Victoria

UNIVERSITY OF WELLINGTON

*Te Whare Wānanga
o te Ūpoko o te Ika a Māui*

EFFORT, IDIOSYNCRATIC RISK AND INVESTMENT UNDER UNCERTAINTY

by

WILLIAM S. TAYLOR

A Thesis

submitted to the Victoria University of Wellington

in fulfilment of the

requirements for the degree of

Doctor of Philosophy

in Economics.

Victoria University of Wellington

2011

Abstract

This thesis is based upon four very simple premises:

1. managers, not shareholders make the investment decisions for the firm;
2. managers do more than just say “yes” or “no” to investments, they can also exert effort that affects the payoff from investment;
3. executive compensation schemes can cause managers to hold more stock than is optimal for diversification purposes; and
4. many investments can be delayed and involve irreversible capital costs as well as uncertain payoffs.

Combining these four premises gives the two central questions this thesis attempts to answer:

1. How does the level of managerial stock-ownership affect the investment decisions managers make for the firm? and
2. given the answer to (1), how does this affect the shareholder’s decision to hire a manager?

In this thesis I use a continuous time “Real Options” framework to answer these questions. The form of the utility function assumed for the manager has a huge impact on the tractability of the modelling. The assumption of Constant Relative Risk Aversion (CRRA) utility as opposed to Constant Absolute Risk Aversion (CARA) causes the manager’s valuation of the cash

flow (the very first step of the modelling) to become wealth dependent. This in itself is an interesting issue, but it also poses interesting numerical issues and makes the later steps of the analysis intractable. Because of this we split the substantive analysis of this thesis into two parts. In the first we assume CARA utility in order to remove wealth dependence from the valuation and obtain a “clean path” to the end goal of a dynamic model of hiring, effort and irreversible investment. In the second we focus on CRRA utility thus allowing the manager’s valuation to depend on his financial wealth. We then explain the resultant numerical issues, and the appropriate approach to their solution.

*“It’s not a problem, it’s a
challenge.”*

- Guthrie (2006)

Acknowledgments

This thesis is dedicated my Granddaddy Archie, Great Uncle Bill and Abuelo¹ José. They have always supported me in everything I have done and never stopped telling me how proud they are of me.

Completing this thesis has been an arduous journey with many ups and many downs. There are many people that I wish thank, which I do so now in no particular order.

I have benefited from the having two outstanding supervisors in Professors Lew Evans and Graeme Guthrie. Lew has always made me step back and look at the big picture, while Graeme has a technical knowledge that I don't think I could find anywhere else in the world. I also have Graeme to thank for the quote on the preface. When discussing a modelling issue I was having, he made the off hand quote that "it's not a problem, it's a challenge". This became the mantra for the rest of my thesis and got me through many low points of my modelling. I doubt he realised at the time, but that off hand comment is probably the reason I have now finished this thesis. So I owe Graeme a huge debt of gratitude for his eloquent (and unintendedly profound!) choice of words.

The New Zealand Institute for Study of Competition and Regulation (ISCR)

¹This is the Spanish word for grandfather for those who don't know!

provided me with a generous scholarship and office space for the first two years of my study before I moved from Wellington to Auckland. Similarly, the School of Economics and Finance at Victoria University has generously provided me with office space when I have needed it in Wellington.

I also need to thank my friends for putting up with my relative solitude during the 2 years I worked full time while also working on my thesis (a horrible idea in hind sight!).

My boss at NERA, James Mellsop respected that weekends were thesis time and thus didn't ask me to work on the weekend when it was avoidable (and still doesn't thankfully!).

Contents

List of Figures	xiii
I Motivation and Framework	1
1 Introduction	2
2 Literature Review	16
2.1 Introduction	16
2.2 Standard Theory of Investment	18
2.3 Real Options Analysis (ROA)	20
2.3.1 Overview	20
2.3.2 The Option to Delay Investment	22
2.3.3 Criticisms: Is There Really Value in Waiting?	24
2.4 Non-Traded Assets and Idiosyncratic Risk	29
2.4.1 General literature on non-traded options/stocks	30

2.4.2	Valuation of non-traded cashflows	31
2.4.3	Real Options	34
2.4.4	Implications for Well Diversified and Risk Neutral Share- holders	41
2.5	Effort	41
3	General Setup	44
3.1	Introduction	44
3.2	The Principal-Agent framework	46
3.3	The Shareholder's problem	47
3.4	The Manager's problem	48
3.4.1	General setup	48
3.4.2	The Manager's Utility Function	49
3.5	Stochastic Structure	51
3.5.1	Market Asset	51
3.5.2	Investment State Variable	51
3.6	Inter-temporal Wealth Equations	53
3.6.1	Stage 2 : Post Investment	53
3.6.2	Stage 1 : Pre Investment	53
3.7	The General Hamilton-Jacobi-Bellman Equation	54
3.8	Special Case: The Model with No Project (Merton (1969)) . .	54
3.8.1	The Manager's Problem	55
3.8.2	Solutions	55

II	Clean path	57
4	CARA Utility with SBM Cash flow	58
4.1	Introduction	58
4.2	Utility Valuation of the Cashflow	60
4.2.1	Solution	60
4.2.2	Comparative statics	62
4.3	The Static Investment Problem	67
4.4	Summary	70
5	Managerial Effort	71
5.1	Introduction	71
5.2	Setup	73
5.3	The Manager's Optimal Level of Effort	75
5.4	The Static Investment/Effort Decision	83
5.5	Summary	86
6	The Shareholder's Static Hiring Problem	88
6.1	Introduction	88
6.2	Setup	91
6.3	Results	92
6.3.1	Managerial Skill/Scope for Cost Savings (B)	94
6.3.2	Managerial Skill (λ)	94
6.3.3	Risk Aversion (γ)	96

6.3.4	Managerial Wealth (W)	99
6.3.5	Level of Cash Flow (Y)	100
6.3.6	Cash Flow Volatility (σ_y)	101
6.4	Conclusion	105
7	Effort and the Timing Option	107
7.1	Introduction	107
7.2	Setup	109
7.3	Numerical Solution Method	110
7.3.1	Description of algorithm	110
7.3.2	Boundary conditions	112
7.4	Results	114
7.5	Summary	131
8	The Shareholder's Dynamic Problem	134
8.1	Introduction	134
8.2	Framework	136
8.3	Payoff from hiring a manager	138
8.3.1	Setup	138
8.3.2	Solution	140
8.4	The Static Hiring Decision	145
8.4.1	Setup	145
8.4.2	Solution	145

8.4.3	Comparative statics	150
8.5	The Dynamic Hiring Decision	158
8.5.1	Setup	158
8.5.2	Solution	159
8.5.3	Comparative statics	162
8.6	Impact on shareholder wealth	165
8.7	Conclusion	173

III Direct wealth effects 176

9 CRRA Utility with GBM Cash Flow 177

9.1	Introduction	177
9.2	The Model	181
9.3	Measures of Value	183
9.4	Limiting Behavior of the Value Function	186
9.5	Numerical Solution Method	188
9.5.1	Description of Algorithm	188
9.5.2	Finite Difference Approximations	189
9.5.3	Boundary Conditions	190
9.6	Numerical Results	193
9.6.1	Review of Munk's (2000) Numerical Results	193
9.6.2	Our Results	196
9.7	Manager's subjective discount rate	206

9.8	Adding Effort to the CRRA Model	208
9.8.1	Solution method	208
9.8.2	Results	209
9.9	Summary	211
IV	Conclusion and Appendices	214
10	Summary and Conclusion	215
A	ROA Valuation Using the CAPM	220
A.1	Setup	220
A.2	Numerical Solution	221
B	Additional Comparative statics: Chapter 5	224
B.0.1	Investment Cost Ceiling (A)	224
B.0.2	Aversion to effort (θ)	225
B.0.3	Correlation Co-efficient (ρ)	226
B.0.4	Cash Flow Growth Rate (μ)	227
C	Additional Comparative Statics: Chapter 9	230
	Bibliography	235

List of Figures

1.1	Executive compensation over time	3
1.2	Compensation Mix	4
1.3	Incentive Mix	4
1.4	Thesis structure	15
4.1	Subjective Value as a function of α and ρ	63
4.2	Subjective Value as a function of α and σ_y	65
4.3	Subjective Value as a function of α and γ	65
4.4	Subjective Value as a function of ρ and σ_y	66
4.5	Investment Threshold (Y^*) as Function of α	68
4.6	Investment Threshold (Y^*) as Function of W	69
5.1	Investment Cost vs Effort for various levels of λ	74
5.2	The effect of W and Y on optimal effort	77
5.3	The effect of α and γ on optimal effort	78

5.4	The effect of α and γ on optimal effort when $\rho = 1$	80
5.5	The effect of θ and λ on optimal effort	81
5.6	The effect of A and B on optimal effort	81
5.7	The effect of σ_y and ρ on optimal effort	82
5.8	The effect of σ_m on optimal effort	83
5.9	Investment Threshold (Y^*) as a Function of α and γ	84
5.10	Investment Threshold (Y^*) as a Function of W and γ	85
6.1	Calculating α^*	93
6.2	Comparative Statics For B	95
6.3	Comparative Statics For λ	96
6.4	Comparative Statics For γ	97
6.5	The effect of α and γ on optimal effort	98
6.6	Comparative Statics For W	99
6.7	Comparative Statics For Y	100
6.8	Comparative Statics For σ_y , $\rho = 0$	102
6.9	Comparative Statics For σ_y , $\rho = 0.5$	103
6.10	Comparative Statics For σ_y , $\rho = -0.5$	104
7.1	Comparison of the “Now or Never” and Dynamic Value functions	115
7.2	Utility Indifference Valuation of the Option and the Project .	117
7.3	Waiting to invest: Optimal Portfolio Investment (π^*), Portfolio Contours and Investment Thresholds	119
7.4	σ_m and the manager’s behaviour	121

7.5	B and the manager's behaviour	125
7.6	σ_y and the manager's behaviour	127
7.7	Investment Threshold when $\sigma_y = 0$	128
7.8	λ and the manager's behaviour	129
7.9	α and the manager's behaviour	130
7.10	γ and the manager's behaviour	132
8.1	Payoff from hiring a manager ($S^M(W, Y, t)$) at $t = t_1$ and $t = t_n$	141
8.2	$S^M(W, Y, t)$ vs $P^M[\alpha, Y]$ and $P^{CAPM}[Y]$ vs $F^{CAPM}(Y, t)$. . .	142
8.3	$S^M(W, Y, t)$ vs $P^M[\alpha, Y]$ for $W = -50$ and $W = 50$	144
8.4	Value function ($V^{\text{static}}(W_1, Y_1)$) given optimal hiring decision .	146
8.5	$V^{\text{static}}(W_1, Y_1) - S^M(W_1, Y_1, 1)$	146
8.6	The Static Hiring decision: Base Case Parameters	147
8.7	κ and the static hiring decision	151
8.8	B and the static hiring decision	152
8.9	σ_m and the static hiring decision	153
8.10	σ_y and the static hiring decision	154
8.11	γ and the static hiring decision	155
8.12	α and the static hiring decision	157
8.13	The dynamic vs static hiring decision: the effect of flexibility .	160
8.14	σ_m and the dynamic hiring decision	163
8.15	B and the dynamic hiring decision	164
8.16	ρ and the dynamic hiring decision when $B = 70$	165

8.17	Difference in shareholder's payoff: $\{B=70, B=90\}$	167
8.18	Difference in shareholder's payoff: $\{\lambda=1, \lambda=0.5\}$	168
8.19	Difference in shareholder's payoff: $\{\gamma=0.5, \gamma=2\}$	169
8.20	γ and the static hiring decision	170
8.21	Difference in shareholder's payoff: $\{\alpha=0.2, \alpha=0.05\}$	171
8.22	The manager's investment decision: $\{\alpha=0.2, \alpha=0.05\}$	172
8.23	Difference in shareholder's payoff: $\{\kappa=0.1, \kappa=0.01\}$	173
9.1	The Value Function, Optimal Consumption and Optimal Investment	197
9.2	Value Function, Policy Functions and Income Multipliers for $Z_{max} = 1000$	199
9.3	Key Variables at $Z = 500$	202
9.4	Key Variables at $Z = 100$	203
9.5	Comparison of Hybrid and Numerical Boundary Conditions Income Multipliers for $Z_{max} = 200$ and $Z_N = 2000$	205
9.6	UI Discount Rate vs α	207
9.7	e^* as a function of W and Y	210
9.8	Investment threshold in the CRRA/GBM model with effort .	211
B.1	Comparative Statics For A	225
B.2	Comparative Statics For θ	226
B.3	Comparative Statics For ρ	227
B.4	Comparative Statics For μ	228

B.5	Optimal Effort ($\hat{e}$) and the Investment Cost ($I(\hat{e})$) when α^* is endogenous	229
C.1	κ and the dynamic hiring decision	231
C.2	σ_y and the dynamic hiring decision	232
C.3	γ and the dynamic hiring decision	233
C.4	α and the dynamic hiring decision	234

Part I

Motivation and Framework

Chapter 1

Introduction

In the last 40 years we have witnessed a startling shift in the way that managers are paid to run firms. Figure 1.1 shows that in the 1970's CEOs of S&P 500 firms were paid almost entirely in cash. In the 1980's equity-linked compensation begins to increase at a moderate rate before skyrocketing in the 1990's to the point where equity based compensation now dwarfs cash compensation.

Interestingly, the level of equity-linked compensation drops sharply in 2002. Figure 1.2 shows that this downward trend has continued in recent years with "long-term incentives" falling. It is not only the proportion of total compensation that is linked to equity that has decreased in recent years, Figure 1.3 shows that there has also been a compositional shift away from stock options towards restricted stock.

Figure 1.1: *Executive compensation over time*

Source: Jensen, Murphy, and Wruck (2004)

Additionally, Bettis, Coles, and Lemmon (2000) empirically examined executive compensation schemes and found that 92% of the firms in their sample have policies restricting trading by insiders and 78% have explicit blackout periods during which trading is prohibited by insiders. Similarly, Kole (1997) found that 79 of 371 Fortune 500 firms in her sample have such restricted stock plans. The average minimum holding period before any shares can be sold ranges from 31 months for firms with a medium level of research and development to 74 months for firms with a high level of research and development. For more than a quarter of the plans, the stock cannot be sold before retirement.

When presented with evidence showing that executive compensation has gone from being almost entirely cash in the 1970's to around 50% restricted stock and options in 2004, and that this equity exposure often has trading restric-

Figure 1.2: *Compensation Mix*

Source: Wall Street Journal and Mercer Consulting (2004)

Figure 1.3: *Incentive Mix*

Source: Wall Street Journal and Mercer Consulting (2004)

tions, the questions that instantly spring to mind are why has this happened? and what are the implications of large managerial ownership stake? As one would expect, the first question can only be partially answered without reference to the second question. This endogeneity is expected as the implications of managerial ownership are no doubt one of the main drivers determining the amount of stock that managers are given.

The primary justification of equity-linked compensation is an agency problem between the owners and managers of the firm. In their seminal work on this subject, Jensen and Meckling (1976) appealed to Adam Smith to explain the nature of this problem, and since this description's relevance has endured to this day, I will repeat it here

The directors of such [joint-stock] companies, however, being managers rather of other people's money than of their own, it cannot well be expected that they should watch over it with the same anxious vigilance with which the partners in a private copartnery watch over their own. Like the stewards of a rich man, they are apt to consider attention to small matters as not for their master's honor, and very easily give themselves a dispensation from having it. Negligence and profusion, therefore, must always prevail, more or less, in the management of such a company

The essence of this and other arguments is that if there is imperfect information¹ and the manager is not the sole owner of the firm, the manager does not bear the full cost of any inefficient actions he undertakes (e.g. diverting resources to themselves or simply shirking). This means that the manager has different incentives from shareholders and thus his optimal actions differ from those that will maximize the utility of shareholders. The analysis

¹This is essential, because in a world of perfect information shareholders could costlessly monitor managers. This would allow shareholders to write contracts specifying the action the manager must take in every state of the world, thus eliminating any agency problem.

of Jensen and Meckling (1976) predicts that as the information asymmetry between managers and shareholders becomes more severe it becomes easier for managers to shirk or consume perks. Therefore a higher proportion of managerial ownership becomes optimal.

It is an intuitively simple and appealing idea that if managers are given stock, then they will “think like shareholders”. However, managerial ownership can also have unintended, perverse consequences. For example, the “no skin the game” problem described by Jensen, Murphy, and Wruck (2004), the “zero cost of equity” problem for option grants described by Jensen (2001) and the reduced incentive for option holding managers to pay out dividends identified by Lambert, Larcker, and Larcker (1989).

The emphasis on incentive alignment has led many academics to believe that the important factor for executive compensation is the percentage of the firm’s outstanding share capital owned by the CEO, as opposed to the fraction of their wealth or pay that is made up of equity instruments². When examining the empirical evidence on the effects of executive compensation, Abowd and Kaplan (1999) posed the question as to whether the increased incentives imposed on executives might be doing more harm than good. They posit that this could happen by making managers overly cautious. They believe this is supported by the work of Hall and Liebman (1998), who found that CEOs can face large fluctuations in their wealth which will make risky projects unappealing even if they are profitable. Another way of putting this is that by giving managers large equity stakes, we may be unintentionally making them undiversified.

This thesis is based upon four very simple premises:

1. managers, not shareholders make the investment decisions for the firm;
2. managers do more than just say “yes” or “no” to investments³, they

²For example see Jensen and Murphy (1990a) and Jensen and Murphy (1990b).

³As mentioned in the fourth bullet point below, they can also decide to “wait” when

can also exert effort that affects the payoff from investment;

3. executive compensation schemes can cause managers to hold more stock than is optimal for diversification purposes; and
4. many investments can be delayed, involve irreversible capital costs and have uncertain payoffs.

While these issues have been examined in detail on their own, the interplay of all four has not. The first two issues have been well documented in the agency literature on executive compensation⁴. A general conclusion to come from this literature is that by giving managers equity exposure we can align their incentives with those of shareholders. If managers “think like shareholders” then the fact that they maximise their own utility rather than shareholder wealth is irrelevant, they will simply maximise the market value of the firm’s equity. Maximising the market value of equity will maximise shareholder wealth and thus incentives are aligned.

The third premise is often ignored when arguments are made that giving managers equity exposure is a good thing. By causing managers to hold more stock than is optimal, it is likely that they will care about the firm’s idiosyncratic risk, something that it is assumed shareholders can diversify away. The result of this is that while shareholders will care about the market value of the firm, managers will have their own subjective valuation of the firm which will incorporate the firm’s idiosyncratic risk. That is, there are two sides to the managerial compensation “coin” - effort and diversification.⁵

investment can be delayed.

⁴See Jensen, Murphy, and Wruck (2004) for a thorough review of the executive compensation literature.

⁵A lack of diversification may also cause managers to engage in “firm level diversification”. This hypothesis was first developed by Amihud and Lev (1981), who presented the diversification argument in the context of managers who have risky firm-specific human capital which is non-traded. They therefore would wish to engage in diversifying mergers to reduce this risk. They point out that from a shareholder’s point of view this is subopti-

The fourth premise is explored in the developing literature on “Real Options”. This literature has found that when the investment decision is irreversible and the payoff is uncertain, there may be “value in waiting”.⁶

Combining these four premises gives the two central questions this thesis attempts to answer

1. How does the level of managerial stock-ownership affect the investment decisions managers make for the firm? and
2. given the answer to (1), how does this affect the shareholder’s decision to hire a manager?⁷

Incorporating all four premises into one model is complex and thus a decision must be made on a framework that is both tractable yet adequately captures the dynamics at play. The “end goal” as such is to have a model in which both the manager and the shareholder make a dynamic decision (hiring for the shareholder and investing/exerting effort for the manager). To reach this end goal a number of intermediate steps must be taken:

1. the first step is to determine the manager’s valuation of a cash flow he is constrained to own a fixed proportion of;
2. with this in place we can examine his decision to invest and exert effort in a now or never setting (i.e. investment cannot be delayed);
3. having established the manager’s static investment/effort decision we can examine the shareholder’s hiring decision

mal, as they can adequately diversify the firm’s risk themselves. There is strong evidence showing that this type of diversification destroys value for shareholders, e.g. Berger and Ofek (1995), Lins and Servaes (1999), Lamont and Polk (2002) and Comment and Jarrell (1995). Furthermore, May (1995) found that CEOs with more of their wealth invested in the firm tended to diversify more.

⁶See Dixit and Pindyck (1994)

⁷Alternatively, “how does the shareholder compensate the manager?”

4. having examined the static hiring and investment decisions, we can then extend the modelling of the manager to incorporate the ability to delay investment (and thus the decision concerning timing and effort);
5. finally we can examine the shareholder's decision of whether or not to hire the manager when investment can be delayed. This can be done in both the context of a static (now or never) and dynamic (i.e. delayable) decision of whether or not to hire.

