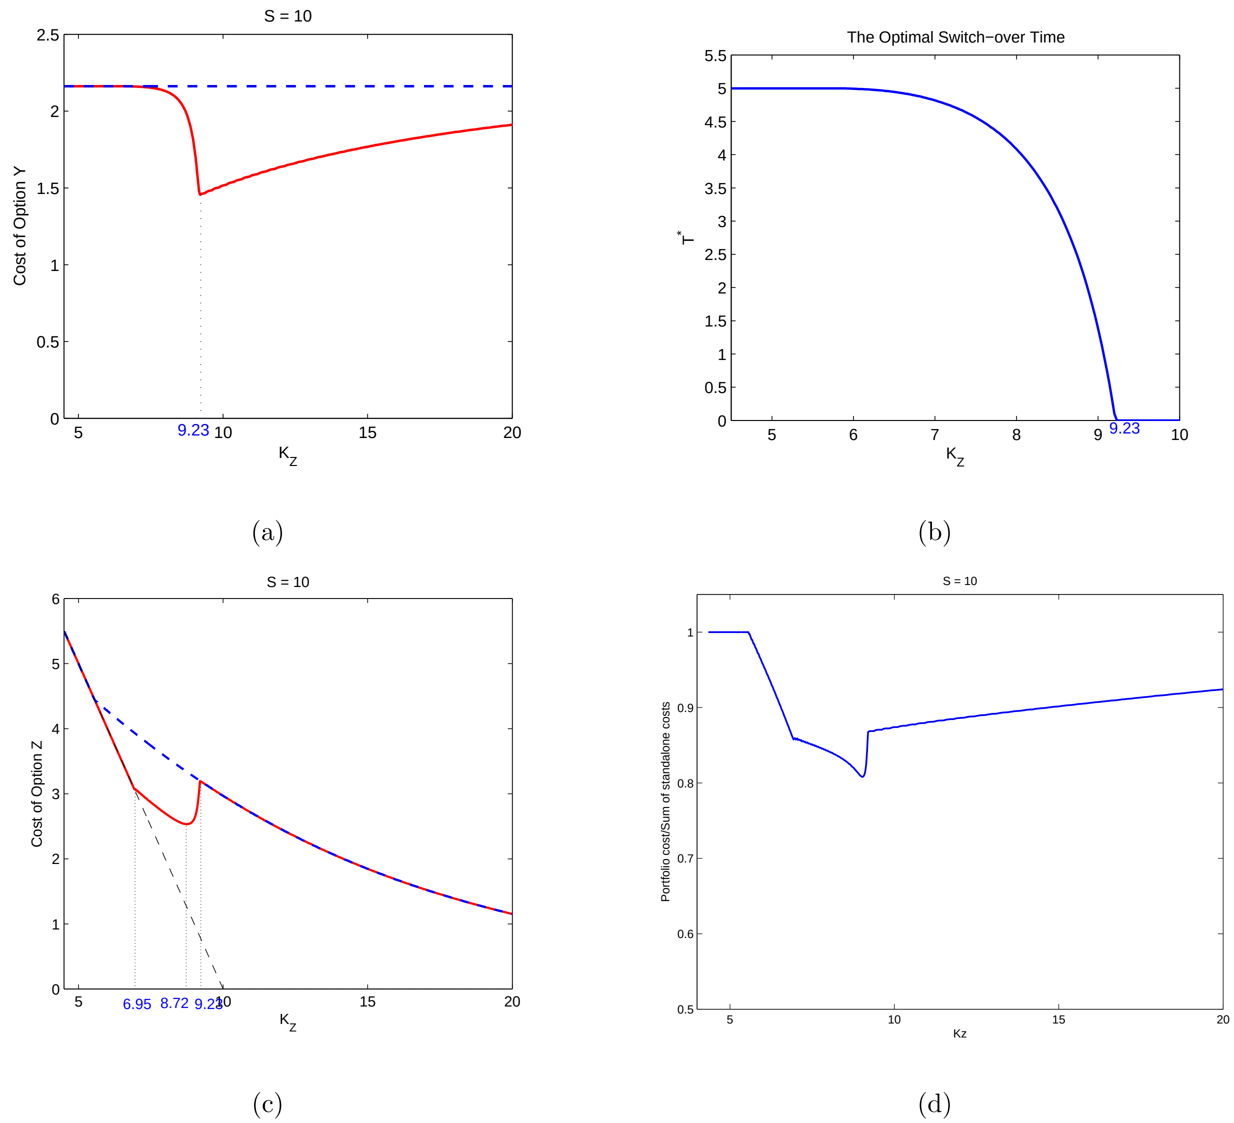

Figure 7: *Top left graph: Cost of option Y, Top right graph: Switchover time, T^*, Bottom left graph: Cost of option Z, Bottom right graph: Ratio of cost of portfolio to sum of stand-alone costs; all vs strike price of option Z. Option Y has $K_Y = 10, T_Y = 5$; option Z has $T_Z = 10$. Parameter values: $\gamma = 0.2$, $r = 0.05$, $q = 0$, $\eta = 0.2$, $\rho = 0$.*