In this thesis I use a continuous time framework to model these steps. As it turns out, the form of the utility function assumed for the manager has a huge impact on the tractability of the modelling. The assumption of Constant Relative Risk Aversion (CRRA) utility as opposed to Constant Absolute Risk Aversion (CARA) causes the manager's valuation of the cash flow in the first step to become wealth dependent. This in itself is an interesting issue, but it increases the dimensionality of the problem, and introduces some very interesting numerical issues while also making the later steps of the process intractable. Because of this we split the substantive analysis of this thesis into two parts. In the first we assume CARA utility in order to remove wealth dependence from the valuation and obtain a "clean path" to the end goal of a dynamic model of hiring and investment. In the second we focus on CRRA utility thus allowing the manager's valuation to depend on his financial wealth. We then explore the resultant numerical issues.

We find that a manager who is constrained to hold a large portion of his wealth in the firm will value the firm's projects less than a shareholder because he is undiversified and thus cares about idiosyncratic risk. The result of this is that the manager may pass up projects that are $NPV > 0$ from shareholders' perspective. We find this result holds under both the wealth dependent Constant Relative Risk Aversion (CRRA) and wealth independent Constant Absolute Risk Aversion (CARA) utility functions. While previous work has already shown that in this setting the manager's valuation depends

upon his risk aversion and the level of idiosyncratic risk⁸, we show that his valuation of the *entire* project depends on the *proportion* of the firm he owns. Put simply, the more the manager owns of the firm, the less diversified he is and thus the lower he values the firm's cashflow. Another way to think about this is that the discount rate used by the manager will depend on his ownership stake in the firm. The additional feature introduced by a CRRA utility function is that the more wealth the manager has, the less of an issue diversification becomes. In essence, in the static valuation sense a rich manager thinks more like a shareholder. Intuitively, as the project becomes a smaller fraction of the manager's wealth, the less the constraint to hold the project affects the manager's diversification.

This is of course only one side of the coin. Managers are given equity⁹ under the presumption that they can have some positive effect on the value of the firm. The fact that too much stock can make the manager undiversified means that shareholders face a trade off when deciding on the optimal amount of equity to give the manager. To understand this trade off we extend the CARA model to allow the manager to exert effort at the time of investment to reduce the investment cost. This model allows us to examine how the manager's optimal level of effort changes in response to changes in various parameters. Interestingly we find that the manager's optimal level of effort is highly non-linear in the proportion of the firm he owns and that this relationship depends heavily on how risk averse the manager is.

Having modelled the manager's optimal level of effort in the situation where he is constrained to own a proportion of the firm he runs, it is now possible to model the shareholder's choice of an optimal level of managerial ownership. In contrast to the standard executive compensation literature, the manager's

⁸See Miao and Wang (2007), Munk (2000) and Svensson and Werner (1993).

⁹Note that by using the term equity we are not limiting the discussion to stock grants. Managers also obtain equity through stock option grants. The issue we are interested in is the level of managerial ownership, not the manner in which it is obtained.

effort is fully reflected in the shareholder's payoff.¹⁰ The shareholder therefore trades off incentivising the manager to exert more effort and making him less diversified. The level of managerial ownership that maximises the payoff to the shareholder varies significantly across different parameter values. We therefore conduct comparative statics analysis to see how the optimal level of managerial ownership varies across different parameters. Interestingly we find that market and project specific parameters do not have a very large effect on the optimal level of managerial ownership. By contrast, parameters that are specific to the manager, such as managerial skill, wealth, risk aversion and dislike of effort, have a very significant effect on the optimal level of managerial ownership. This is because these are the parameters that have the largest impact on the shareholder's payoff. This is interesting as it suggests that the type of firm the manager works for is relatively unimportant when determining his compensation; it is the unobservable characteristics of the manager that matter. Given that managerial characteristics are generally unobservable this poses a challenging problem for executive compensation design.

The real options literature has shown that in an environment of uncertainty there can be significant value in waiting. While the effect of non-tradeability has been examined in a real options setting¹¹, the issue of partial ownership by the manager and his ability to exert effort has not been examined in a dynamic setting where it is possible to delay investment. In this situation the fact that the manager is maximising his own utility rather than the market value of the project turns out to be very important. The manager effectively incurs two costs when investing, a financial cost¹² and the utility cost of

¹⁰I.e. there is no "noise" in the effort signal as in the standard model pioneered by Holmstrom (1979).

¹¹See Miao and Wang (2007) for a CARA/Simple Brownian Motion (SBM) model with consumption and Henderson (2007) for a CRRA/Geometric Brownian Motion (GBM) model without consumption.

¹²The manager is a shareholder and thus has to pay his proportion of the dollar investment cost.

the effort he exerts. The manager faces a direct trade off between these two costs in that reducing the financial cost increases the utility cost. Because the manager has diminishing marginal utility over wealth, the richer he is the less the financial cost means to him. Therefore for high enough levels of wealth he will not exert any effort. While the manager wishes to avoid effort because it is costly, he is also exposed to wealth shocks.¹³ If the manager invests now and then subsequently receives a negative wealth shock, he will wish he had exerted more effort when he invested as money is relatively more important to him now. This fear of slacking off and then regretting it *ex post* leads the manager to delay investing in the region where he would exert little effort if he invested. That is the manager wants to be very sure that if he slacks off he will not regret it and so he waits to invest. This is effectively the “bad news principle” as outlined by Dixit and Pindyck (1994) and discussed in Chapter 2. Given that the shareholder’s objective is to maximize the market value of the firm, he will simply want the manager to exert “maximum” effort.¹⁴ There is thus a disconnect between the manager’s choice of effort and what the shareholder desires. More interesting though is that the manager’s “effort related” desire to wait is decreasing with market volatility, contrary to the standard real options result that the value of waiting increases with uncertainty.

The shareholder can make the investment decision but incurs monitoring costs prior to investment. The shareholder thus waits for a manager with desirable characteristics and/or favorable cash flow outcomes. Unsurprisingly, the shareholder’s decision of whether or not to hire is heavily influenced by the likelihood that the manager will exert effort upon investment. In both the static and dynamic hiring models, the more likely it is the manager will exert effort upon investment (and thus hiring is *ex post* optimal), the more likely it is that a manager will be hired immediately. When the ability to

¹³This is due to his holdings of risky assets.

¹⁴In the context of this thesis, “maximum” means exerting enough effort to reduce the investment cost to its lower bound.

delay hiring is introduced, the “bad news principle” can be appealed to in order to explain the shareholder’s behaviour. In essence, when the manager has a “choice”¹⁵ over his level of effort, the shareholder will delay his hiring decision in order to avoid outcomes where the manager doesn’t exert effort upon investing.

While it is important to understand the hiring *decision* a shareholder would make when faced with a certain type of manager, it also important to understand what type of manager a shareholder would *prefer* to hire given the choice. To understand what type of manager the shareholder desires we must examine the payoff to the shareholder for different levels of the relevant managerial parameters. The typical result of this exercise is that the shareholder will generally prefer a manager who is less risk averse, more skilled and who requires a lower level of firm ownership. However, when one moves past generalisations, there are situations where the opposite is true and where the shareholder is indifferent.

The thesis is in four parts. This is graphically represented in Figure 1.4. This flow chart is a useful reference while reading this thesis and thus will be reproduced at the beginning of each chapter with the current chapter colored in green instead of blue.

In Part I we review the relevant literature (Chapter 2) and set out the general modelling framework (Chapter 3).

In Part II we set out the “clean path”. As will become clear shortly, the focal point of Part II is Chapters 7 and 8. The preceding chapters enable the analysis in those chapters to be carried out. Chapter 4 presents the manager’s valuation of the firm’s cash flow in a “now or never” setting when he has CARA utility and the cash flow follows Simple Brownian Motion (SBM), Chapter 5 extends the model of Chapter 4 to allow the manager to exert

¹⁵By “choice” we are referring to the area where the manager’s wealth is not so large that he ever exerts effort and not so small that he always exerts “maximum” effort. The concept of maximum effort is explained in Chapter 5.

effort at the time of investment to reduce the financial cost of investing, Chapter 6 then extends the model of Chapter 5 to examine the shareholder's problem of determining the optimal level of managerial ownership, Chapter 7 then extends the model of Chapter 5 to examine the manager's effort and investment decisions in a dynamic setting where it is possible to delay investment and Chapter 8 extends the model of Chapter 7 to examine the shareholder's hiring decision.

In Part III we examine the numerical issues and qualitative effects of assuming a CRRA utility function while Part IV contains our conclusions and appendices.

Figure 1.4: *Thesis structure*

Chapter 2

Literature Review

2.1 Introduction

In this chapter we review the theories of investment, valuation and managerial effort relevant to the problem we wish to model as outlined in Chapter 1. We begin in Section 2.2 by covering the standard Net Present Value (NPV) framework for investment and the Capital Asset Pricing Model (CAPM) which is the de facto valuation model used in standard NPV analysis. Section 2.3 reviews Real Options Analysis and contrasts it with the standard NPV framework. With the standard NPV and ROA theories in place, Section 2.4 discusses the work that extends these frameworks to a situation where the market is incomplete due to an inability to trade the asset (i.e. when a CEO has restricted stock and options). Finally work related to managerial effort is discussed in Section 2.5.

2.2 Standard Theory of Investment

The standard proposition concerning investment to come from financial economics is that firms will undertake projects with a Net Present Value (NPV) greater than zero. NPV is defined as the discounted present value of the project's revenues, minus the discounted present value of the project's costs.¹

This approach to valuation is generally referred to as Discounted Cash Flow (DCF) analysis and is generally accepted as the standard method for valuing projects. The area where contention generally arises is in the selection of the discount rate used to discount the project's risky cashflows. The standard approach is to use a variant of the Capital Asset Pricing Model (CAPM). While there are theoretical alternatives to the CAPM² as well as empirical questions about the empirical validity of the CAPM,³ it presents a tractable, intuitively sensible relationship between risk and returns. The CAPM thus serves as a useful benchmark for the analysis that will be conducted in this thesis.

At the heart of the CAPM is a simple proposition: investors can diversify away a portion of an asset's risk through holding it in a portfolio with other assets whose movements are not perfectly correlated. This results in a theoretically optimal holding of each asset in order to properly diversify away its "idiosyncratic" risk. Because idiosyncratic risk can be eliminated, investors do not require a premium for this and thus it is not priced by the market. The risk that cannot be eliminated through diversification is known as "sys-

¹See Brealey, Myers, and Allen (2006) and Copeland, Weston, and Shastri (2005) for text book expositions.

²Such as the Arbitrage Pricing Theory (APT) of Ross (1976).

³A review of the many studies of the empirical validity of the CAPM is beyond the scope of this thesis and thus the reader is pointed to Copeland, Weston, and Shastri (2005) for a good review of the major papers on the subject. It is also worth mentioning that the many works of Eugene Fama and Kenneth French (see, e.g. Fama & French (1992, 1995, 1996)) while not explicitly testing the CAPM, show that a three factor model better explains the cross section of expected returns.

tematic” or market risk. Given that we will be using the CAPM valuation as the benchmark for this thesis, it is useful to review the CAPM valuation of the cashflows we will be considering. The two different cashflow processes we will be considering in this thesis are Geometric Brownian Motion (GBM) and Simple Brownian Motion (SBM)⁴. In continuous time, a perpetual cash flow Y_t that follows these processes will evolve according to the following processes

$$GBM : dY_t = \mu_y Y_t dt + \sigma_y Y_t d\xi$$

$$SBM : dY_t = \mu_y dt + \sigma_y d\xi$$

where μ_y and σ_y are the drift and volatility of the respective processes and $d\xi$ is a wiener process.⁵ Under the assumptions of the CAPM, the continuous time valuations of the cash flows are

$$GBM : \frac{Y_t}{r + \rho \frac{\sigma_y}{\sigma_m} (\mu_m - r) - \mu_y}$$

$$SBM : \frac{Y_t}{r} + \frac{\mu_y - \rho \sigma_y \Phi}{r^2}$$

where r is the risk free interest rate, μ_m and σ_m are the expected return and volatility of the market portfolio⁶ respectively, ρ is the correlation between the cashflow and the market portfolio and Φ is the Sharpe ratio of the market portfolio.⁷ What these two expressions say is that the cash flow is valued using a discount rate that reflects the “systematic” or market risk of the cash flow. The key parameter that governs how much systematic risk the cash flow has is ρ in both cases. This is because the correlation of the cash

⁴The relative merits of these different processes will be discussed in Chapter 3 when we discuss the general modelling framework and assumptions we will be making.

⁵Note that μ_y and σ_y have slightly different meanings when we talk about a GBM and SBM process given that volatility and drift are multiplicative under the assumption GBM but additive in SBM.

⁶The market portfolio is assumed to follow GBM.

⁷Where the Sharpe ratio is a measure of an asset’s risk-return trade off and is defined as $\Phi = \frac{\mu_m - r}{\sigma_m}$.

flow with the market represents how much of the cashflow's risk is caused by general market movements and thus cannot be diversified away. To illustrate this point it is useful to consider the case where $\rho = 0$ and thus where the cashflow has no systematic risk. In this case:

$$\begin{aligned} GBM : & \quad \frac{Y_t}{r - \mu_y} \\ SBM : & \quad \frac{Y_t}{r} + \frac{\mu_y}{r^2} \end{aligned}$$

From the above equation we can see that when $\rho = 0$, the CAPM valuation depends only on the current level of the cashflow (Y_t), the riskless interest rate (r) and the drift term of the cashflow. While the cashflow is risky (i.e. $\sigma_y > 0$), because all of the cashflow's risk can be diversified away the risk is not priced.

Bringing everything together, in the simple case where there is a lump sum cost of investment I ⁸, investment will occur in the standard CAPM/NPV framework if

$$\begin{aligned} GBM : & \quad \frac{Y_t}{r + \rho \frac{\sigma_y}{\sigma_m} (\mu_m - r) - \mu_y} \geq I \\ SBM : & \quad \frac{Y_t}{r} + \frac{\mu_y - \rho \sigma_y \Phi}{r^2} \geq I \end{aligned}$$

2.3 Real Options Analysis (ROA)

2.3.1 Overview

In recent years a literature has emerged that recognizes and corrects for a number of assumptions that often do not hold in practice but which are im-

⁸i.e. The firm pays I and immediately starts receiving Y every period.

plicitly made when conducting Standard NPV analysis. Copeland, Weston, and Shastri (2005) note that empirically managers often invest in projects that have a negative NPV, something that is irrational in the standard framework. They point out this is because standard NPV analysis typically ignores at least five options that are embedded in many projects:

1. The Expansion Option

If a project turns out to be much more profitable than expected the manager can expand the scale of the project

2. The Extension Option

Similar to the expansion option, if the project turns out to be particularly profitable the manager can extend the life of the project

3. The Contraction Option

In direct contrast to the expansion option, if the project turns out to be less profitable than expected the manager can reduce the scale of the project.

4. The Abandonment Option

If the project turns out to be significantly unprofitable, the manager can abandon the project to avoid future losses.

5. The Delay Option

If investment is irreversible, the manager may wish to delay investment in order to receive more information concerning the profitability of the project.

Incorporating the value of these options into our analysis explains why it is possible to observe managers investing when the standard NPV is negative. To see why, we can think of investment in terms of cost benefit analysis. Cost benefit analysis states that an action should only be undertaken if the

benefits exceed the costs. The benefits in this situation are the present value of the project's revenues plus any options that are created by investing.⁹ and the costs are the financial investment cost and any options that are destroyed by investing¹⁰ Thus when options are incorporated the manager should only invest if

$$\text{Present Value of Net Revenues} + \text{Options Created} \geq I + \text{Options Destroyed} \quad (2.1)$$

Thus it is entirely possible that the standard NPV rule will be violated (i.e. $NPV \not\geq 0$), but investment will occur because the optionality of the project means that the total benefit of investing exceeds the total cost.

2.3.2 The Option to Delay Investment

Because the real options portion of this thesis is going to focus on the delay option, it is important to analyze it in some detail. When analyzing the option to delay investment, real options analysis (ROA) typically makes the following *explicit* assumptions¹¹

1. The investor is risk neutral¹² *or* the asset can be perfectly spanned *or* the investor is properly diversified.
2. The investment is irreversible.
3. Investment is not a “now or never” proposition and can be deferred.
4. There is uncertainty regarding the payoff or costs of the project.

⁹Such as the expansion, extension, abandonment and contraction options

¹⁰Such as the delay option.

¹¹ There are of course other implicit assumptions that are made which we will discuss separately.

¹²A risk neutral investor cares only about the expected payoff and not its distribution. Risk Neutral investors therefore discount future cash flows using the risk free rate.

The main effect of the first assumption is that idiosyncratic risk doesn't matter as the assumption of risk neutrality or perfect spanning means that risk neutral pricing can be used,¹³ while a risk averse investor who is properly diversified will only care about systematic risk and thus the CAPM applies.¹⁴ Given that the primary purpose of this thesis is examining situations where idiosyncratic risk matters, for our purposes the distinction is not crucial as both methods share similar general results. For there to be value in deferring investment, the second assumption is required. This is because if the investment was completely reversible, a firm could invest now and restore itself to its pre-investment position if it subsequently receives bad news. In other words nothing is lost by investing now. The requirement of the third assumption is obvious: if by deferring the investment decision the firm loses the ability to invest there is no value in waiting. The final assumption is required so that something is actually gained from waiting. If there is no uncertainty regarding the profitability of the project then no information is gained by waiting and thus waiting is pointless.¹⁵

The main conclusion from the literature concerning the option to wait is that when there is uncertainty there can be value in waiting. The result of this is that the firm will invest only when the present value of the project's net revenues exceeds the investment cost by a strictly positive amount.¹⁶ An important implication of this is that the hurdle rate used by the firm will be higher than its Weighted Average Cost of Capital (WACC) to reflect the fact that the investor needs to be compensated for the delay option that is

¹³See Dixit and Pindyck (1994) for an overview.

¹⁴ McDonald and Siegel (1986) present the first treatment on valuing the option to wait when the investor is properly diversified.

¹⁵There is a present value related reason to wait when there is no uncertainty which is discussed in Section 5.1.A of Dixit and Pindyck (1994). This is however simply a matter of timing the investment to maximise the present value, rather than waiting to gain new information.

¹⁶For textbook expositions of this see Dixit and Pindyck (1994) and Trigeorgis (1996).

destroyed once investment occurs.¹⁷ More importantly though, is the fact that the value of the option to wait, and thus the investment threshold and hurdle rate used, is increasing in the volatility of the project's cashflow. This is significant as the standard CAPM valuation of the cashflow, and thus the investment threshold, does not depend on the volatility of the cash flow.

2.3.3 Criticisms: Is There Really Value in Waiting?

The basic proposition that there is value in waiting has lead to a proliferation of academic work¹⁸ as well as growing acceptance among practitioners of the importance of using ROA when making investment decisions. This has of course led to closer scrutiny of the implicit assumptions made when conducting standard ROA. As Triantis (2005) points out, ROA often reflects “perfection” rather than economic reality. Many of the criticisms raised do in fact reduce the value of waiting, but as we will discuss, once the model's assumptions are relaxed to more closely reflect economic reality, there is still value in waiting, it is just not as large as the standard model suggests. In what follows we examine the main objections that are raised against ROA.

Competition Destroys the Value of Waiting

The intuition behind this argument is very simple: by waiting, a firm provides competitors with the ability to invest first and obtain any available first-mover profits.¹⁹ However, it is easy to over-state this argument as it requires some quite restrictive assumptions. In particular, it assumes that all

¹⁷Empirical evidence that managers do often use a hurdle rate that is higher than the WACC is presented by Poterba and Summers (1995), Meier and Tarhain (2007) and Chirinko and Schaller (2009).

¹⁸See Schwartz and Trigeorgis (2004) for a collection of some of the most important contributions to the Real Options literature.

¹⁹Bulan, Mayer, and Sommerville (2009) provide empirical evidence supporting this using real estate markets.

investments can be characterised as an “early bird gets the worm” situation. While there may indeed be advantages to investing before ones competitors in some cases, there are other situations where it is beneficial to let these competitors go first. For example, suppose the proposed investment is in a new and highly uncertain market or technology: by letting a competitor go first and observing its fortunes, much of the firm’s investment risk can be eliminated. When such second mover advantages are present, competition actually reinforces the value of waiting.²⁰

Many investments are likely to contain elements of both first and second mover advantages, so the total effect of competition on the value of waiting depends on which effect dominates. Only when pre-emption is essential (i.e. the first mover effect dominates) will the strategic importance of waiting become negligible. Recent empirical work shows that the benefits from pre-emption are dependent on industry structure, but the relationship is not a one way street. In fact, it has been found that while competitive industries do indeed invest faster than monopolistically competitive firms, firms in the least competitive industries actually invest the fastest.²¹ This suggests that pre-emption is much more important in less competitive industries, which makes sense given that in a less competitive industry the “prize” from winning the investment “race” will often be greater since there are fewer agents to share it with.