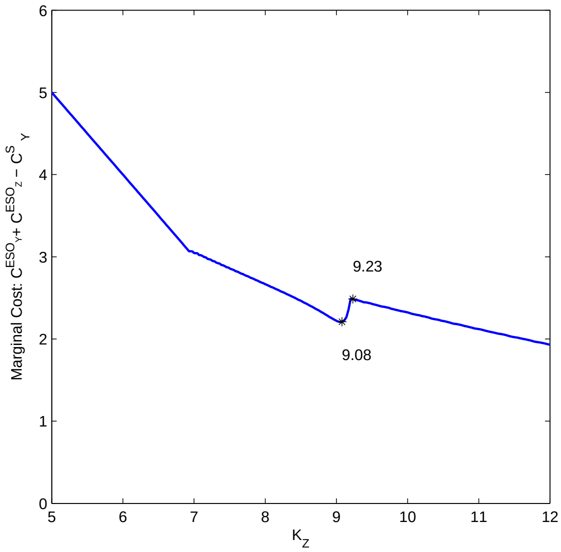

Figure 8: *Incremental cost of new grant Z in portfolio with existing ATM option Y, as a function of K_Z. Option Y has $K_Y = 10$, $T_Y = 5$. Parameter values: $\gamma = 0.2$, $r = 0.05$, $q = 0$, $\eta = 0.2$, $\rho = 0$.*

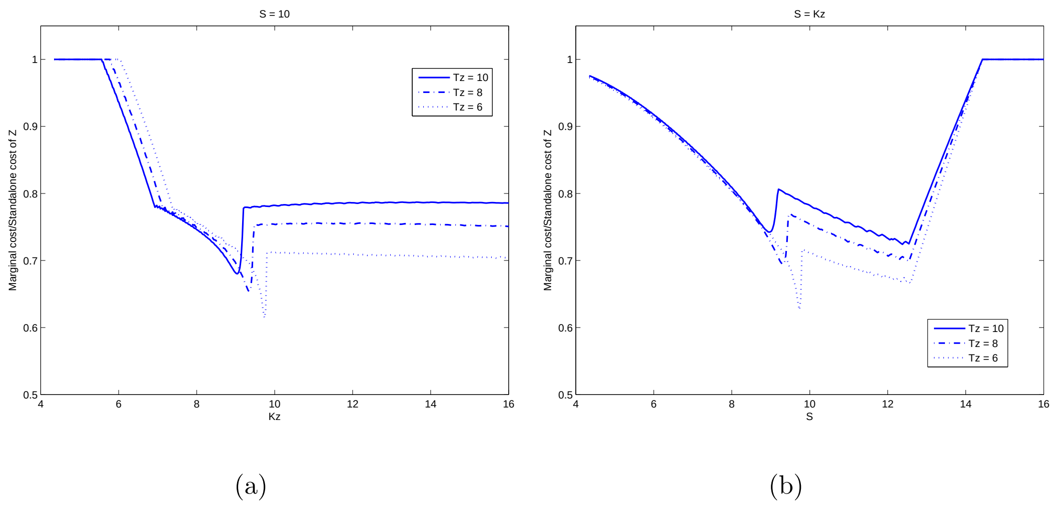

Figure 9: *Ratio of incremental cost of new grant Z in portfolio with existing option Y, as a proportion of stand-alone cost of Z for different T_Zs. Left graph vs K_Z with $S = 10$, right graph vs S for ATM option grant of Z. Option Y has $K_Y = 10$, $T_Y = 5$. Parameter values: $\gamma = 0.2$, $r = 0.05$, $q = 0$, $\eta = 0.2$, $\rho = 0$.*

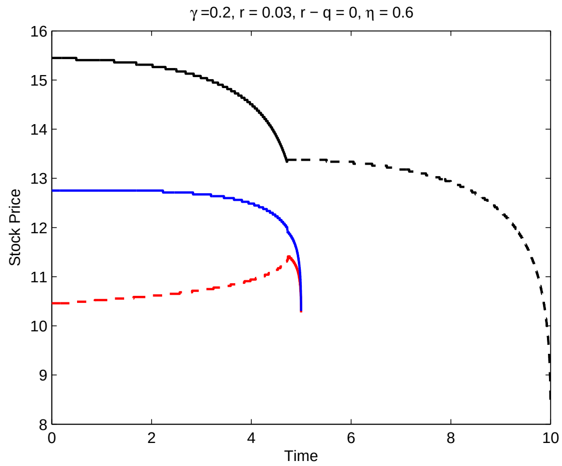

Figure 10: *Exercise thresholds vs time for portfolio with three options: X, Y and Z. We have $K_X = K_Y = 10$, $T_X = T_Y = 5$ and $K_Z = 8$, $T_Z = 10$. Parameter values: $\gamma = 0.2$, $r = 0.03$, $q = 0.03$, $\eta = 0.6$, $\rho = 0$.*