In a slightly different context (i.e. looking beyond first and second mover advantages), Novy-Marx (2007) shows that options still have value in competitive settings when there is cross-sectional variation in firm size and investment is lumpy and Guthrie (2010) shows that there is option value when firms have different cost structures in competitive settings.

²⁰The key paper to present this result is Childs, Ott, and Riddiough (2002) who show that competition does not destroy the value of waiting when the first firm to invest reveals information about the true state of demand. More generally, see Boyer, Gravel, and Lasserre (2004) for a survey of real options models involving strategic competition.

²¹See Akdoğan and MacKay (2008).

Credit Constraints

Investment must be paid for, whether out of internal funds (retained profits) or external funds (sale of new securities). Clearly, the ability to invest depends heavily on the availability of such financing, and on its price and its terms. Firms that must pay a high price to obtain new external finance (perhaps because they are deemed to be high risk or there are information asymmetries) will rely on internal funds, and therefore on the profitability of existing assets.

NPV and most real options models ignore the financing problem, instead simply assuming that investment will be paid for somehow. But firms that rely on internal financing run the risk that this funding may not be available in the future: for example, an adverse shock to profits may deplete internal funds to the extent that investment becomes impossible. In this situation, waiting is less valuable because of the risk that the investment opportunity may in effect disappear.

Although the possibility of financial constraints weakens the advantages of waiting to invest, these advantages do not disappear entirely. Even a severely cash constrained firm benefits from acquiring new information - the optimal waiting time is simply shorter than if it were unconstrained.²²

Investments take “Time to Build”

In the standard real options world, investment occurs instantaneously at the commencement of the project; subsequent investment is not required. In practice most projects take “time to build” - that is, they require implementation and construction over a period of time. As a result, typical investment expenditure is ongoing rather than a one-off lump sum.

Investments that begin, and then take time to complete, can of course be

²²See Boyle and Guthrie (2003) and Boyle and Guthrie (2006).

abandoned before completion if new information suggests that this would be the optimal strategy. Such projects are, in effect, more reversible than simple lump-sum projects, since the remaining investment cost can be avoided by abandoning the project. But the more reversible the project, the lower the incentive to delay its launching in order to acquire more information about its prospects. A longer implementation period thus decreases the value of waiting to invest.²³

This phenomenon is exacerbated if the risks surrounding the project's ultimate cost are primarily of the 'technical' variety (uncertainty about the time needed for completion and the quantity of inputs required).²⁴ These kinds of risks are generally only resolved by having construction commence, so that delaying investment provides no potential for additional information and hence is not valuable. By contrast, input price risk (uncertainty about the price of inputs) remains, whether or not construction is currently active. So there is value in waiting to gain new information about this even for projects that take time to build.

In the case of major infrastructure projects (for example the building of transmission investment or a new stadium for a sports event), a certain level of capacity must be in place at a known date in the future.²⁵ When a project has a long or uncertain construction period, the value of waiting is reduced because waiting may leave insufficient time for completion by the required date.

Mean Reversion

Real options models typically assume that shocks to the value of the project value follow a "random walk" process and that these shocks are permanent. Thus the value of a project can be subject to repeated adverse shocks, and the

²³See Milne and Whalley (2000).

²⁴See Pindyck (1993).

²⁵See Boyle, Guthrie, and Meade (2006).

incentive to delay investment arises from a desire to minimise the probability of such an outcome.

However, some projects are more accurately thought of as mean-reverting: when their value lies above or below a long-run mean, it tends to revert back towards that mean.²⁶ As a result, negative shocks tend to be followed by positive shocks, which lowers the potential magnitude of adverse outcomes and thus would seem to lessen the value of waiting.

But this is not the whole story. If mean project value exceeds the investment cost, then even if the project value is currently low, the project is likely to eventually be worth more than it costs. In such a case, the value of waiting can be even greater than in the standard situation.²⁷

Costly Information

The decision on whether to invest or delay requires calculation of a project's profitability. But unlike the holders of financial derivatives (on which real options theory is based), investment managers cannot continuously re-evaluate project profitability. The complexity of most real world projects means that such calculations are time consuming and costly.

Costly evaluations lower the value of waiting in two ways.²⁸ First, evaluation costs directly increase the project's cash outflows and thus lower its value. When the acquisition of further information about the project is costly, the opportunity cost of delaying investment to acquire this information is greater and hence the incentive to wait is lower. Second, and more subtly, higher evaluation costs lead to less frequent evaluations (since doing so continuously would be prohibitively expensive), and therefore to a lower probability of

²⁶This is particularly true for projects whose value depends on commodity prices.

²⁷Dixit and Pindyck (1994) value the option to wait when the project's value follows a mean reverting stochastic process.

²⁸See Guthrie (2007).

choosing the best time to invest. Since much of the value of waiting stems from having the flexibility to invest on exactly the right date, this effect reduces the incentive to wait in the first place. This is particularly important when project value is mean-reverting, as the temporary nature of value shocks means that getting the timing of investment exactly right is crucial.

Summary

Although these real-world factors drive a wedge between the true value of waiting and that predicted by simple theoretical models, the size and sign of this wedge is frequently unclear. Even when the waiting value is unambiguously smaller than its theoretical counterpart, it is extremely unlikely to be zero.

2.4 Non-Traded Assets and Idiosyncratic Risk

The central idea behind this thesis is that the managers of firms are forced to hold too much stock in the companies they work for and thus are unable to eliminate idiosyncratic risk. If managers are able to trade their stock then this is not a problem as they could simply sell down their stake to the level that is optimal for diversification purposes. While a fair amount of work has examined the valuation of non-traded stocks and derivatives in incomplete markets, non-traded stochastic cashflows have received comparably less attention. The purpose of this section is thus to review the literature concerning the valuation of non-traded cash flows in both a static and real options context.

We begin in Section 2.4.1 by briefly reviewing some of the general literature on the valuation of non-traded stocks and options.

We therefore begin by reviewing the literature regarding the valuation of

a non-traded stochastic cashflow, before proceeding to discuss the recent literature concerning the valuation of the delay option when the underlying asset is non-traded. This is sensible because it is difficult to understand how an option on a non-traded asset is valued without first discussing how the non-traded asset itself is valued. The key insight that we shall find is that because idiosyncratic risk cannot be eliminated, the valuation is no longer preference free. Therefore factors such as the manager's level of risk aversion and personal financial wealth can have significant impacts on the decisions made.

2.4.1 General literature on non-traded options/stocks

The majority of the work concerning the valuation of non-traded assets has focused upon how the presence of a non-traded asset affects an agent's optimal consumption and investment in risky assets. In terms of consumption/savings decisions, the general result of this literature is that the classical two-fund separation result no longer holds.²⁹ This is because an investor's demand for risky assets no longer depends solely on its mean and variance, but also on its ability to hedge the risks associated with their non-traded income.³⁰ From a valuation perspective, since the demand for risky assets depends on the ability to hedge non-traded income, many standard asset pricing models will not hold.³¹ The presence of unhedgeable risks is likely to

²⁹The two-fund separation result originates from the work of Tobin (1958) and states that the decision concerning choice of the optimal mix of risky assets can be separated from the choice of how much money to invest in the portfolio of risky assets. Thus the decision for the investor is simply how much money to allocate to the risky fund of assets and how much to allocate to the risk free asset.

³⁰See Campbell and Viceira (2002) for discrete time models of consumption/saving decisions in the presence of non-traded income.

³¹For example He and Pages (1992) show that Merton's Intertemporal Capital Asset Pricing Model (ICAPM) doesn't hold in the presence of non-traded income, but a version of Breeden's Consumption CAPM with modified Euler equations does hold.

reduce the valuation an agent places on an asset relative to the case where the asset is freely traded. For example, Kahl, Liu, and Longstaff (2003) have shown that an agent's valuation of restricted stock is much less than the value he would place on it were it tradeable, and Ingersoll (2006), Meulbroek (2001) and Hall and Murphy (2002) have developed models showing that the value of an employee stock option can be anywhere between 50-70% of its Black-Scholes value.

2.4.2 Valuation of non-traded cashflows

The starting point for the valuation of a non-traded cashflow in continuous time is the simple model of Merton (1969). In this model the investor has either CARA or CRRA, can freely trade in risky assets as well as the risk free asset and does not have a non-traded source of income. This model has since been extended in a number of ways to examine the role of idiosyncratic risk and non-tradability.

Merton (1971) extended the classic no-income model of Merton (1969) by examining the situation where a CARA investor receives an income stream that follows a Poisson process. He finds that investors treat the certainty equivalent value of their lifetime income as an addition to their current wealth. Thus investors consume a constant fraction of their "total wealth" and invest a constant fraction of their financial wealth into risky assets based upon their mean and variance. In this situation two-fund separation still holds as a Poisson process means that shocks to income are entirely idiosyncratic.

He and Pages (1992) isolated the effect of non-tradability by examining the case where an investor has a perfectly-spanned income stream that he cannot trade or borrow against. They found that in this case the liquidity constraints alter the effective planning horizon of the investor, causing the investor to smooth consumption over time more than they would in the absence of liquidity constraints. He and Pages also show that as wealth goes to

infinity, the optimal consumption and savings policies converge to the policies that would be chosen without the constraints. Intuitively this means that as income becomes a smaller part of total wealth, the constraints on income become irrelevant. Under the same general assumptions, El Karoui and Jeanblanc-Picqué (1998) derived a closed form solution for consumption as a function of wealth and income. They used this to show that at zero wealth, a smaller fraction of income is consumed relative to the unconstrained case. This result is broadly consistent with the consumption-smoothing result of He and Pages (1992).

Duffie, Fleming, Soner, and Zariphopoulou (1997) relaxed the assumption of complete markets, and showed that when the investor has CRRA utility and the income process follows GBM, the Hamilton-Jacobi-Bellman (HJB) equation is a second order non-linear partial differential equation. Crucially, the assumption of CRRA utility means that the valuation of the cash flow now depends on the agent's financial wealth. The key contribution of this paper was to show that the dimensionality of the HJB equation can be reduced by transforming the model into a function of one variable (the ratio of wealth to income $Z \equiv \frac{W}{Y}$). Using this transformed model they were able to prove that the value function is a constrained-viscosity solution to the HJB equation. While they were unable to find analytical solutions they proved that as the ratio of wealth to income reached infinity, the investor would behave as if he was in a Merton(1969, 1971) world with *no income*.³² This result is very similar to the limiting result of He and Pages (1992). The implication is that an investor (or manager) who has sufficiently large financial wealth is indifferent to the parameters of the cashflow.

Making the same assumptions about the investor's opportunity set as Duffie, Fleming, Soner, and Zariphopoulou (1997), Koo (1998) derives the implicit value the agent places on the income stream and uses this to characterize

³²This stems from the fact that the model of Duffie, Fleming, Soner, and Zariphopoulou (1997) differs from the framework of Merton (1969,1971) only through the addition of non-traded income.

the optimal policies.³³ This implicit value is defined as the marginal rate of substitution between income and wealth. Koo shows that the implicit value of the income stream is, in general, less than the complete markets valuation, consistent with the stock and option pricing models mentioned previously. As with Duffie, Fleming, Soner, and Zariphopoulou (1997), Koo was unable to analytically solve the model. However, he did prove that the value function, optimal policies and implicit valuation of income converge to their complete market counterparts as the ratio of wealth to income goes to infinity.

Building upon the work of Duffie, Fleming, Soner, and Zariphopoulou (1997) and Koo (1998), Munk (2000) makes two major contributions. The first is to recognize that the implicit valuation derived by Koo is not an entirely satisfactory measure of value. This is because the implicit value is by definition a marginal value and thus does not offer an entirely satisfactory account of the value of the entire stream to the investor.³⁴ Munk therefore derives the utility indifference value of the income stream.³⁵ While the implicit and utility indifference valuations differ initially, Munk proves that they converge as the ratio of wealth to income goes to infinity. Munk's second contribution is to show how a simple converging numerical method can be used to solve the HJB equation numerically. He uses the Markov chain approximation method and confirms that the valuation of the income stream is much lower than the complete markets valuation. Using the numerical solution to the HJB equation he is also able to compute the agent's optimal consumption and investment in the risky asset. He confirms the analytical results of Duffie, Fleming, Soner, and Zariphopoulou (1997) by showing that when the ratio of wealth to income is very large, optimal consumption and investment are constant across different levels of the income stream. More interesting is the

³³These policies refer to the consumption and portfolio decisions of the agent.

³⁴The valuation of the entire stream is what is relevant to this thesis since we are analysing investment decisions.

³⁵See Henderson (2009) for an overview of the utility indifference pricing method.

behavior of optimal consumption and investment when wealth is relatively low. Munk finds that both consumption and investment increase rapidly with income initially, and then become proportional to income as income becomes relatively large.

Miao and Wang (2007) take a slightly different approach by making assumptions that simplify the analysis considerably compared to the Duffie, Fleming, Soner, and Zariphopoulou (1997) framework. By assuming CARA utility and thus eliminating wealth effects they are able to derive the following closed form solution for the investor's valuation of a non-traded SBM cashflow

$$G[Y] = \left(\frac{Y}{r} + \frac{\mu - \rho\sigma_y\Phi}{r^2} \right) - \frac{\gamma\phi^2}{2r^2} \quad (2.2)$$

which is simply the CAPM valuation of an SBM cashflow presented in Section 2.2, with an additional term that decreases the valuation of the cashflow depending on the manager's risk aversion (γ) and the level of idiosyncratic risk (ϕ).

2.4.3 Real Options

The major papers concerning the valuation of the option to wait when the underlying asset is non-traded take two different approaches. Henderson (2007) assumes that the project value is stochastic and thus the investor receives a lump sum upon investment. This is quite different from assuming that upon investment the investor begins to receive a risky cashflow. The result of this is that investing actually eliminates the investor's exposure to the project's risk. While this type of situation does describe many investments,³⁶ it is not the primary concern of this thesis.

³⁶The primary example is that of an apartment development which is then sold off upon completion as opposed to being managed by the developers. In this situation the developer is bearing the risk prior to completion, but once the project is completed the investors who purchased units then bear all the risk.

The analysis of Miao and Wang (2007) is more complete than that of Henderson (2007) in that the authors analyse both the case where the payoff from investment is a lump sum and where it is a stochastic cashflow, allowing a comparison between the two. Given that our primary interest is the analysis of cashflows, the majority of this section will focus on this paper. However, as will be seen shortly, the approach of Henderson (2007) is unique and thus it is important to analyse the difference in methodology between the two papers. A third paper by Hugonnier and Morellec (2007b) also analyzes the value of the option to wait when the underlying asset is non traded, but as we shall see their analysis is incomplete.

Henderson (2007)

Henderson (2007) models the problem of an entrepreneur who has an infinite horizon CARA utility function defined over financial wealth (x in her notation)

$$U(x) = -\frac{1}{\gamma}e^{-\gamma x}, \gamma > 0 \quad (2.3)$$

By defining utility over wealth, consumption has been eliminated from the model, thus simplifying the analysis. The entrepreneur is the sole owner of a firm whose only asset is the right to invest in a project that will pay the lump sum V_t at the time investment of the lump sum cost I is paid. The entrepreneur also has access to a risky asset (whose price is denoted P_t) which partially spans movements in V_t . Because the firm's asset is non-tradable and P_t only partially spans the risk associated with V_t , we are in an incomplete market and thus idiosyncratic risk matters. V_t and P_t are assumed to evolve according to the following GBM

$$\begin{aligned} \frac{dV}{V} &= \eta(\xi dt + dW) + rdt \\ \frac{dP}{P} &= \sigma(\lambda dt + dB) + rdt \end{aligned}$$

where $\xi \equiv \frac{v-r}{\eta}$ and $\lambda \equiv \frac{\mu-r}{\sigma}$ are the Sharpe ratios of the two assets. dW

and dB are the two Weiner processes whose correlation is denoted by $-1 \leq \rho \leq 1$. The novelty of Henderson's approach is the derivation of what the author refers to as a "time consistent utility function". This approach is taken because of the issues that arise when trying to value cash flows at intermediate times when the investor has no consumption. The author shows that the time consistent utility function that is maximised at the time of investment is

$$U_\tau(x) = -\frac{1}{\gamma} e^{-\gamma r \tau x} e^{\frac{1}{2} \lambda^2 \tau}$$

where τ is the time that investment occurs. This differs from the standard utility function in two ways: the $-r\tau$ term converts the investment payoff into today's money and the $e^{\frac{1}{2} \lambda^2 \tau}$ makes an adjustment for the fact that the investor's wealth has been optimally invested in the risky asset. CARA utility's lack of a wealth effect in valuation combined with the use of time consistent utility allows Henderson to obtain a closed form solution for the investor's value function which is then used to determine a certainty equivalent valuation of the option.

Her results are very interesting when compared to the standard results from the real options literature. The primary findings are the following

1. In general if $|\rho| < 1$ and $\gamma > 0$ then the value of the option and thus the investment threshold is less than the McDonald and Siegel (1986) and risk neutral values.
2. The value of the option to wait and thus the investment threshold is increasing in $|\rho|$.
3. The value of the option to wait, and thus the investment threshold, is decreasing in γ .
4. It is possible for the value of the option to decrease with uncertainty.
5. The risk neutral valuation and the model of McDonald and Siegel

(1986) are contained as special cases of the model when $\gamma \rightarrow 0$ (McDonald and Siegel (1986)) and $\rho \rightarrow 1$ (risk neutral).

The intuition behind the first three results is fairly straightforward. In general when an agent is forced to bear idiosyncratic risk he values the project less than the situation where he is not exposed to that risk. Thus it makes sense that he values the option less in this situation than he would if he were not exposed to this risk. As $|\rho| \rightarrow 1$, we are approaching a situation where the risky asset perfectly spans V_t and thus the project no longer has any idiosyncratic risk. Therefore it is sensible that the agent's valuation of the option increases. The manager's dislike of idiosyncratic risk is dependent on his risk aversion and thus it is unsurprising that as his risk aversion increases (and thus his dislike of idiosyncratic risk increases) his valuation of the option decreases. The fourth result is slightly more complex. Because the agent dislikes idiosyncratic risk, an increase in volatility has two effects: idiosyncratic volatility increases and the convexity of the payoff increases³⁷. The first effect reduces the value of the option while the second increases the option value via the standard channel. Depending on parameter values either effect can dominate.

Based upon this the author argues that the use of either the McDonald and Siegel (1986) or risk neutral investment rules can lead to sub-optimal under-investment, as this model predicts that investment can occur much earlier than the standard models suggest. This conclusion relies upon the manager being the sole owner of the firm. If the manager is a minority shareholder and the other investors have different risk preferences to the manager or are properly diversified while the manager is not, the conclusion is likely to be quite different. If the manager makes a decision based upon the incomplete model, then from a shareholder's perspective he may be investing too early which would result in over investment as far as shareholders are concerned. Acknowledging the difference between managers and shareholders will be the

³⁷That is, the downside that is avoided by not investing increases.

primary concern of this thesis.

One thing that Henderson does not consider is the distinction between a lump sum payoff and a stochastic cashflow upon investment. When a lump sum is assumed, the payoff from investment is independent of risk preferences. In this situation risk aversion decreases the value of the option and not the payoff. Thus it is perhaps unsurprising that investment is sped up. As we shall see shortly from the analysis of Miao and Wang (2007), if the payoff is instead a stochastic cashflow, then risk aversion will also affect the manager's valuation of the payoff. Therefore the net effect on investment timing depends on whether the effect of risk aversion is greater on the option or the payoff.

Miao and Wang (2007)

Similarly to Henderson (2007), Miao and Wang assume that the investor has CARA utility and thus there is now wealth dependence in the problem. However they assume that utility is defined over consumption (C_t) and thus takes the following form

$$U(C_t) = -\frac{1}{\gamma}e^{-\gamma C_t}, \gamma > 0$$

By defining utility over consumption, wealth serves a purpose at intermediate points in time and thus a time consistent utility function does not need to be derived as it was in Henderson (2007). However, the inclusion of consumption does result in the investor's HJB equation becoming a second order non-linear differential equation which has no analytical solution. Therefore the HJB is solved numerically using the projection method implemented with collocation.³⁸ The authors examine both the case where the payoff from investment is a lump sum and the case where the payoff is a stochastic cashflow. They examine both of these for the cases of "self insurance" (the investor only has access to the risk free asset) and when there is a risky asset

³⁸See Judd (1998) for an overview of this method.

that can be used to partially hedge the project. Given that the focus of this thesis is not on self-insurance, we focus solely on the partial spanning case.

In contrast to Henderson (2007), Miao and Wang assume that the investment state variable (which represents either a lump sum received at the time of investment or a cash flow) follows SBM (as opposed to Henderson's assumption of GBM). To capture the difference between systematic and idiosyncratic risk they assume that the state variable evolves as follows

$$dX_t = \alpha_x dt + \rho \sigma_x dB_t + \epsilon_x d\tilde{B}_t$$

where B and $\tilde{B}$ are independent wiener processes and ϵ_x is the project's idiosyncratic volatility, which can be defined as

$$\epsilon_x = \sqrt{1 - \rho^2} \sigma_x$$

For the lump sum case the results of Miao and Wang are consistent with those of Henderson (2007) in that they find the value of the option to wait is lower than the risk neutral valuation and is decreasing in both risk aversion (γ) and the level of idiosyncratic volatility (ϵ_x^2).

If the investment payoff is instead a flow, then there are two competing effects on investment timing. The first effect is the same as for the lump sum case in that risk aversion and idiosyncratic risk decrease the value of the option, something that encourages investment. The second effect is that after investment the manager is receiving a stochastic cashflow. The fact that he cannot fully eliminate the cashflow's idiosyncratic risk means that risk aversion and idiosyncratic risk decrease the manager's valuation of the cash flow, something that discourages investment. The net effect on investment timing thus depends on which effect dominates. The authors then go on to show that the project value effect dominates the option effect and thus idiosyncratic risk and risk aversion delay investment when we have flow payoffs. The main implication of this result is that statements cannot be made about the effect of idiosyncratic risk on investment timing without first considering the nature of the investment payoff.

Hugonnier and Morellec (2007)

The paper of Hugonnier and Morellec (2007b) shares elements of both the self insurance model of Miao and Wang (2007) and the model of Henderson (2007). They consider an agent with CRRA utility defined over wealth who is the owner of a firm that has the option to invest in a project that will deliver a GBM cashflow upon completion. The manager has no consumption and does not have access to a risky asset to partially hedge the project's cashflow. Interestingly they do not employ the time-consistent utility approach of Henderson despite the fact that there is no consumption in the model. The agent has access to the risk free asset and is assumed to have the cost of investment (I) invested in the riskfree asset at all times prior to investment, thus generating an income stream of rI at every point in time prior to investment. They find that risk aversion causes the manager to delay investment relative to the McDonald and Siegel (1986) case. They posit that this is because by investing the agent is exchanging a risk free cashflow (rI) for a risky cash flow and thus the agent wishes to delay exposing himself to idiosyncratic risk. This is the same result as found by Miao and Wang (2007) when the payoff from investment is a cashflow. However, in this case the result is a direct result of the setup of their model. Hugonnier and Morellec (2007b) have modelled the problem essentially ignoring the fact that the agent is holding an option prior to investment and thus is exposed to idiosyncratic risk through that option. Thus rather than investment being delayed because the effect of idiosyncratic risk is greater on the project value than on the option value, it is delayed to put off transitioning from a risk free state to a state where they are exposed to risk.

2.4.4 Implications for Well Diversified and Risk Neutral Shareholders

The general conclusion to come from the models concerning non-tradability, risk aversion and idiosyncratic risk is that exposing an agent to an asset's idiosyncratic risk will lower his valuation of that asset. In a "now or never" investment setting this means that managers are likely to value the firm's projects less than shareholders and thus could pass up projects that have a positive NPV from a properly diversified/risk neutral shareholder's perspective.

When the analysis is expanded to projects where delay is possible, we can have either over or under investment from a well diversified/risk neutral shareholder's perspective. If the project's payoff is of a lump sum nature, then it is likely that managers will invest sooner than would be optimal from a shareholder's perspective. Conversely, if the project is a flow then it is likely that the manager will invest later than is optimal from a shareholder's perspective. Thus constraining managers to hold large amounts of their company's stock can actually have adverse impacts on shareholder wealth. That is, giving managers stock has the potential to make them think *less* like shareholders.

2.5 Effort

Managers are generally given stock because they are believed to have the ability to positively influence the value of the firm. Work relating to the optimal incentive level that shareholders should give managers to motivate them to exert effort has shown that the optimal incentive level decreases with firm risk.³⁹ These models generally focus on a manager who can directly affect the stock price of the firm he works for by exerting effort, but the

³⁹See Murphy (1999) for a survey of this literature.

stock price is subject to significant noise. As the level of noise increases, the manager's optimal level of effort decreases because the firm's stock price performance depends less on his level of effort. It is therefore not surprising that as the level of noise increases, the optimal incentive decreases in this context. As will become clear later, a key distinction between my thesis and this work is that there is no "noise" in my model. That is, the manager's effort is fully reflected in the shareholder's payoff.

The majority of this work however does not make any attempt to separate the impact of the firm's systematic and idiosyncratic risk on the optimal incentive level. The exception to this is a paper by Jin (2002), who examines the optimal incentive given to the manager when he cannot trade the stock of his own firm but can trade the market portfolio. He finds that idiosyncratic risk decreases the optimal incentive for the same reasons as in the standard model, while systematic risk has no effect on the optimal incentive level. The reason that systematic risk has no effect is that the manager can simply adjust his holding of the market portfolio to eliminate any changes in the systematic component of the firm's risk. This analysis however doesn't take into account how changes in the level of both systematic and idiosyncratic risk can affect the manager's valuation of the firm. This is likely to affect the manager's optimal level of effort, and thus the contract provided to the manager. In this thesis I will model a slightly different problem that incorporates the manager's valuation problem as well as his effort decision. I do this by allowing effort to affect the investment cost of the firm's project. Once the manager's optimal level of effort is determined, we can then model the shareholder's problem of how much stock to give the manager.

A related problem that has not been discussed is what impact effort has in a real options setting. The real options papers that have examined the difference between shareholders and managers and how to contract to correct for these differences have generally focused on inducing optimal timing from a manager who has private information. Wonder (2006) presents a model

of private information where the manager can divert the firm's assets to his own use and shows that the optimal contract is a call option on the project payoff. Grenadier and Wang (2005) take a different approach and allow the manager to exert effort to alter the probability distribution of the private portion of the payoff. In their model effort alters the chance of getting a "good" project instead of a "bad" project. In this setting, since the manager also receives private information by waiting, his option is more valuable than the shareholders' and thus the manager will delay investment more than a shareholder would. Finally, Hugonnier and Morellec (2007a) present a model where the manager is undiversified and thus values the project differently to the shareholders but faces the possibility of a control challenge if his behaviour deviates too far from value maximising behaviour. They find that risk aversion and idiosyncratic risk significantly speed up investment, but this is mitigated by the threat of a control challenge. This model however has a very ad-hoc link between the manager's wealth and the project in that his wealth is simply scaled up or down by a constant at the time of investment. This constant depends on whether he is replaced and if the project has a negative NPV.

Chapter 3

General Setup

3.1 Introduction

The purpose of this chapter is to set out the modelling framework we will be using in the rest of this thesis. This serves the purpose of consolidating any common assumptions in a single point of reference. We begin this chapter in Section 3.3 by setting out the principal agent problem between the manager and shareholder. Following this we describe the specific assumptions made concerning the shareholder's problem in Section 3.3 and the manager's problem (including possible utility functions) in Section 3.4. We next set out the assumed stochastic structures for the market asset and the firm's cashflow in Section 3.5. The final part of the framework we need to lay out is the manager's inter-temporal wealth equations, which we do in Section 3.6. These equations set out the dynamics of the manager's financial wealth before and after investment.

The next step is to combine these different assumptions to derive the manager's general Hamilton-Jacobi-Bellman (HJB) equation. This is done in Section 3.7 and serves as the basis for the various models we will solve throughout

out this thesis.

The final section of this chapter (Section 3.8) sets out a special case of the general model, corresponding to the manager's value function when there is no project/cashflow. This is the well known model from Merton (1969). The solution to this model is important as it represents the manager's outside option, i.e. the payoff he gets from investing/not accepting the job.¹

3.2 The Principal-Agent framework

The principal in this thesis is the representative shareholder.² The shareholder owns the right to a project which delivers a stochastic revenue stream Y_t upon paying the investment cost I .³ The shareholder can either manage the project himself or delegate the decision on whether or not/when to invest to a manager (the agent). If the shareholder chooses to hire a manager, the manager receives a proportion α of the project which he cannot trade. The cost of this project is funded by the firm's shareholders in proportion to their ownership stake⁴ and thus the manager must pay a cost of αI as soon as the firm invests.

The potential for agency problems in this thesis arises by making the following key assumptions:

- The manager can exert effort to reduce the investment cost I , but the shareholder cannot;

¹Managers of course have the option to accept another job. For the purposes of tractability we assume the manager has no outside employment opportunity.

²As set out in Section 3.3, for simplicity we assume that there is one representative shareholder. That is, shareholder unanimity holds.

³It is assumed that upon investment the project is completed immediately, i.e there is no "Time to Build".

⁴That is, the firm is entirely equity financed.

- The shareholder incurs a monitoring cost κ while waiting to invest but the manager does not; and
- The shareholder values the firm's cash flows using the Capital Asset Pricing Model (CAPM) whilst the manager does not since he is not properly diversified.

No skill and proper diversification provide a simple point of reference for the manager's position relative to shareholders. We could have allowed shareholders to have some skill and be relatively un-diversified, but this would increase the complexity of the analysis for no discernible benefit.

Similarly, we assume that the shareholder incurs a positive monitoring cost while the manager's monitoring cost is normalised to zero. The manager could have been given a positive monitoring cost that is less than the shareholder's, but the same analysis of relative monitoring costs can be achieved by simply normalizing the manager's monitoring cost to zero.

To focus on the agency issues caused by differences in diversification, ability to exert effort and monitoring costs, we assume that there is no asymmetric information. Our analysis thus differs from the traditional principal agent literature in that the manager does not have any private information. This means that the shareholder is able to observe (amongst other things) the manager's effort, skill and financial wealth.⁵

3.3 The Shareholder's problem

In this thesis, we model the shareholder's problem in two different general ways, based upon the complexity of the manager's problem that underlies the shareholder's decision. In Chapter 6 we consider the shareholder's decision

⁵In a way this is a weakness of the framework, but on the same note it allows the intertemporal issues we are interested in to be isolated.

when the manager makes a now-or-never investment decision. In this setting there is a closed form solution for the manager's problem and thus we can allow the shareholder to choose how much of the firm's stock to give the manager. Specifically, the shareholder chooses the *level* of the variable α which is the percentage of the firm's stock that is given to the manager. In this setting, choosing $\alpha = 0$ corresponds to not hiring a manager and the shareholder managing the project.

In the model of Chapter 8 things are more complicated. The manager makes a decision in a dynamic setting where investment can be delayed and this problem must be solved numerically.⁶ This means that the proportion of the firm that the manager owns must be treated as exogenous from the shareholder's perspective. Therefore in Chapter 8 the shareholder's decision is a one-off decision of whether or not to hire, given an exogenous "cost" of hiring the manager, where the cost is the exogenous portion of the firm that is given to the manager if hired. In addition, we also allow the shareholder to delay his decision of whether or not to hire a manager/invest. Thus the shareholder can wait for a manager with desirable characteristics.

3.4 The Manager's problem

3.4.1 General setup

The manager's objective is to maximize expected utility of lifetime consumption, defined as

$$E \left[\int_0^\infty e^{-\beta dt} U(C_t) dt \right]$$

where β is the manager's time preference. The manager owns a non-traded exogenous fraction α of a firm whose investment decisions he controls.⁷ The

⁶This analysis is carried out in Chapter 7.

⁷In reality, managers would generally also have a fixed component to their income. However, we are focused on the incentive effect of ownership and thus have normalised

manager also has an exogenous initial endowment of wealth W_t and can invest in risk free bonds and a risky market asset. Because the manager cannot trade his shares in the firm he cannot trade away the firm's idiosyncratic risk. Therefore we are in a situation of incomplete markets. Prior to investment the manager thus makes the following decisions at every point in time: the dollar value of investment in the risky asset (π_t), consumption (C_t), investment in the risk free bonds ($W_t - \pi_t$), how much effort (e) to exert if investment occurs and whether or not the firm will invest in the project. Because the investment decision is assumed to be irreversible, the manager's only decisions post investment are the asset allocation and consumption/savings decisions.

Given the above structure, we can think of the manager as being in one of two states at any point in time: in State 2 the firm has invested and he is receiving his share of the income stream and in State 1 the firm is yet to invest. State 1 can be characterised as either a dynamic setting where investment can be deferred or a static setting where the manager must make a now or never investment decision.

3.4.2 The Manager's Utility Function

Constant Absolute Risk Aversion (CARA)

$$U(C) = -\frac{1}{\gamma}e^{-\gamma C} \quad (3.1)$$

The CARA utility function (3.1) is characterized by increasing relative risk aversion. This suggests that people become more averse to risks involving a proportion of their wealth as their wealth increases. Evidence presented by Campbell and Viceira (2002), suggests that this may not be the case. They argue that the long run behavior of the economy suggests that there is no persuasive link between relative risk aversion and wealth. This is based upon the fact that risk premia and interest rates (which are essentially the prices

this to zero.

of relative risk) have remained relatively constant over the last two centuries despite massive increases in per-capita consumption and wealth.

Despite the aforementioned intuitive shortcomings, this utility function does have some empirical support. Bliss and Panigirtzoglou (2004) have shown that observed option prices and forecasts show support for investors having a CARA utility function. From a modelling perspective this utility function seems to be the most tractable based upon previous work.⁸ This stems from the fact that CARA utility eliminates wealth effects, which reduces the dimensionality of problems, making them easier to solve.

Constant Relative Risk Aversion (CRRA)

$$U(C) = C^\gamma \tag{3.2}$$

As previously pointed out, the arguments of Campbell and Viceira (2002) support a utility function with constant relative risk aversion and thus this utility function is appealing in that respect. However, using a constant relative risk aversion utility function allows for wealth effects. This has an intuitive appeal as it is possible that an agent's wealth level will have an impact on his investment decisions. However, the added dimensionality of the problem may make the analysis intractable.

CARA utility also exhibits decreasing absolute risk aversion which is an intuitively plausible description of behavior. Therefore in selecting a utility function we face a trade off between tractability (CARA) and economic plausibility (CRRA).

⁸Closed form solutions to incomplete markets based investment timing problems have been found by Miao and Wang (2007) and Henderson (2007).

3.5 Stochastic Structure

The following assumptions are made concerning the stochastic structure of the investment opportunity set.

3.5.1 Market Asset

The value of the market asset, M_t , evolves according to the following Geometric Brownian Motion (GBM)

$$dM_t = \mu_m M_t dt + \sigma_m M_t d\xi_t \quad (3.3)$$

GBM is the generally preferred specification for stock prices since it doesn't permit negative values and thus captures the limited liability nature of traded companies. It is also more tractable than other potentially more realistic models such as those that assume volatility is stochastic.

3.5.2 Investment State Variable

Because the assumption regarding the evolution of the investment state variable has a significant effect on the complexity of the problem, we will consider the following processes for Y_t (i.e. the project's cashflow).

Simple Brownian Motion (SBM)

$$dY_t = \mu_y dt + \rho \sigma_y d\xi_t + \phi d\eta_t \quad (3.4)$$

Note that $d\eta_t$ is an additional Brownian motion. Here the state variable can become negative, which in the case of a stochastic income stream is interpreted as the project making a loss. While this specification simplifies the analysis substantially, the state variable can become arbitrarily large and

negative so that we need to assume that the manager will continue the project even if it is incurring substantial losses.⁹

Geometric Brownian Motion (GBM)

$$dY_t = \mu_y Y_t dt + \rho \sigma_y Y_t d\xi_t + \phi Y_t d\eta_t \quad (3.5)$$

The assumption of GBM means that the state variable cannot become negative. Furthermore, since changes to the state variable are multiplicative, once the state variable reaches zero it stays there forever. This implies that once the cashflow or project value reaches zero the project rights will be worthless.

Variables related to idiosyncratic risk

Given that we are interested in separating the effects of systematic and non systematic risk, the following variables are defined for both the GBM and SBM cases:

- $-1 \leq \rho \leq 1$: the correlation of the firm's cash flow with the market portfolio
- $\phi = \sqrt{1 - \rho^2} \sigma_y$: the idiosyncratic component of the project's volatility
- $d\eta_t$: the Wiener process governing the idiosyncratic component of the project cash flow.
- $d\xi_t$: the Wiener process governing the systematic component of the project cash flow.
- $(d\xi_t)(d\eta_t) = 0$

⁹This problem could be solved by introducing an abandonment option. This however would significantly complicate the analysis.

3.6 Inter-temporal Wealth Equations

Given the characterization of the manager's problem and the stochastic structure in which the problem is framed, we can express the manager's pre and post-investment intertemporal wealth equations as follows.

3.6.1 Stage 2 : Post Investment

In this state the manager has already incurred the investment cost αI and is currently receiving his share of the stochastic income stream αY_t each period. The budget constraint is therefore

$$dW_t = ((rW_t + \pi(\mu_m - r)) - C_t + \alpha Y_t)dt + \pi\sigma_m d\xi_t \quad (3.6)$$

Intuitively, the manager's wealth increases each date due to his investments in the risk free asset, risky asset and the project. Each date it also decreases because he spends C_t on consumption.

3.6.2 Stage 1 : Pre Investment

To allow for the situation where the manager may wish to invest in the project but does not have the money to do so, we do not constrain the manager to hold his share of the investment cost at every point in time prior to investment.¹⁰ Therefore the pre-investment intertemporal wealth equation is simply the post-investment equation without the stochastic cashflow

$$dW_t = ((rW_t + \pi(\mu_m - r)) - C_t)dt + \pi\sigma_m d\xi_t \quad (3.7)$$

¹⁰Put another way, the manager is not liquidity constrained. If we did constrain the manager to hold αI at every point in time prior to investment, this would be accomplished by constraining the manager to hold αI in the risk free asset at every point in time prior to investment.

3.7 The General Hamilton-Jacobi-Bellman Equation

While the specific form of the manager's value function depends on the assumptions concerning the investment state variable, the form of the investment payoff, the manager's utility function and whether or not investment has occurred, in order to avoid repetition it is useful to derive the general form of the manager's value function. Given that the manager maximizes the expected lifetime utility of consumption over an infinite horizon, we can use dynamic programming to express the manager's value function as follows

$$J(W, Y, t) = \max_{C, \pi} [U(C_t)dt + e^{-\beta dt} E[J(W(t+dt), Y(t+dt), t+dt)]] \quad (3.8)$$

Using standard methods it is straightforward to show that the Hamilton-Jacobi-Bellman (HJB) for the manager's problem is

$$\beta J(W, Y) = \max_{C, \pi} \left[U(C) + E[J_t + J_w dW + J_y dY + \frac{1}{2} J_{ww} dW^2 + \frac{1}{2} J_{yy} dY^2 + J_{wy} dY dW] \right] \quad (3.9)$$

Note that the exact form this equation takes depends on whether dW takes the form specified in Equation (3.7) or Equation (3.6).

3.8 Special Case: The Model with No Project (Merton (1969))

In the case where the manager does not own any of the firm or if the project rights have expired, then lifetime utility no longer depends on the investment state variable. Therefore the problem reduces to the classic case studied by Merton (1969)

3.8.1 The Manager's Problem

In this situation the manager's intertemporal budget constraint is given by (3.7). Using M superscripts to denote that this is the Merton (1969) value function, we can simplify (3.9) down to:

$$\begin{aligned}
 \beta J^M(W) &= \max_{C,x} [U(C) + J_w^M (rW + x(\mu_m - r) - C) \\
 &\quad + \frac{1}{2} J_{ww}^M (x\sigma_m W)^2] \\
 FOC_C : U'[C] &= J_w^M \\
 FOC_\pi : \pi^* &= -\frac{J_w^M (\mu_m - r)}{J_{ww}^M \sigma_m^2}
 \end{aligned} \tag{3.10}$$

This is the standard system from Merton (1969). The consumption first order condition is the familiar envelope condition which states that in equilibrium the marginal utility of consumption must be equal to the marginal utility of deferring a unit of wealth. The portfolio first order condition is simply a continuous time analogue to standard mean variance portfolio theory in that the dollar amount invested into the risky asset is positively related to the risk premium of the asset $(\mu_m - r)$ and negatively related to the variance of the asset (σ_m^2)

3.8.2 Solutions

Because the solution for the manager's intertemporal value function is dependent on the form of his utility function, we will present the solutions for both CRRA and CARA utility. However, as the solution method for this problem is well known from previous work, the derivations will not be repeated here.

CARA Utility

When utility is assumed to have the functional form given in (3.1), Merton (1969) has shown that the solution to (3.10) is

$$J^M(W) = -\frac{1}{\gamma r} e^{-\gamma r(W + \frac{\Phi^2}{2\gamma r^2})}$$

where Φ is the Sharpe ratio of the market asset.¹¹ Thus optimal consumption and investment can be expressed as

$$C^* = r \left(W + \frac{\Phi^2}{2r^2\gamma} \right)$$

$$\pi^* = \frac{\mu_m - r}{r\gamma\sigma_m^2}$$

CRRA Utility

When utility is assumed to have the functional form given in (3.2), Merton (1969) has shown that the solution to (3.10) is

$$J^M(W) = A^{\gamma-1} W^\gamma$$

where

$$A \equiv \frac{\gamma}{1-\gamma} \left[\frac{\beta}{\gamma} - r - \frac{1}{2} \left(\frac{\mu_m - r}{\sigma_m} \right)^2 \frac{1}{1-\gamma} \right]$$

Therefore optimal consumption and investment in the risky asset can be expressed as

$$C^* = AW$$

$$\pi^* = \frac{\mu_m - r}{(1-\gamma)\sigma_m^2} W$$

¹¹Thus $\Phi = \frac{\mu_m - r}{\sigma_m}$

Part II

Clean path

Chapter 4

CARA Utility with SBM Cash flow

4.1 Introduction

As outlined in Chapter 1, there are a number of steps that must be completed before we answer the primary questions this thesis seeks to answer. In this chapter we will complete the first step. The question this chapter examines is “how does a manager value a cash flow he is constrained to own a proportion of?”. Having established the manager’s valuation of his share of the cashflow, we also briefly examine the manager’s static (“now-or-never”) investment decision.

Previous work has shown that the assumption of CARA utility causes the manager’s valuation of the income stream to be independent of his level of wealth,¹ this simplifies the analysis considerably and thus CARA utility will be the starting point for our analysis. Given that Miao and Wang (2007) have obtained a closed form solution for a CARA investor that owns a SBM cash-flow, we will use their solution procedure for the case of partial ownership. This dual assumption of CARA/SBM will continue through to Chapter 8

¹See Miao and Wang (2007)

and is the basis of the “clean path”.

The layout of this chapter is as follows: Section 4.2 sets out the solution (4.2.1) and comparative statics (4.2.2) for the manager’s utility valuation of the cash flow, Section 4.3 sets out the manager’s now-or-never investment decision and Section 4.4 summarises the findings of this chapter.

4.2 Utility Valuation of the Cashflow

4.2.1 Solution

In this state the manager is already receiving an income stream that evolves according to Simple Brownian Motion and faces the Stage 2 intertemporal budget constraint. Using (3.4), (3.1) and (3.6) we can simplify the general HJB equation (3.9) down to the following

$$\begin{aligned} \beta J^2(W, Y) = \max_{C, \pi} & \left[-\frac{1}{\gamma} e^{-\gamma C} + J_w^2 (rW + \pi (\mu_m - r) - C + \alpha Y) + J_y^2 \mu_y \right. \\ & \left. + \frac{1}{2} J_{ww}^2 (\pi \sigma_m)^2 + \frac{1}{2} J_{yy}^2 (\phi^2 + \rho^2 \sigma_y^2) + J_{wy}^2 (\pi \rho \sigma_y \sigma_m) \right] \end{aligned}$$

Taking the first order conditions for consumption and investment in the risky asset yields

$$\begin{aligned} FOC_C : C^* &= -\frac{\ln(J_w^2)}{\gamma} \\ FOC_\pi : \pi^* &= -\frac{J_w^2 (\mu_m - r)}{J_{ww}^2 \sigma_m^2} - \frac{\rho J_{wy}^2 \sigma_y}{J_{ww}^2 \sigma_m} \end{aligned}$$

Following Miao and Wang (2007) we assume² that $\beta = r$ and guess a solution for the value function of the following form

$$J^2(W, Y) = -\frac{1}{\gamma r} e^{-\gamma r (W + \alpha G[Y] + \frac{\Phi^2}{2\gamma r^2})} \quad (4.1)$$

²Miao and Wang (2007) note that this assumption is necessary to obtain a closed form solution.

which is simply the Merton (1969) no income value function plus the unknown function $\alpha G[Y]$. By setting up the problem in this way, $\alpha G[Y]$ is the implicit valuation of the manager's stake in the firm. Conversely, $G[Y]$ can be interpreted as the manager's implicit valuation of the entire firm given that he is constrained to own a proportion α of the firm. Using this guess and substituting in the first order conditions allows the HJB equation to be simplified down to

$$Y - rG(Y) + \frac{1}{2}r\alpha\gamma(\rho^2 - 1)G'(Y)^2\sigma_y^2 + \left(\mu_y + \frac{\rho(r - \mu_m)\sigma_y}{\sigma_m}\right)G'(Y) + \frac{1}{2}G''(Y)\sigma_y^2 = 0$$

Following Miao and Wang, we make an initial guess for $G[Y]$ of

$$G[Y] = \left(\frac{Y}{r} + A\frac{\mu_y - \rho\sigma_y\Phi}{r^2}\right) - A\frac{\gamma\phi^2}{2r^2}$$

where A is a constant to be solved for. If we substitute this guess into the differential equation and solve for A we get the following

$$A = \frac{2\rho(r - \mu_m)\sigma_y + \sigma_m(\alpha\gamma(\rho^2 - 1)\sigma_y^2 + 2\mu_y)}{2\rho(r - \mu_m)\sigma_y + \sigma_m(\gamma(\rho^2 - 1)\sigma_y^2 + 2\mu_y)}$$

Substituting this back into our initial guess for $G[Y]$ yields the following solution

$$G[Y] = \left(\frac{Y}{r} + \frac{\mu_y - \rho\sigma_y\Phi}{r^2}\right) - \alpha\frac{\gamma\phi^2}{2r^2} \quad (4.2)$$

This is identical to the solution obtained by Miao and Wang (2007) except that the third term is now scaled by α . Given that the third term deals with the idiosyncratic risk of the project and all the variables in that term are positive constants we can see that the manager's subjective valuation of the *entire* project is decreasing in α . The intuition behind this is that the more of the firm the manager is constrained to hold, the more idiosyncratic risk he is forced to bear as he is less diversified. Therefore the manager will effectively

value the cash flow using a higher discount rate. From Equation (4.2) we can also make the observation that as $\alpha \rightarrow 0$ the valuation of the project approaches the CAPM valuation and thus does not depend on risk aversion (γ) or idiosyncratic risk (ϕ).

4.2.2 Comparative statics

We now conduct comparative statics to determine how the manager's valuation of cashflow varies with particular parameters. The main purpose of this is to see whether the effect of certain parameters differs from the standard effects predicted by the CAPM.³ The base case parameters we will use are as follows:

Table 4.1: *Base case parameters for Chapter 4*

$$\begin{array}{lll} \gamma = 1 & r = 0.1 & \mu_m = 0.15 \\ \beta = r & \sigma_m = 0.3 & \mu_y = 0.1 \\ \sigma_y = 0.3 & \rho = 0.0 & I = 100 \end{array}$$

These parameters differ slightly from those used by Miao and Wang (2007). Where the parameters are different, this has been done to maintain some level of consistency with the parameters used in the numerical analysis of Chapter 9.⁴

Given that the volatility of the project (σ_y) and the correlation of the project (ρ) appear in both the second and third terms⁵ of Equation 4.2, their effect

³As discussed in Chapter 3, the CAPM is the benchmark valuation we use for the shareholder.

⁴Loosely speaking, the market parameters from Chapter 9 (which are taken from Munk (2000)), are used in conjunction with the manager-specific parameters from Miao and Wang (2007). Because different processes are used for the cashflow (GBM and SBM) we cannot have perfectly consistent parameters between Chapter 9 and the rest of the thesis.

⁵Recall that $\phi \equiv \sqrt{1 - \rho^2} \sigma_y$

on the value of the project will depend crucially on the level of α . To help understand this, Figure 4.1 plots the manager's subjective value of the project as a function of both α and ρ .

Figure 4.1: *Subjective Value as a function of α and ρ*

Calculated using the base case parameters

From Figure 4.1 we can see that the effect of ρ on $G[Y]$ depends on whether α is high or low. When α is low, $G[Y]$ decreases linearly in ρ whereas when α is large $G[Y]$ initially decreases in ρ but then begins to increase. The intuition here is that when α is small the manager's valuation of the project approaches the CAPM valuation. Therefore the value of the project decreases linearly with ρ since this increases the project's "beta".⁶ On the other hand, when α is large the manager is very exposed to the firm and thus cares more about the firm's idiosyncratic risk. To understand why the valuation is non-linear in ρ when α is large, recall that $\phi = \sqrt{1 - \rho^2} \sigma_y$. Therefore ρ appears in the third term of equation (4.2) in a nonlinear fashion. It is easy to show that the third term of this equation is increasing in ρ until $\rho = 0$ after which this

⁶In the CAPM, beta is the measure of systematic risk. This should not be confused with the parameter β in our model representing the agent's time preference.

term is decreasing in ρ . It is also interesting to note that the value of the project decreases with respect to α except when $|\rho| = 1$. This is because in this situation we are in a complete market and thus the project has no idiosyncratic risk.

Similarly, to understand how the effect of σ_y on the subjective valuation depends on α , Figure 4.2 plots the manager's subjective valuation as a function of σ_y and α when $\rho = 0.5$. From Figure 4.2 we can see that the effect of σ_y also depends strongly on how much of the firm the manager owns. We can see that when α is small, σ_y has an insignificant effect on $G[Y]$. However when α is large the effect of σ_y is large and non-linear.⁷ This is because when α is small the manager does not care about the project's idiosyncratic risk and thus the only effect of σ_y on the manager's valuation of the project is through its effect on the project's beta. When α is large the manager also cares about the project's idiosyncratic risk and thus σ_y has a significant effect on the manager's subjective valuation. As shown by the third term in Equation (4.2), idiosyncratic risk has a non-linear effect on project value.

Figure 4.3 plots the subjective valuation as a function of α and γ . The intuition behind the effect of γ as α is varied is similar to that for σ_y and ρ . Figure 4.3 shows that when α is very small the manager's valuation of the project approaches the CAPM valuation. Given that the CAPM valuation is independent of preferences it is unsurprising that γ has little effect when α is low. The larger α is, the more idiosyncratic risk the manager is exposed to. Given that the higher risk aversion is, the more a manager cares about idiosyncratic risk, it is expected that the value of the project decreases with γ . It is also unsurprising that this relationship is stronger when α is higher since idiosyncratic risk is larger.

In the CAPM, the only effect that total project volatility has is through its

⁷Note that this graph is drawn with $\rho = 0.5$ and thus σ_y has a negative effect on $G[Y]$. If a value of ρ close to -1 was chosen the same analysis would apply except σ_y would have a positive effect.

Figure 4.2: *Subjective Value as a function of α and σ_y*

Calculated using the base case parameters and $\rho = 0.5$

Figure 4.3: *Subjective Value as a function of α and γ*

Calculated using the base case parameters

effect on the beta of the project. To understand the distinction between systematic risk (beta in the CAPM) and idiosyncratic risk (ϕ) in the current model, Figure 4.4 plots the manager's subjective valuation of the project as function of ρ and σ_y .

Figure 4.4: *Subjective Value as a function of ρ and σ_y*

Calculated using the base case parameters

Because investors only care about systematic risk in the CAPM, when $\rho = 0$, σ_y will have no effect on the value of the project since this risk is entirely idiosyncratic. Figure 4.4 illustrates that in the current model idiosyncratic risk matters. When $\rho = 0$, σ_y has a negative effect on the subjective valuation of the project. In Figure 4.4 as $\rho \rightarrow 1$ there is a significant negative relationship between the subjective value and total volatility, whereas when $\rho \rightarrow -1$ the relationship is positive. This is the standard result from the CAPM (the “CAPM effect”). However in the current setting the manager requires compensation for both systematic and idiosyncratic risk. Therefore the net effect of σ_y on the subjective valuation of the project will depend on the relative split between idiosyncratic and systematic risk.

The idiosyncratic risk of a project is by definition independent of movements in the market asset and thus the *effect* of idiosyncratic risk on the valuation of the project is independent of ρ and thus always negative. However, in our modelling framework the *level* of idiosyncratic risk depends on ρ , because ρ determines how much of the project's risk can be hedged using the risky asset. As $\rho \rightarrow |1|$, the project has no idiosyncratic risk and thus the CAPM effect of σ_y dominates the idiosyncratic effect, whereas as $\rho \rightarrow 0$ the project's risk is entirely idiosyncratic and the idiosyncratic effect dominates.

4.3 The Static Investment Problem

If investment in the project can be characterized as a now or never decision, then we can think of the manager as choosing between not investing (i.e. receiving the Merton (1969) value function) or paying αI and receiving the Stage 2 value function. The manager's value function can thus be expressed as

$$\begin{aligned} V(W, Y) &= \max[J^2(W - \alpha I, Y), J^M(W)] \\ &= \max\left[-\frac{1}{\gamma r} e^{-\gamma r(W + \alpha(G[Y] - I) + \frac{\Phi^2}{2\gamma r^2})}, -\frac{1}{\gamma r} e^{-\gamma r(W + \frac{\Phi^2}{2\gamma r^2})}\right] \end{aligned} \quad (4.3)$$

We use this function to characterize the situations when a manager does and does not invest for a range of parameter values. Given that we have already examined the valuation the manager places on the cashflow, we know the effect that certain parameters will have on the manager's decision of whether or not to invest. Anything that reduces the value of the project makes investment less attractive and vice versa. However, whether or not these effects are significant to the investment decision is the more important question. The main variables of interest are those which are manage-specific and thus we will examine the significance of wealth, ownership and risk aversion on the investment threshold.

Figure 4.5: *Investment Threshold (Y^*) as Function of α*

Calculated using the base case parameters and $\gamma = 10, 5, 2, 1, 0.5$. In this graph the y-axis represents Y and the x-axis represents α . The top line plots the threshold for $\gamma = 10$ while the bottom line plots the threshold for $\gamma = 0.5$.

To understand the manager's investment decision, Figure 4.5 plots the manager's investment "threshold" Y^* as a function of α for various levels of γ . The threshold Y^* is the level of the cashflow where the manager invests once $Y > Y^*$. From Figure 4.5 we can see that the effect of α on the investment threshold depends heavily on the manager's risk aversion. For context, the CAPM investment threshold for the base case parameters is 9. When risk aversion is very low we find that the investment threshold is almost flat (and approaches the CAPM threshold) whereas when risk aversion is large the investment threshold increases significantly in α . This happens because when γ is small the manager approaches risk neutrality and thus does not care about idiosyncratic risk.

Figure 4.6 plots the manager's investment threshold as a function of W and Y for various levels of γ . When we examine the effect of wealth on the investment threshold we unsurprisingly find that the investment decision is

Figure 4.6: *Investment Threshold (Y^*) as Function of W*

Calculated using the base case parameters and $\gamma = 10, 5, 2, 1, 0.5$. In this graph the y-axis represents Y and the x-axis represents W . The top line plots the threshold for $\gamma = 10$ while the bottom line plots the threshold for $\gamma = 0.5$

independent of the manager's wealth level. This stems from the fact that the manager's valuation of the cash flow is wealth independent.⁸ The interesting feature of Figure 4.6 is that as risk aversion increases, the gap between the investment threshold selected by the manager and that predicted by the CAPM (9 in Figure 4.6) gets larger.

⁸It is also of course dependent on the assumption that the manager faces no liquidity constraints.

4.4 Summary

This chapter has focused on answering the question “how does a manager value a cash flow he is constrained to own a proportion of?”.

We have seen that the constrained valuation is often non-linear and the effects of certain parameters differ markedly from those predicted by the CAPM. The main result to come from this chapter is that as α increases, the manager’s subjective valuation of the project decreases since this exposes him to more idiosyncratic risk. As shown in Figure 4.5, this means that the higher α is, the *less* likely it is that the manager will invest.

Chapter 5

Managerial Effort

5.1 Introduction

In the model of Chapter 4 the only effect the manager has on the firm is choosing whether or not investment takes place. In reality managers are generally only hired because they have the *ability* to positively¹ affect the value of the firm in some way.

Given that the previous model only captures the negative effects associated with managerial ownership, we now extend the model of Chapter 4 to the case where the manager can exert effort to reduce the cost of investment, thereby increasing the firm's investment payoff.

We will thus complete the next step in our progression by understanding what factors influence the manager's decision of how much effort to exert. The impact of the manager's effort decision on his investment decision are then examined in a “now-or-never” setting.

This chapter is set out as follows: Section 5.2 outlines the framework and

¹Managers can of course also negatively affect the value of the firm. Our focus however will be on positive impact the manager can have on the firm.

assumptions used to introduce effort into the model of Chapter 4, Section 5.3 examines the solution for the manager's optimal level of effort, Section 5.4 examines the manager's investment decision in a now-or-never setting given that he can exert cost reducing effort at the time of investment and Section 5.5 summarises the results of this chapter.

5.2 Setup

We assume that the manager can affect the investment cost in the following way

$$I[e] = \exp(-\lambda e)A + (1 - \exp(-\lambda e))B \quad e \geq 0 \quad (5.1)$$

By defining the investment cost function in this way, the manager's effort is fully reflected in the investment cost and therefore the eventual payoff to the shareholder. In other words there is no "noise" in the Holmstrom (1979) sense.²

While we could have allowed for e to become negative (thus allowing the manager to shirk)³, we will restrict ourselves to the case of strictly positive effort in order to focus on the issues of managerial effort and idiosyncratic risk.

By defining the cost of investment in this way we can think of A as the base cost of investment if no effort is exerted. The parameters B and λ govern two different aspects of managerial skill. B governs the absolute amount by which the investment cost can be reduced by the manager exerting effort.⁴ We can thus think of B as characterizing both the manager's ability to reduce the investment cost and the natural scope for cost reduction in the project being considered. The aspect of managerial skill that λ represents is subtly different

²To introduce noise into this setup one would include a random error term in the investment cost function.

³Although arguably $e = 0$ could be characterised as shirking.

⁴Thus if $A = \$100m$ and $B = \$80m$ then there is the potential for $\$20m$ of cost savings.

from B . While B determines how much the investment cost can be reduced through effort, λ tells us how much effort is required to reach an arbitrary level of cost reduction. Thus all other things being equal a higher λ means that effort is more effective at reducing the cost of investment. Figure 5.1 illustrates the impact of λ on the effectiveness of effort. We can see that when λ is very large it takes almost no effort to reduce the investment cost to B , whereas when λ is quite small it takes a significant amount of effort to reach B .

Figure 5.1: *Investment Cost vs Effort for various levels of λ*

Calculated using $\lambda = 0.01, 0.1, 0.5, 1, 2, 3$, where the top line represents $\lambda = 0.01$ and the bottom curve is $\lambda = 3$

It is important to note that there are two extreme cases when effort has no effect on the investment cost. When $A \approx B$, either the manager doesn't have the skill or there is physically no scope for cost savings. Therefore effort has no effect on the investment cost. Similarly if $\lambda = 0$, even if there is significant scope for cost savings,⁵ the manager lacks the skill to reduce the investment cost.

The exponential nature of the cost function also imposes two intuitively

⁵That is, B is quite low relative to A .

appealing attributes. Firstly, the manager will experience diminishing returns to effort meaning that the initial units of effort he exerts will have a large effect on the investment cost whereas later units will have a smaller effect. Secondly, so long as $B \geq 0$ we avoid the unrealistic situation of a negative investment cost.

Like the previous models, we solve this problem backwards by first solving the manager's problem given that he has decided to invest and then determining whether or not investing is the utility maximizing strategy for the manager. We already know that if the manager invests he receives the payoff $J^2(W - \alpha I, Y)$. We assume that at the time of investment the manager can exert effort which has a benefit defined by Equation (5.1) and imposes a lump sum cost on the manager. We can thus express the manager's payoff from investing as

$$J^e(W, Y) = J^2(W - \alpha I[e], Y) - \theta e \quad (5.2)$$

where θ can be thought of as the manager's dislike of effort. By setting up the payoff this way we are implicitly assuming that the cost of effort is independent of wealth and thus is physical dislike of effort as opposed to a financial representation of the cost of effort.⁶

5.3 The Manager's Optimal Level of Effort

Using Equations (5.2) and (4.1) we can simplify the payoff function to the following for the SBM/CARA case

$$J^e(W, Y) = -\frac{1}{\gamma r} e^{-\gamma r(W + \alpha(G[Y] - I[e]) + \frac{\Phi^2}{2\gamma r^2})} - \theta e \quad (5.3)$$

⁶This implies that Bill Gates dislikes exerting effort just as much as we do if we have the same θ . However, as Equation (5.4) shows, his optimal level of effort will still be different from mine due to the differences in our wealth.

Substituting in the solution for $G[Y]$ and taking the first order condition for effort allows us to solve for the manager's optimal level of effort

$$\hat{e} = \frac{\Omega - 2r \log \left(\frac{\theta}{(A-B)\alpha\lambda} \right) + \frac{2rW}{\lambda} \left(e^{\frac{\eta^2}{2r} - \frac{\alpha\gamma\rho\sigma_y\eta}{r} - \frac{\alpha^2\gamma^2\sigma_y^2}{2r} + \frac{\alpha^2\gamma^2\rho^2\sigma_y^2}{2r} + rW\gamma - Br\alpha\gamma + Y\alpha\gamma + \frac{\alpha\gamma\mu}{r} r\gamma\theta \right)}{2r\lambda} \quad (5.4)$$

where

$$\Omega \equiv 2(B\alpha - W)\gamma r^2 - \eta^2 + 2\alpha\gamma(\delta\eta\rho - \mu - rY) + \alpha^2\gamma^2\delta^2(1 - \rho^2)$$

Given the complexity of this expression, we will conduct graphical analysis to determine the effect of the firm and manager-specific parameters on the level of effort exerted by the manager. We use the same base case parameters as outlined in Section 4.2.2, with the addition of “effort” specific parameters as outlined in Table 5.1

Table 5.1: Base case parameters for Chapter 5

$$\begin{array}{llllllll} \gamma = 1 & r = 0.1 & \mu_m = 0.15 & \mu_y = 0.1 & A = 100 & \lambda = 1 & W = 1 \\ \beta = r & \sigma_m = 0.3 & \sigma_y = 0.3 & \rho = 0.0 & B = 80 & \theta = 0.1 & Y = 10 \end{array}$$

The first parameters we examine are the manager's wealth (W) and the level of the cash flow (Y). Figure 5.2 plots the manager's optimal level of effort as a function of W and Y . It shows optimal effort ($\hat{e}$) is a decreasing function of both the manager's outside wealth and the level of the cash flow. This is easily understood intuitively given that CARA utility exhibits diminishing marginal utility from wealth. Because increasing effort increases the payoff from investment, the utility gain from that payoff is large when W is small and small when W is large. Therefore we observe that effort decreases with W . Similarly when Y is very large, the utility benefit from exerting effort is small and thus we also see that effort decreases with Y .

The next parameter we wish to examine is α . As it turns out, the effect of α depends heavily on the manager's level of risk aversion (γ). Figure 5.3 therefore plots optimal effort as a function of α and γ .

Figure 5.2: *The effect of W and Y on optimal effort*

Calculated using the base case parameters

Figure 5.3 shows that the effect of α on optimal effort is highly non-linear and depends very strongly on the manager's risk aversion (γ). To understand the effect of α on optimal effort we need to recognize that increasing α has three conflicting effects on optimal effort:

- The manager obtains a greater direct benefit from effort since he owns more of the firm and thus exerts more effort;
- The manager is "richer" and thus exerts less effort; and
- The manager is less diversified and thus exerts more effort to offset his lower valuation of the cashflow.

Starting from the observation that the manager bears the full cost of any effort he exerts, but has to share the benefit of his effort with shareholders, we can see that increasing α increases the benefit the manager receives from exerting effort. Therefore α has a positive effect on effort. On the other hand,

Figure 5.3: *The effect of α and γ on optimal effort*

Calculated using the base case parameters

increasing α also scales up the manager's exposure to Y . We have already demonstrated that effort decreases with Y and thus it is unsurprising that scaling up the manager's exposure to Y decreases effort. In a sense the more of the firm the manager owns, the "richer" he is and thus his marginal utility of an increase in Y exposure is lower. The third effect relates to how diversified the manager is and is highly dependent on how risk averse the manager is. As discussed in Chapter 4, a higher value for α decreases the manager's valuation of the cashflow. To offset this decrease in value the manager exerts more effort.

We can see from Figure 5.3 that when γ is small the first two effects dominate; effort initially increases due to the first effect but as α gets large the marginal utility effect starts to kick in and thus effort tapers off and eventually decreases. However when γ is quite large we find that the diversification effect begins to dominate the first two and effort will actually start to increase in α after initially increasing then decreasing. This is interesting as the relationship between optimal effort and α is highly non-linear and heavily

influenced by the manager's risk aversion.

Another interesting aspect of Figure 5.3 is the effect of risk aversion on effort. For low levels of α effort decreases with γ . It is easy to show in the standard Merton (1969) model that the marginal utility of wealth is decreasing in γ . Thus an increase in γ decreases the utility payoff from effort and thus decreases optimal effort. However as mentioned previously, risk aversion is a significant determinant of whether under-diversification affects the manager's valuation of the cashflow. Thus we can see that when α is large and the manager is undiversified, optimal effort starts to increase because an increase in γ starts to significantly reduce the manager's valuation of the project.

To confirm our intuition that it is a lack of diversification driving the behaviour in the region where γ is high, we consider the case where $\rho = 1$. In this situation the project has no idiosyncratic risk and thus diversification is not a factor. If idiosyncratic risk is the cause of the change in behaviour in the region where γ is large, then when there is no idiosyncratic risk the diminishing marginal utility effect should dominate. That is, when γ is large effort should be decreasing in γ and α . We can see from Figure 5.4 that this is the case.

While θ and λ can respectively be thought of as the cost and benefit of effort, Figure 5.5 shows they do not have the opposite effects on effort as one might expect. As would be expected, optimal effort decreases as the cost of effort (θ) rises. As λ approaches zero optimal effort is zero as effort has no effect on the investment cost. For very small levels of λ we do find that optimal effort increases with λ , but this is only for a very small range. Outside of this range optimal effort actually decreases with λ . The intuition for this is that when λ is very small, an increase in λ means that effort actually has an effect on the investment cost. This means that effort is worthwhile and thus optimal effort increases. However, once λ becomes large enough, optimal effort starts to decrease in λ . This is because we are in a state where only a small amount

Figure 5.4: *The effect of α and γ on optimal effort when $\rho = 1$*

Calculated using the base case parameters

of effort is required to drive the investment cost close to B .⁷ Once we are in this state, the cost of driving the investment cost down to B is small enough that the manager will always exert enough effort to do so. Therefore once we are in this state optimal effort is decreasing in λ because an increase in λ decreases the effort required to drive the investment cost down to B .

If we now turn to the parameters concerning the investment cost, it is unsurprising to find that the upper (A) and lower boundaries (B) have opposite effects on optimal effort. Figure 5.6 shows that as A increases, the benefit from exerting effort increases and thus optimal effort increases. This effect begins to diminish as A becomes large due to the fact that the amount of effort required to reduce the cost to B does not change a great deal as A increases. The opposite logic applies to B where we find that when B is very small optimal effort is relatively insensitive to changes in B , but begins to rapidly decrease as B approaches A . In the case where $A \approx B$, optimal effort is driven towards zero as would be expected given effort has no effect on the

⁷Note that due to the exponential nature of the cost function, only with $e = \infty$ will $I[e] = B$.

Figure 5.5: *The effect of θ and λ on optimal effort*

Calculated using the base case parameters

investment cost.

Figure 5.6: *The effect of A and B on optimal effort*

Calculated using the base case parameters

Given the discussion in Section 4.2, it is unsurprising that in Figure 5.7 the effect of σ_y on optimal effort depends on the sign of ρ . When ρ is positive, σ_y has a negative effect on the manager's subjective valuation. Because σ_y has a negative effect on the manager's valuation the manager exerts more effort.

The opposite occurs when ρ is negative since σ_y increases the manager's valuation in this situation. We also see that as ρ increases optimal effort increases because the manager's valuation decreases as ρ increases.

Figure 5.7: *The effect of σ_y and ρ on optimal effort*

Calculated using the base case parameters

The final parameter we consider is the volatility of the market asset (σ_m). Figure 5.8 plots the manager's optimal level of effort against σ_m . Figure 5.8 shows that as σ_m increases so does the manager's optimal level of effort. This is interesting because the base case assumption of $\rho = 0$ means that σ_m has no effect on the manager's valuation of the project. Indeed, it is easy to show that regardless of the assumption concerning ρ , there is a positive relationship between optimal effort and σ_m . This occurs because a *ceteris paribus* change in σ_m reduces the market asset's Sharpe ratio. In other words the attractiveness of the manager's portfolio investments has reduced. This has the effect of reducing the manager's utility which all other things equal causes the marginal utility of wealth to increase. This increase in the marginal utility of wealth means that the utility benefit of effort has increased. Therefore we witness the optimal level of effort increasing as σ_m increases.

Figure 5.8: *The effect of σ_m on optimal effort*

Calculated using the base case parameters

5.4 The Static Investment/Effort Decision

Having analysed the manager's optimal effort decision, we are now in a position to examine the “now or never” investment decision the manager makes in a static setting. If the manager decides to invest, we have already solved for his optimal level of effort, and thus know that the payoff from investing is simply Equation (5.3) with the expression for the optimal level of effort (Equation (5.4)) substituted in. Conversely if the manager decides not to invest then he is simply in the Merton (1969) world where he is not receiving the cash flow nor does he have the option to invest. Given this characterization, the manager's investment problem can be expressed as

$$V^e(W, Y) = \max[J^2(W - \alpha I[\hat{e}], Y) - \theta \hat{e}, J^M(W)] \quad (5.5)$$

Because α and W both affect the manager's optimal level of effort ($\hat{e}$), their effects on the investment threshold are considerably different compared to

the no effort case where their effects were linear.

Figure 5.9: *Investment Threshold (Y^*) as a Function of α and γ*

Calculated using $\gamma = 5, 2, 1, 0.5$. In this graph the y-axis represents Y and the x-axis represents α . The top line plots the threshold for $\gamma = 10$ while the bottom line plots the threshold for $\gamma = 0.5$.

Figure 5.9 shows that for low levels of risk aversion the investment threshold decreases in α for the majority of the range and then gradually increases. If we recall Figure 5.3, for small levels of γ , effort is generally increasing in α because the manager benefits more from exerting effort. The fact that effort increases with α results in an investment threshold that declines as α increases. On the other hand, when γ is quite large, diversification is much more important to the manager. Therefore, for high values of γ , as α gets large the manager's valuation of the project decreases. This causes the investment threshold to increase, after initially decreasing due to the effect mentioned previously.

Figure 5.10 plots the manager's investment threshold as a function of W

Figure 5.10: *Investment Threshold (Y^*) as a Function of W and γ*

Calculated using $\gamma = 5, 2, 1, 0.5$. In this graph the y-axis represents Y and the x-axis represents W . The top line plots the threshold for $\gamma = 10$ while the bottom line plots the threshold for $\gamma = 0.5$.

and Y for various levels of risk aversion. In this graph we can see that for each level of risk aversion there is effectively a “high” threshold and “low” threshold which correspond to high wealth and low wealth. This is a direct result of using an investment cost function which has an upper (A) and lower (B) bound and the fact that effort decreases in W . I.e. the low wealth threshold roughly corresponds to the $I[e] \approx B$ and the high wealth threshold corresponds to $I[e] = A$.

If we focus on the right hand side of the graph (high W) where no effort is exerted, we observe the same general pattern as in the no effort model (e.g. Figure 4.6). That is, for higher values of γ the investment threshold is higher since the manager values the project less.

The most interesting feature of Figure 5.10 is the shape of the transition be-

tween the “high” and “low” thresholds for different levels of γ . The general pattern we observe is that for higher values of γ , the transition is steeper. Put another way, a smaller change in W is required to move from the low threshold (high effort) to the high threshold (low effort). This relationship occurs because the CARA utility function is strictly negative and thus faces an asymptote at zero. Given that an increase in γ increases the concavity of the utility function, as γ increases the level of utility approaches the asymptote much quicker (i.e. the function becomes flatter). The result of this is that for large values of γ the marginal utility of wealth and thus the utility benefit of effort get driven close to zero quite quickly. This is why in Figure 5.10 we see that for high levels of γ the change from a low investment threshold (i.e. high effort) to a high investment threshold (i.e. low effort) occurs quite suddenly. This is in stark contrast to the case of low risk aversion where effort is exerted over a much broader range of W .

5.5 Summary

In this chapter we examined the constrained manager’s decision of how much effort to exert at the time of investment. The manager’s optimal level effort is highly non-linear in the proportion of the firm he owns and depends on how risk averse he is. In other words, more is not necessarily better when it comes to motivating the manager to exert effort. Importantly, manager-specific parameters really matter.

Perhaps the more interesting feature of this chapter is the manager’s investment decision. In the model of Chapter 4, the manager’s financial wealth (W) did not affect his investment decision because under CARA utility the valuation of the cash flow is not wealth dependant. However, once effort is introduced into the model the manager’s investment decision depends on his wealth. This is because the manager’s effort decision depends on his wealth and thus so does the investment cost. Therefore, despite wealth having no

direct impact on the manager's valuation, his actions depend on his wealth and therefore the manager's wealth indirectly impacts the value of the firm.

Chapter 6

The Shareholder's Static Hiring Problem

6.1 Introduction

This chapter represents an intermediate point of analysis in this thesis. The model of Chapter 5 is not dynamic, but it does allow us to model the trade off a shareholder faces between making the manager less diversified and giving him the incentive to exert effort in a now-or-never setting. While not directly comparable to the final model in Chapter 8,¹ this serves as a useful point of comparison to the results of that chapter.

The empirical trend demonstrated in Chapter 1 suggests that there has been a “more is better” attitude towards equity linked compensation. However we have demonstrated in Chapter 4 that a large α can reduce the manager's valuation of the project. This can cause the manager to pass up investment opportunities that would be wealth increasing² from the shareholder's per-

¹See Section 3.3 for a discussion of the different approaches taken to the shareholder in this thesis.

²That is, from a shareholder's perspective the project would have a positive Net Present

spective. We also found in Chapter 5 that increasing α does not necessarily cause the manager to exert more effort.

These observations naturally lead to the question of what is the optimal level of managerial ownership? To answer this question we examine the shareholder's problem of choosing how much of the firm to give to the manager given that the manager can exert effort and is also restricted from trading his shares in the firm. Given that the proportion of the firm that the shareholder gives to the manager is likely to depend on the "type" of manager, our focus will be on manager-specific parameters.³

While our focus is on the manager, the literature surveyed in Section 2.5 shows that the optimal incentive given to managers decreases with the level of "noise" in a firm's share price. The corollary to this in our model is σ_y and thus we will also examine the affect this has on optimal compensation in our model,⁴ though we again note that this is not "noise" in the traditional (i.e. Holmstrom (1979)) sense.

We begin this chapter by setting up the shareholder's problem in Section 6.2. The difference between the manager and shareholder in this setting is that shareholders are properly diversified and value the project using the Capital Asset Pricing Model (CAPM). The manager on the other hand is not diversified and can also exert effort to reduce the cost of the project. With the shareholder's problem framed, we then solve for the level of managerial ownership (α^*) that maximises the payoff to the shareholder in Section 6.3. Section 6.4 summarises the conclusions from this chapter.

Value (NPV).

³E.g. the manager's risk aversion and "skill".

⁴Recalling again that there is no assymetric information in our model.

6.2 Setup

We assume that there is a representative shareholder who uses the CAPM to value the firm's cashflows. This means that shareholders only care about systematic risk, in contrast to the constrained manager who we have shown previously cares about both the systematic and non-systematic component of the project's risk. The CAPM valuation of a Simple Brownian Motion (SBM) cashflow is simply the manager's valuation of the cash flow (Equation (4.2)) with $\alpha = 0$

$$\frac{Y}{r} + \frac{\mu - \rho\sigma_y\Phi}{r^2}$$

We assume that the shareholder has no skill (this means that either $A = B$ or $\lambda = 0$) and thus the cost of investment if no manager is hired is simply A .⁵ Therefore the payoff to the shareholder from investing himself is

$$P^{CAPM} = \frac{Y}{r} + \frac{\mu - \rho\sigma_y\Phi}{r^2} - A \quad (6.1)$$

On the other hand if a manager is hired then the shareholder gives up α of the firm, but the investment cost is now determined by Equation (5.1) and thus depends on the manager's optimal effort level. This gives the following payoff if a manager is hired

$$P^M[\alpha] = (1 - \alpha) \left(\frac{Y}{r} + \frac{\mu - \rho\sigma_y\Phi}{r^2} - I[\hat{e}] \right) \quad (6.2)$$

Crucially, we assume that the manager's level of effort is observable.⁶ Equation (6.2) shows that if a manager is hired there is a trade off between diluting the shareholder's claim of the company, and increasing the total payoff

⁵While we could have assumed that the shareholder has some scope to reduce the investment cost, the purpose of our analysis is to examine the shareholder's decision given that the manager is more skilled than the shareholder. The simplest way to achieve this is to assume the shareholder has no skill.

⁶Equivalently, we assume A, B, λ and θ are observable.

through the manager exerting effort. Thus if a manager is hired, the shareholder will choose α^* to maximize Equation (6.2). If hiring a manager does not leave the shareholder any better off than if they did not, then they will not bother. Thus the shareholder's problem is to choose α^* , subject to the manager's participation constraint and his own "hiring constraint". We can thus formally represent the shareholder's problem as

$$\max_{\alpha} \quad (1 - \alpha) \left(\frac{Y}{r} + \frac{\mu - \rho\sigma_y\Phi}{r^2} - I[\hat{e}] \right) \quad (6.3)$$

subject to the following constraints

$$PC : \quad J^2(W - \alpha^* I[\hat{e}], Y) \geq J^M(W) \quad (6.4)$$

$$HC : \quad (1 - \alpha) \left(\frac{Y}{r} + \frac{\mu - \rho\sigma_y\Phi}{r^2} - I[\hat{e}] \right) > \frac{Y}{r} + \frac{\mu - \rho\sigma_y\Phi}{r^2} - A \quad (6.5)$$

The hiring constraint (HC) states that the shareholder will hire a manager only if this will lead to a higher payoff.⁷ The manager's participation constraint (PC) simply states that the manager will invest with a given α^* only if his utility from doing so is greater than the utility he would receive if he did not invest.

6.3 Results

By substituting the expressions for $I[\hat{e}]$, $\hat{e}$, $J^2(W - \alpha^* I[\hat{e}], Y)$ and $J^M(W)$ into Equations (6.3) - (6.5), we can solve for α^* . Given the non-linearity of the resulting constrained maximization problem, we cannot obtain a closed form solution for α^* and thus we must numerically maximise Equation (6.3) subject to the participation and hiring constraints. This is done using a numerical maximisation procedure which simply compares the payoff to the

⁷Although it is not explicitly incorporated into our model, it is generally costly to hire a manager and thus we assume that the payoff with a manager must be strictly greater than the payoff without a manager to account for this cost.

shareholder from investing himself (P^{CAPM}) to that of hiring a manager ($P^M[\alpha^*]$) for every level of α holding all other parameters constant. This is graphically demonstrated in Figure 6.1 where P^{CAPM} and $P^M[\alpha^*]$ are calculated for the base case parameters from Chapters 4 and 5.

Figure 6.1: *Calculating α^**

In this graph the dashed line represents $P^M[\alpha]$ while the solid line represents P^{CAPM} .

In Figure 6.1, α^* is simply the peak of the dashed line (i.e. the maximum of $P^M[\alpha]$) since this is greater than P^{CAPM} . Calculating α^* for these parameters yields a figure of 6.5%, which is a fairly significant proportion of the company for the manager to hold. This figure is however specific to the parameters we have chosen and thus it is important so see not only whether α^* is very sensitive to the parameters chosen, but also the specific effect that each parameter has on α^* .

The hiring and the participation constraints make understanding the result of comparative statics relatively complex. We will therefore plot three graphs for each parameter we examine. To understand whether or not the hiring constraint is binding (i.e. the shareholder is better off by hiring a manager), we plot the shareholder's payoff as a function of the variable in question with ($P^M[\alpha]$) and without (P^{CAPM}) a manager. This is panel (a) in each figure. To understand whether or not the manager's participation constraint is

binding we plot the manager's utility as a function of the variable in question when he is hired ($J^e(W, Y)$) and when he is not hired/chooses not to invest ($J^M(W)$). This is panel (b) in each figure. The third graph (panel (c)) plots α^* as a function of the parameter in question and completes the story told by the previous two graphs.

Consistent with the rest of this thesis we focus on manager-specific parameters as well as volatility.

6.3.1 Managerial Skill/Scope for Cost Savings (B)

As our focus is on manager-specific parameters, we begin examining the investment cost floor (B) which is a measure of the natural scope for cost savings as well as the manager's skill. Figure 6.2 shows three graphs which we must analyse in order to understand the effect of B (or any other parameter) on α^* .

We can see from Figure 6.2 (a) that as B increases, $P^M[\alpha^*]$ decreases because the manager's ability to reduce the investment cost is diminishing. Similarly Figure 6.2 (b) shows that the payoff to the manager decreases as B increases and thus we witness $J^e(W, Y)$ falling as well. From Figure 5.6 we know that optimal effort decreases with B . Therefore as B increases we can see in Figure 6.2 (c) that α^* increases to offset this. Also, as B approaches A we can see that we eventually get to the point where $P^M[\alpha^*] = P^{CAPM}$ and thus the "hiring" constraint starts to bind and α^* hits zero.

6.3.2 Managerial Skill (λ)

From Figure 5.5 we know that for very small levels of λ the manager exerts no effort because it is almost entirely ineffective. Therefore the hiring constraint binds and we witness $\alpha^* = 0$. For low values of λ , a small increase in λ causes a very large increase in effort and thus we witness a large jump in α^* once

Figure 6.2: *Comparative Statics For B*

 (a) Solid Line: $P^M[\alpha^*]$, Dashed Line: P^{CAPM}

 (b) Solid Line: $J^2(W - \alpha^* I[\hat{e}], Y)$, Dashed Line: $J^M(W)$

the hiring constraint no longer binds. As λ increases further, α^* falls. This occurs because λ is a measure of the manager's skill and tells us how much effort is required to drive the investment cost down to B . This means that as λ increases the manager can achieve an investment cost close to B with lower effort. The result of this is that the manager needs less incentive to exert effort and thus α^* falls. This is why we see the shareholder's payoff rising in Figure 6.3 (b) and the manager's payoff falling in Figure 6.3 (a).

Figure 6.3: *Comparative Statics For λ*

 (a) Solid Line: $P^M[\alpha^*]$, Dashed Line: Payoff^{CAPM}

 (b) Solid Line: $J^2(W - \alpha^* I[\hat{e}], Y)$, Dashed Line: $J^M(W)$

6.3.3 Risk Aversion (γ)

Figure 6.4 shows that the relationship between α^* and γ is highly non-linear and at the same time α^* does not vary much with γ . To understand the intuition behind the relationship between α^* and γ shown in Figure 6.4, it is necessary to have an understanding of Figure 5.3, which we will reproduce here as Figure 6.5 for clarity's sake.

Figure 6.4: Comparative Statics For γ

 (a) Solid Line: $P^M[\alpha^*]$, Dashed Line: P^{CAPM}

 (b) Solid Line: $J^2(W - \alpha^* I[e], Y)$, Dashed Line: $J^M(W)$

As Figure 6.5 shows, the relationship between γ and optimal effort is complicated and depends quite heavily on α . Therefore it is unsurprising that there is such an unusual relationship between γ and α^* . For the lower levels of α that we are dealing with here, an increase in γ decreases optimal effort and thus we witness that $P^M[\alpha^*]$ is decreasing in γ . For the range of α^* under consideration, we know from Figure 6.5 that optimal effort is decreasing in γ . Thus as γ increases the shareholder will wish to motivate the manager to

Figure 6.5: *The effect of α and γ on optimal effort*

exert more effort.

The effect that α has on effort depends upon risk aversion and this is the reason that we witness α^* increase, decrease and then increase in γ . This occurs because when γ is small an increase in α increases effort because the marginal utility of wealth is large (thus effort is very beneficial). Thus α^* increases to offset the decrease in effort caused by the increase in γ . As γ increases we move into the region where increasing α actually decreases optimal effort due to the fact that increasing α makes the manager “richer” and thus decreases marginal utility making effort less worthwhile. The result of this is that in this region decreasing α will increase effort and we thus witness α^* falling. As γ starts to get very large we are in the region where the manager really cares about diversification. Because an increase in α makes the manager value the project less in this region, increasing α increases optimal effort and thus α^* begins to increase with γ .

Despite the very complex effects that γ has, it is interesting to note that the absolute change in α^* does not appear to be very large over a wide range of γ . This suggests that the fact that managerial risk aversion is unobservable

in practice is not that important for the optimal contract.

6.3.4 Managerial Wealth (W)

Figure 6.6: *Comparative Statics For W*

(a) Solid Line: $P^M[\alpha^*]$, Dashed Line: P^{CAPM}

(b) Solid Line: $J^2(W - \alpha^* I[\hat{e}], Y)$, Dashed Line: $J^M(W)$

Because the manager has diminishing marginal utility, as W increases the manager's optimal level of effort decreases. We thus witness $P^M[\alpha^*]$ decreasing as W increases in panel (a) of Figure 6.6. To offset the manager's decreasing level of effort, α^* rises until $P^M[\alpha^*] = P^{CAPM}$ in Figure 6.6 (a) at which point the hiring constraint binds and α^* drops to zero.

6.3.5 Level of Cash Flow (Y)

Figure 6.7: Comparative Statics For Y

(a) Solid Line: $P^M[\alpha^*]$, Dashed Line: P^{CAPM}

(b) Solid Line: $J^2(W - \alpha^* I[\hat{e}], Y)$, Dashed Line: $J^M(W)$

Panel (a) of Figure 6.7 shows that for low values of Y we are in a situation where $P^{CAPM} < 0$ and thus the shareholder gives away the whole project ($\alpha^* = 1$). Panel (b) of Figure 6.7 shows that in this region $J^2(W - \alpha^* I[\hat{e}], Y) < J^M(W)$ and thus the manager's participation constraint is binding and he would not even invest. As Y increases the participation constraint stops binding (as shown in panel (b)) and $P^M[\alpha^*]$ becomes positive (as shown in panel (a)). Because P^{CAPM} is increasing in Y , α^* decreases as the man-

ager becomes unnecessary. Because the manager's optimal level of effort is decreasing in Y and α^* is not increasing sufficiently to offset this, we see that as Y gets large we eventually reach a situation where $P^M[\alpha^*] = P^{CAPM}$. Thus the hiring constraint binds causing α^* to drop to zero in panel (c) of Figure 6.7.

6.3.6 Cash Flow Volatility (σ_y)

The effect of σ_y on P^{CAPM} depends explicitly on ρ because this determines whether or not an increase in σ_y increases ($\rho > 0$) or decreases ($\rho < 0$) the project's beta. An increase in the project's beta decreases P^{CAPM} and thus we see in Figure 6.9 that when ρ is positive α^* is increasing in σ_y as the shareholder wishes to offset the fall in P^{CAPM} by motivating the manager to exert more effort. Figure 6.10 shows that the opposite logic applies when $\rho < 0$ and thus α^* is decreasing in σ_y . Under the CAPM investors do not price idiosyncratic risk and thus when $\rho = 0$ (and the project's risk is entirely idiosyncratic) P^{CAPM} does not depend on σ_y . Because P^{CAPM} doesn't depend on σ_y , neither does α^* .

Figure 6.8: *Comparative Statics For σ_y , $\rho = 0$*

(a) Solid Line: $P^M[\alpha^*]$, Dashed Line: P^{CAPM}

(b) Solid Line: $J^2(W - \alpha^* I[\tilde{e}], Y)$, Dashed Line: $J^M(W)$

Figure 6.9: *Comparative Statics For σ_y , $\rho = 0.5$*

(a) Solid Line: $P^M[\alpha^*]$, Dashed Line: P^{CAPM}

(b) Solid Line: $J^2(W - \alpha^* I[\tilde{e}], Y)$, Dashed Line: $J^M(W)$

Figure 6.10: *Comparative Statics For σ_y , $\rho = -0.5$*

(a) Solid Line: $P^M[\alpha^*]$, Dashed Line: Payoff^{CAPM}

(b) Solid Line: $J^2(W - \alpha^* I[\tilde{e}], Y)$, Dashed Line: $J^M(W)$

6.4 Conclusion

In the standard models analysing the optimal contract for a risk averse manager, the manager can affect the stock price of the firm in an ad hoc fashion by exerting effort. These standard models predict that the optimal incentive for the manager is decreasing in the level of volatility. However, unlike the present situation, in these models the shareholder can only observe a noisy signal of the manager's effort.

Efforts to separate the effect of idiosyncratic and systematic risk have shown that it is idiosyncratic risk that matters when a manager is constrained to hold the stock of his firm but can trade the market portfolio.⁸ The effect in these models is still the same though because the more volatile a firm's share price is, the more likely it is that the manager's effort just gets "lost in the noise". This result is however driven by the assumption that effort is unobservable. Crucial to our results is the assumption that effort is observable.⁹

Given the typically ad-hoc nature of the manager's assumed impact on the firm, we attempted to provide a more detailed analysis of the manager and his impact on the firm. By starting from the firm's cashflows we are able to introduce the manager's valuation problem and the implications that this has for the shareholder. Allowing effort to reduce the cost of investment means we now have an explicit relationship between effort and the value of the firm that can be analysed in detail.

With respect to the previous literature on executive compensation, the primary finding of this chapter is that α^* does not unambiguously decrease with the level of the firm's volatility (σ_y). We find that when the manager's subjective valuation and an explicit channel for effort are acknowledged, the effect of volatility on α^* depends crucially on the sign of ρ as this determines

⁸See Jin (2002).

⁹As discussed in Section 3.3, this is a weakness of our setup. It does however allow us to isolate the agency issues caused by idiosyncratic risk and effort.

the impact of volatility on the manager's valuation of the firm. Analysis of the impact that the manager's level of skill and the scope for cost reductions have on the optimal contract also yields some interesting insights. Unsurprisingly the lower the investment cost can be driven, the higher the optimal incentive level. More interesting though is that the optimal contract is practically insensitive to the ceiling on the investment cost (i.e the cost when no effort is exerted). The finding that a more skilled manager (high λ) is given a lower level of α is also initially surprising, until one considers that this means he can achieve cost reductions with less effort and thus needs less incentive to work hard.

The personal characteristics of the manager also yield some surprising results. The relationship between the manager's risk aversion and the optimal level of managerial ownership is very complex. Despite this, optimal firm ownership for a low γ manager is not that different from that for a high γ manager. This suggests that the real world un-observability of managerial risk aversion is not of critical importance when determining managerial compensation.¹⁰ The fact that the more a manager dislikes effort, the more incentive he needs to exert effort is unsurprising. Similarly, rich managers need more incentive to exert effort than poor managers, because the financial benefit of effort in utility terms is much lower when you have more money.

The key message of this chapter is that determining the optimal amount of the firm a manager should own in order to maximise the value for shareholders is a very complex problem. It requires specific examination of the project the manager will be managing as well as the personal characteristics of the manager - managerial share ownership is not a one size fits all proposition.

¹⁰Though as Figure 6.4 shows, the shareholder's payoff is materially effected by γ .

Chapter 7

Effort and the Timing Option

7.1 Introduction

The analysis of the manager's investment decision has so far been in a static setting, in which the manager makes a “now or never” decision of whether or not to invest. The literature on real options has shown that investment decisions can be quite different from what static models of investment predict. Specifically in an environment of uncertainty there is “value in waiting” as this allows the agent to gain new information and thus resolve some of the uncertainty around the profitability of investment.

As discussed in Chapter 2, the majority of the work acknowledging the role of managers has focused on the information asymmetry between shareholders and managers, rather than the positive impact a manager can have. Given the significant and independent impacts that effort and the ability to wait have on investment behaviour, we will now extend the framework of Chapter 5 in order to determine whether the interplay of these factors has any interesting implications for managerial behaviour.

This chapter thus allows us to analyse how the manager's decision of whether

to invest or wait is affected by the ability to exert effort.

The rest of this chapter is set out as follows: Section 7.2 sets up the model, Section 7.3 outlines the numerical solution method used to solve the model, Section 7.4 discusses the results of the model and Section 7.5 summarises the results of this chapter.

7.2 Setup

In this situation the manager is waiting to invest and thus at every point in time he chooses his consumption, portfolio investment and whether or not the firm will invest. If the manager decides to invest he pays his proportion of the investment cost ($\alpha I[e]$), exerts the optimal level of effort determined by Equation (5.4) and moves to Stage 2 (i.e. the post-investment state). Thus the investment payoff is simply the Stage 2 value function when the manager can exert effort that we analysed in Chapter 5. Therefore the investment payoff is simply Equation (5.3) evaluated at $\hat{e}$

$$J^e(W, Y) = -\frac{1}{\gamma r} e^{-\gamma r(W - \alpha I + \alpha G[Y] + \frac{\Phi^2}{2\gamma r^2})} - \theta \hat{e} \quad (7.1)$$

Conversely if the manager chooses to defer the firm's investment, his intertemporal budget constraint is given by Equation (3.7). Following the same process as for Stage 2 and using the 1 superscript to denote the Stage 1 value function we can simplify the General HJB (Equation (3.9)) down to the following when the manager is waiting to invest

$$\begin{aligned} \beta J^1(W, Y) = \max_{C, \pi} & \left[-\frac{1}{\gamma} e^{-\gamma C} + J_t^1 + J_w^1 (rW + \pi (\mu_m - r) - C) + J_y^1 \mu_y \right. \\ & \left. + \frac{1}{2} J_{ww}^1 (\pi \sigma_m)^2 + \frac{1}{2} J_{yy}^1 (\phi^2 + \rho^2 \sigma_y^2) + J_{wy}^1 (\pi \rho \sigma_y \sigma_m) \right] \end{aligned} \quad (7.2)$$

with the following first order conditions

$$FOC_C : C^* = -\frac{\ln(J_w^1)}{\gamma}$$

$$FOC_{\pi} : \pi^* = -\frac{J_w^1(\mu_m - r)}{J_{ww}^1\sigma_m^2} - \frac{\rho J_{wy}^1\sigma_y}{J_{ww}^1\sigma_m}$$

Therefore at every point in time the manager will compare the solution to the HJB (Equation (7.2)) with the payoff from investment (Equation (7.1)) and if the payoff is greater he will invest, otherwise he will wait. Given the non-linear nature of the HJB for this problem, it must be solved numerically.

7.3 Numerical Solution Method

7.3.1 Description of algorithm

The complexity of the differential equation that needs to be solved (Equation (7.2)) means that some modifications to standard finite difference techniques are required to obtain a stable solution. While not as severe as that considered in Chapter 9, the choice of boundary condition can have a large impact on the stability of the solution. We also impose a finite length on the option to delay.¹

The solution procedure used is a “policy iteration”² finite-difference algorithm. The numerical algorithm begins by first defining a grid in (W, Y) space for each point in time (t_i) , where the grid co-ordinates are defined as

$$\begin{aligned} dW &= \frac{W_{max} - W_{min}}{W_N} \\ W_i &= W_{min} + dW(i - 1) \\ dY &= \frac{Y_{max} - Y_{min}}{Y_N} \\ Y_i &= Y_{min} + dY(i - 1) \end{aligned}$$

¹The option expires after a certain period of time at which point the manager faces a now-or-never decision.

²See Judd (1998) for a description of policy iteration algorithms.

$$dt = \frac{t_{max} - t_{min}}{t_N}$$

$$t_i = t_{min} + dt(i - 1)$$

where $(Y_{max}, W_{max}, t_{max})$ are the maximum values for the grid, $(W_{min}, Y_{min}, t_{min})$ are the minimum values and (W_N, Y_N, t_N) are the number of steps. Thus one can either think of the grid as t_N two dimensional matrices with dimensions of $W_N \times Y_N$, or a single three dimensional matrix with dimensions $W_N \times Y_N \times t_N$.

The finite difference approximation of Equation (7.2) is obtained using the following approximations for the derivatives.

$$J_t \approx \frac{J[W_i, Y_i, t_i] - J[W_i, Y_i, t_{i-1}]}{dt}$$

$$J_W \approx \frac{J[W_{i+1}, Y_i, t_i] - J[W_{i-1}, Y_i, t_i]}{2dW}$$

$$J_Y \approx \frac{J[W_i, Y_{i+1}, t_i] - J[W_i, Y_{i-1}, t_i]}{2dY}$$

$$J_{WW} \approx \frac{J[W_{i+1}, Y_i, t_i] - 2J[W_i, Y_i, t_i] + J[W_{i-1}, Y_i, t_i]}{dW^2}$$

$$J_{YY} \approx \frac{J[W_i, Y_{i+1}, t_i] - 2J[W_i, Y_i, t_i] + J[W_i, Y_{i-1}, t_i]}{dY^2}$$

$$J_{WY} \approx \frac{J[W_{i+1}, Y_{i+1}, t_i] - J[W_{i+1}, Y_{i-1}, t_i] + J[W_{i-1}, Y_{i+1}, t_i] + J[W_{i-1}, Y_{i-1}, t_i]}{4dWdY}$$

That is, for derivatives with respect to either Y or W , central difference approximations are used, while for the time derivative (J_t) a backward difference is used. The result of this is that it is possible to solve for the value function at date t_{i-1} ($J^1(W_i, Y_i, t_{i-1})$) as a function of the date t_i value function, consumption and portfolio investment functions. The consumption (C_t) and portfolio investment (π_t) functions are the “policy functions” for the “policy iteration” algorithm.

With that in mind, the steps to the algorithm are as follows:

1. The value function at the last possible date for investment (t_{N+1}) is

calculated as $J^1(W_i, Y_i, t_{N+1}) = \max[J^M(W_i), J^e(W_i, Y_i)]$. In other words the manager makes a now-or-never decision between investing ($J^e(W_i, Y_i)$) and not investing ($J^M(W_i)$).

2. The policy functions are calculated at t_{N+1} using $J(W_i, Y_i, t_{N+1})$.
3. Boundary conditions are imposed along the four boundaries $(W_1, W_{N+1}, Y_1, Y_{N+1})$
4. From here the algorithm progressively steps backwards in time, using the solutions for the value function and policy functions at t_i to calculate the value function (and thus the policy functions) at t_{i-1} .

7.3.2 Boundary conditions

Boundary condition at Y_{N+1}

Along this boundary Y is very large and thus the boundary condition we impose is simply that the manager invests immediately. This is a common approach in numerical option valuation work.³ The condition along this boundary is

$$J^1(W_i, Y_{n+1}, t_{i-1}) = J^e(W_i, Y_{n+1}) \quad (7.3)$$

Boundary condition at Y_1

Along this boundary Y is very small and thus the boundary condition we impose is simply that the manager does not invest/abandons the project. In other words he receives the Merton (1969) value function outlined in Section 3.8.

$$J^1(W_i, Y_1, t_{i-1}) = J^M(W_i, Y_1) \quad (7.4)$$

³For financial (as opposed to “real”) option valuation, the corresponding assumption is that the option is exercised immediately.

Boundary condition at $W_{N+1}, W_{i=1}$

Unlike the boundary conditions for Y , there is no clear theoretical condition that can be imposed to govern behaviour at the two W boundaries. In this situation a numerical boundary condition is typically imposed. The simplest method used in this regard is to set the finite difference approximation for the second derivative equal to zero.⁴ For example, we could evaluate the second derivative with respect to W at t_N , set it equal to zero and solve for the value of $J(W_{N+1}, Y_i, t_i)$ as follows

$$0 = \frac{J(W_{N+1}, Y_i, t_i) - 2J(W_N, Y_i, t_i) + J(W_{N-1}, Y_i, t_i)}{dW^2}$$

$$\rightarrow J^1(W_{N+1}, Y_i, t_i) = 2J(W_N, Y_i, t_i) - J(W_{N-1}, Y_i, t_i)$$

However, using this boundary condition along the W boundaries yields unstable solutions. We therefore implement boundary conditions related to the slope of the Merton (1969) value function. More specifically, it can easily be shown that the following relationship holds between the first and second derivatives of the Merton value function.

$$J_{WW}^M = -\gamma r J_W^M \quad (7.5)$$

Using the finite difference approximations outlined above for J_{WW}^W and J_W^M this becomes

$$\frac{J(W_{i+1}, Y_i, t_i) - 2J(W_i, Y_i, t_i) + J(W_{i-1}, Y_i, t_i)}{dW^2} = -\gamma r \frac{J(W_{i+1}, Y_i, t_i) - J(W_{i-1}, Y_i, t_i)}{2dW} \quad (7.6)$$

To derive the condition at W_1 , we evaluate Equation (7.6) at W_2 and then solve for $J[W_1, Y_i, t_i]$. This yields the following

⁴In other words the boundary condition is a linear projection of the previous values.

$$J^1(W_1, Y_i, t_i) = \frac{(dWr\gamma + 2)J(W_3, Y_i, t_i) - 4J(W_2, Y_i, t_i)}{dWr\gamma - 2} \quad (7.7)$$

Similarly we can derive an expression for $J^1[W_{N+1}, Y_i, t_i]$ by evaluating Equation (7.6) at W_N .

$$J^1(W_{N+1}, Y_i, t_i) = \frac{(dWr\gamma - 2)J^1(W_{N-1}, Y_i, t_i) + 4J^1(W_N, Y_i, t_i)}{2 + dWr\gamma} \quad (7.8)$$

7.4 Results

Given that our results are based on a numerical solution method, we begin by again setting out our base case parameters, which are identical to those used in previous sections. Using these parameters the model is then solved

Table 7.1: *Base case parameters for Chapter 7*

$$\begin{array}{cccccc} \gamma = 1 & r = 0.1 & \mu_m = 0.15 & \mu_y = 0.1 & A = 100 & \lambda = 1 \\ \beta = r & \sigma_m = 0.3 & \sigma_y = 0.3 & \rho = 0.0 & B = 80 & \theta = 0.1 \end{array}$$

as outlined in Section 7.3. A useful starting point for examining the results of this model is to compare the value function when the manager can delay investment with that of the “now or never” case.

Figure 7.1 plots the manager’s value function for both the now-or-never case (panel (a)), when delay is possible (panel (b)) and the difference between them (panel (c)). As can be seen in Figure 7.1, once the ability to delay is introduced, the value function is no longer kinked. Instead it is smoothed out over the region near the investment threshold. This smoothing out is a result of the option value that is embedded in the project when delay is possible. This option premium is best illustrated by examining the difference between the “now or never” and dynamic value functions. As can be seen

Figure 7.1: Comparison of the “Now or Never” and Dynamic Value functions

Calculated using the base case parameters.

from Figure 7.1 (c), the option premium disappears when Y is very small and very large reflecting the fact that when Y is very small the ability to invest has little value as it is very unlikely that investment will take place. On the other hand when Y is very large there is little value having timing flexibility since it is optimal to invest straight away. The source of this option value is most easily understood in terms of the “Bad News Principle”, which is best described by Dixit and Pindyck (1994) who state that

“... it is the ability to avoid the consequences of “bad news” that leads us to wait.”

In short the manager delays investment for fear of receiving bad news about the profitability of the project after investment has taken place. The direct result of this is that an increase in the magnitude of the potential bad news that can be received will increase the value of waiting to invest. The direct result of there being a value in waiting is that the investment threshold will exceed the standard NPV threshold. In the context of the current model it is difficult to express the NPV and option value because investment has both a financial (αI) and non-financial cost (θe). To express a valuation of the project and the option in dollar terms we turn to the concept known as “Utility Indifference Pricing” which we also implement in Chapter 9. Henderson (2009) defines the utility indifference price as follows

The utility indifference buy (or bid) price p^b is the price at which the investor is indifferent (in the sense that his expected utility under optimal trading is unchanged) between paying nothing and not having the claim C_T and paying p^b now to receive the claim C_T at time T .

In other words, the utility indifference valuation is the dollar amount that would leave the manager indifferent between having the option/project and not having it. Therefore the utility indifference valuation of the project and the option can be found by solving the following equations for UI^{RO} and UI^{NPV}

$$\begin{aligned} J^{RO}(W, Y) &= J^M(W + UI^{RO}) \\ J^E(W, Y) &= J^M(W + UI^{NPV}) \end{aligned}$$

Thus UI^{RO} (UI^{NPV}) is the amount of money which would make an agent indifferent between having the option (project) and being in the Merton (1969) world where he does not have the option. Substituting in the Merton value function and solving for UI^{RO} (UI^{NPV}) yields the following

$$UI^{RO} = \frac{-2W\gamma r^2 - \eta^2 - 2r \log[-r\gamma J^{RO}(W, Y)]}{2r^2\gamma} \quad (7.9)$$

$$UI^{NPV} = \frac{-2W\gamma r^2 - \eta^2 - 2r \log[-r\gamma J^E(W, Y)]}{2r^2\gamma} \quad (7.10)$$

If we plot these two equations for the base case we see the familiar Dixit and Pindyck (1994) story holds⁵ where the option value is strictly positive and exceeds the NPV for relatively low values of Y . The region where the option value exceeds the NPV is the area where the manager waits. For larger values of Y the NPV and the option value converge, representing the fact that the project is so attractive that option holder would invest immediately and thus there is no value in waiting.

Figure 7.2: *Utility Indifference Valuation of the Option and the Project*

Calculated using the base parameters where the solid line= UI^{RO} , and the dashed line= UI^{NPV}

In the now or never setting the investment threshold is simply the point at which the NPV is positive. In the dynamic setting, when investment occurs it must be the case that the value of the option is simply the NPV. Therefore,

⁵See, e.g., Figure 5.3 of Dixit and Pindyck (1994).

we can characterize the dynamic investment threshold $\hat{Y}$ as the point where the option value and the project NPV first become equal.

To understand the dynamics of the model it is useful to examine the manager's optimal investment in the risky asset (π^*) and the investment thresholds. Figure 7.3 plots the manager's optimal investment in the risky asset *while he is waiting to invest* (panel (a)), the contours⁶ of that plot (panel (b)) and the investment thresholds (panel (c)) all as functions of the manager's financial wealth (W) and the level of the cash flow (Y) for the base case parameters. In panels (b) and (c) the purple region is the area where immediate investment is optimal. In panel (c), the tan region is the area where investment is not optimal in a now-or-never world (alternatively, delay is optimal when waiting is possible) and the light blue region is where investment would be optimal in a now-or-never world but delay is optimal when waiting is possible.

For the base case we can see in Figure 7.3 that the investment threshold exhibits the behaviour that would be expected when the option to wait is introduced: the threshold when waiting is allowed is approximately a parallel shift upward of the threshold from the now or never case (e.g. Figure 5.10). However, once the manager's portfolio behaviour is examined the similarities between our model and the standard real options story end.

The thing that immediately stands out in Figure 7.3 is the spike in the optimal investment in the risky asset. Like the static investment decision of Chapter 5, for low levels of wealth the manager will exert enough effort to drive the investment cost down to B and for high levels of wealth the manager will exert no effort. It is the area between zero and “maximum”⁷ effort that the spike occurs. This can be thought of as the region where the

⁶A top-down, two dimensional view of panel (a).

⁷By maximum we simply mean the level of effort that sets $I[e] \approx B$. Given that $I[e]$ is asymptotic around B the investment cost will never actually reach B , and thus there is not a maximum level of effort.

Figure 7.3: *Waiting to invest: Optimal Portfolio Investment (π^*), Portfolio Contours and Investment Thresholds*

Calculated using the base case parameters. For the Investment thresholds the tan region represents the region where investment is not optimal, the light blue region is where now-or-never investment is optimal, but in the dynamic case it is optimal to delay investment and the purple region is where it is optimal to invest in the dynamic setting.

manager has a “choice” over the level of effort he exerts: if the manager is very poor he will exert maximum effort regardless, whereas if the manager is very rich, his marginal utility of wealth is so low that he has no incentive to exert any effort. This choice over effort is the defining feature of the model and the driver of the portfolio spike.

But why does the ability to adjust effort result in the manager investing a large amount of money in the risky asset? The answer is deceptively simple: the manager can exert effort to offset adverse outcomes from his investment in the risky asset. Put more simply, potential future effort acts as a hedge against the manager’s portfolio.⁸ Following this line of logic it is unsurprising to see that within the region where there is a choice over effort, the spike is at its highest point as effort approaches zero (high wealth) and at its lowest

⁸ This is a similar logic to the result of Bodie, Merton, and Samuelson (1992). In this paper investors use the ability to alter their labour supply later in life as a hedge for volatile investment returns. The key distinction between that model and the present setting is that in Bodie et al the agent can continuously vary his labour supply, whereas the manager here makes a one off, irreversible choice over effort when investment occurs.

point when effort is high (wealth is low). This is because if the manager is currently in a situation where he would exert little effort upon investment, he has significant scope to alter his effort to offset any adverse shocks he receives to his portfolio at the time of investment. Conversely if he is in a situation where he would exert maximum effort upon investment, there is little scope to alter effort and thus the spike dissipates.

In the standard real options framework the source of the value of waiting and thus the key driver of the model, is the volatility of the cash flow. As it turns out, in the current framework, despite the fact that we have assumed the cash flow is uncorrelated with the risky asset, it is the volatility of the risky asset that drives the interesting behaviour of the model. More interesting though, is that an increase in market volatility actually reduces the value of waiting, contrary to the standard model.

Figure 7.4 reproduces Figure 7.3 for various levels of σ_m . This is the standard form that will be used for the graphical analysis of this chapter. In Figure 7.4 we can see that in the region where effort would either be zero or at its “maximum”, there is a constant value in waiting. Given that effort is essentially constant in these regions, the value in waiting in these regions is simply related to bad news concerning Y . However, the striking feature of Figure 7.4 is that as σ_m falls, a spike in the investment threshold begins to emerge in the region where the manager has “choice” over effort and this spike peaks near the point where effort would be zero if investment occurred. This then poses two main questions that need to be answered: Why is there value in waiting in this region? and why is this effect only relevant when market volatility is low?

We now need to understand what bad news the manager can receive in the region where he has a choice over effort, compared to the regions that he does not. Given that uncertainty in Y is already accounted for, the only other source of uncertainty the manager faces is wealth uncertainty due to his holdings of the risky asset. But what bad news does the manager receive

Figure 7.4: σ_m and the manager's behaviour

from wealth shocks and how can the consequences of this bad news be avoided by delaying investment? To understand this it is essential to understand the concept of diminishing marginal utility and its effect on the utility valuation of the investment cost. Diminishing marginal utility simply states that as wealth increases an agent will value an additional dollar of wealth less. Thus as wealth increases, the utility valuation of the investment cost decreases and vice-versa when wealth decreases. It is also important to note that the actual cost of investment can be broken into the financial component and the cost of effort exerted which is non-financial and thus does not depend on the manager's level of wealth. Because the cost of effort does not depend on the manager's level of wealth, as wealth decreases spending money on the project becomes relatively more expensive than exerting effort. This is the source of the bad news that the manager can receive since he is exposed to wealth shocks: because investment is irreversible, if the manager invests now and then subsequently receives a negative wealth shock, he will have wished he had exerted more effort as the money he spent on the project is now relatively more valuable to him.

This also explains why the peak of the investment threshold spike (and thus the area where there is the greatest value in waiting) is at the area where effort is very close to zero. The bad news principle states that as the magnitude of the bad news that can be received increases, the value in waiting will increase. To see why the potential bad news is worst when the manager is exerting little to no effort it is useful to compare this point to the case where the manager has exerted almost enough effort to drive the cost down to B . In the case where the manager would exert close to maximum effort upon investment, if he subsequently receives a negative wealth shock this is not really bad news⁹ because he is already exerting maximum effort and thus could not have reduced the investment cost even if he wanted to. In contrast, when effort would be close to zero upon investment, if the manager receives a

⁹Clearly a negative wealth shock is "bad news", just not with respect to the effort decision in this context.

negative wealth shock post investment this is very bad news because he could have paid a large proportion of the investment cost through effort instead of wealth.

Having established why there is value in waiting in addition to the standard model, we must now answer the second question of why this effect dissipates as market volatility increases. This is interesting because in standard models of investment under uncertainty, an increase in volatility typically increases the value of waiting. In other words an increase in volatility increases the magnitude of the bad news that can be received after investment has occurred. Yet we find that the opposite occurs. This initially counter-intuitive result can actually be explained very simply if one thinks about the effective volatility of the manager's wealth rather than the volatility of the risky asset. As the volatility of the risky asset falls, all other things being equal, the Sharpe ratio of that asset will increase.¹⁰ Therefore the risky asset becomes relatively more attractive to the agent and he invests a greater amount of wealth in the risky asset. This means that although the asset has become less risky, the manager is holding much more of it which actually makes his wealth more volatile. To illustrate this we can look at the polar cases of $\sigma_m = 0.3$ and 0.1 in Figure 7.4. For $\sigma_m = 0.3$, the base line amount the manager invests in the risky asset is around 5.5. This leads to an effective volatility of wealth of $\pi\sigma_m = 0.3 \times 5.5 = 1.65$. If we contrast this to the case of $\sigma_m = 0.1$, the baseline amount invested in the risky asset is 50, which gives an effective volatility of wealth of $\pi\sigma_m = 0.1 \times 50 = 5$. Thus the manager's wealth is more than three times as volatile compared to the case when $\sigma_m = 0.3$.

Given that the volatility of wealth determines the magnitude of the bad news that the manager can receive (i.e bigger wealth shocks are worse news), we can see that the manager's endogenous adjustment of his portfolio causes the bad news to be less severe as volatility increases. Therefore the value

¹⁰Where the Sharpe ratio is defined as $S = \frac{\mu_m - r}{\sigma_m}$.

of waiting decreases as volatility increases. In fact, this effect is not limited to the volatility of the market asset, it depends on the Sharpe ratio of the market asset. If the Sharpe ratio of the market asset increases¹¹ the manager will increase his holdings of the risky asset which will make his wealth more volatile and thus increase the magnitude of possible bad news.

The other factor that affects the magnitude of wealth related bad news is the extent to which effort can reduce the cost of investment. The simplest way to think of this is holding A constant and reducing B .¹² Given that the lower B is, the more the cost of investment can be reduced, a lower B means that the magnitude of possible bad news is greater. This is because the more the manager can reduce the cost through effort, the more he will regret not exerting effort if he receives a large negative wealth shock. This is shown in Figure 7.5 where we can see that as B falls the spike in the investment threshold becomes substantially larger. The second thing we notice about Figure 7.5 is that as B falls the portfolio spike also becomes significantly larger. This stems from the fact that the smaller B is, the more effort can be used as a hedge to offset portfolio shocks

If we now turn to the volatility of the cash flow we can check to see whether the standard real options result holds.¹³ In Figure 7.6 we can see that the investment threshold increases as σ_y increases which is what we would expect. We know that in the region where there is no choice over effort, the value from waiting is purely related to bad news concerning Y and thus an increase in σ_y scales up the threshold in these areas. Interestingly, as σ_y gets quite large we can see that the spike starts to disappear, or get drowned out. This suggests that as σ_y gets very large, bad news concerning Y starts to become more important than bad news concerning W . Because an increase in σ_y does not change the effectiveness of effort as a hedge against the risky asset,

¹¹This occurs through an increase in μ_m or a decrease in σ_m or r .

¹²The same logic applies if we hold B constant and increase A .

¹³I.e. whether or not an increase in volatility of the cash flow causes an increase in the value of waiting.

Figure 7.5: B and the manager's behaviour (a) $B = 100$ (b) $B = 90$ (c) $B = 80$ (d) $B = 70$ Calculated using the base case parameters and $\sigma_m = 0.1$

the height of the spike remains relatively constant as σ_y changes. However what we do notice is that the contours begin to extend much further away from the investment threshold as σ_y increases. This happens because when σ_y is large, even if Y is currently quite far below the investment threshold, the manager knows that it is quite likely that it will get there in the future and thus the manager is willing to gamble with his portfolio now.

Another interesting aspect of Figure 7.6 is that when $\sigma_y = 0$ there is still a value of waiting in the region where there is no choice over effort. Given our hypothesis that the value from waiting in this region stems purely from bad news concerning Y , one would expect that in the deterministic case the threshold would be the same as the now or never case, i.e. there would be no value in waiting in this region. However it has been shown that in the deterministic case there can still be a value in waiting if the cash flow has a positive growth rate.¹⁴ This value in waiting stems from the fact that if the manager defers investment, the present value of both the revenues and the cost will decrease, but the present value of the revenues will decrease by less since the cash flow has a positive growth rate. Thus deferring investment can increase the NPV of the project. To check that this is what is occurring here we will set $\mu_y = 0$ and check that there is no non-effort related value in waiting when $\sigma_y = 0$. Figure 7.7 shows that this is the case.

Given the significant impact that the level of B has on the manager's behaviour, one would expect that the manager's skill (λ) and the personal cost to him of effort (θ) would also have a significant effect on optimal behaviour. As it turns out, λ and θ have opposite effects to each other and the impact is not as significant as one would expect. In Figure 7.8 we can see that when $\lambda = 0$ (and thus effort is completely ineffective) we have a flat investment threshold. As λ increases and effort becomes more effective, all that happens is that we get a much larger range where the manager will exert full effort. Therefore the region where the manager has a choice over wealth occurs at

¹⁴For a detailed example see Dixit and Pindyck (1994, pp138-139).

Figure 7.6: σ_y and the manager's behaviour Calculated using the base case parameters and $\sigma_m = 0.1$

Figure 7.7: *Investment Threshold when $\sigma_y = 0$*

Calculated using the base case parameters and $\sigma_m = 0.1$

a much higher level of wealth. Other than the fact that the spike shifts to the right as λ increases, the shape of the investment threshold and portfolio spikes are essentially unchanged. If we carry out the same exercise for θ we find the opposite behaviour in that the spike shifts to the left and eventually disappears as θ gets larger.

In the now or never model of Chapter 5 the proportion of the firm that the manager owns (α) had a significant impact on the manager's investment decision. For the base case parameters α turns out to have little impact on the manager's investment decision. Figure 7.9 shows that as α increases the investment threshold spike is relatively unchanged. This is interesting given the significant impact α has on the investment threshold in the now or never case. The way the spike does change is that it becomes "thinner" as α increases. This occurs because all other things equal an increase in α increases effort. Therefore as W falls, "maximum" effort is reached quicker and thus the spike is not as wide. What does however change quite significantly is the spike in the manager's investment in risky assets. As α increases the portfolio spike becomes significantly larger. This reflects the fact that as α

Figure 7.8: λ and the manager's behaviour Calculated using the base case parameters and $\sigma_m = 0.1$

Figure 7.9: α and the manager's behaviour Calculated using the base case parameters and $\sigma_m = 0.1$

increases, the benefit from effort (and thus the effectiveness of effort as a hedge) increases.

The final parameter of interest is the manager's risk aversion (γ). In the now or never case we know that the area where the manager has a "choice" over effort decreases as γ increases due to the concavity of the value function increasing.¹⁵ In Figure 7.10 we can see that because this region is much smaller, the area where there is an effort-related value in waiting is also much smaller. Given that the manager has less "choice" over effort it is unsurprising that the effort-related value in waiting starts to dissipate. This is because it is the regret over the chosen level of effort that drives the value of waiting in this model. Another interesting feature of Figure 7.10 is that the portfolio spike also decreases as γ increases. Given that the portfolio spike is a result of the manager using effort to hedge his risky asset exposure, as the manager has less choice over effort it becomes a less effective hedge and thus this result is expected. This also gives us the intuitive result that as risk aversion increases the manager's propensity to "gamble" in the risky asset decreases.

7.5 Summary

In this section we examined the dynamic investment decision of a manager who is constrained to hold a portion of the firm's cash flow in his portfolio. We found that interplay of the manager's ability to invest in the market and exert investment cost reducing effort significantly alters the results of a model relative to standard models of investment delay.¹⁶ In the absence of effort,

¹⁵Recall that because the concavity of the value function increases with γ , marginal utility quickly changes from being very high to very low as W changes. Because marginal utility directly determines the benefits of effort, this makes the choice of effort more of an all or nothing decision.

¹⁶See for example Dixit and Pindyck (1994) and the discussion of the delay option in Section 2.3.2.

Figure 7.10: γ and the manager's behaviour

Calculated using the base case parameters and $\sigma_m = 0.1$

this problem has been modelled by Miao and Wang (2007) and our model obtains equivalent results when W is very large and thus the manager will never exert any effort.¹⁷

However, when W takes on more moderate values and the manager has a “choice” over his effort level, our results differ markedly from Miao and Wang (2007). In particular there is a “spike” in the investment threshold in the area where the manager has a “choice” over effort. This results in investment being delayed beyond what standard real options models predict.

¹⁷A similar logic holds if W is very small as the manager will always exert enough effort to drive the investment cost to B . When W is very small we thus obtain the same results as Miao and Wang (2007) with an assumed investment cost of B .

Chapter 8

The Shareholder's Dynamic Problem

8.1 Introduction

We began our analysis by examining the manager's valuation of a cash flow that he cannot trade. This was then extended to allow the manager to exert effort that reduces the cost of investment. With this in place we were then able to examine the shareholder's problem of determining the level of managerial ownership that optimally trades off the detrimental and beneficial effects of managerial ownership. At this point we stepped back and recognised that the manager's effort and investment decisions were static in that he was making a now or never decision. We thus extended the manager's problem to a dynamic setting where he could delay investment.

In other words, we have examined the manager's effort/investment decision in both a static and dynamic setting and we have also examined the shareholder's problem in the static (now-or-never) setting. We have not examined the shareholder's hiring decision in the dynamic setting.

The purpose of this chapter is thus to extend the model of Chapter 7 to examine how the manager's dynamic investment and effort decisions affect

the shareholder's decision of whether or not to hire a manager. While we examine the shareholder's decision to hire given certain parameters, it is also important to determine what "type" of manager the shareholder would prefer. We therefore also examine the impact on shareholder wealth of different parameters.

This chapter is laid out as follows: Section 8.2 sets out the framework we use to model the shareholder's hiring decision, Section 8.3 analyses the shareholder's payoff from hiring a manager, Section 8.4 analyses the shareholder's decision when hiring is a now-or-never decision, Section 8.5 analyses the shareholder's decision when hiring can be delayed, Section 8.6 conducts comparative statics analysis of the effect of various parameters on shareholder wealth and Section 8.7 summarises the results of this chapter.

8.2 Framework

The shareholder has the choice to hire a manager or run the project himself. If the shareholder chooses to run the project himself, then the shareholder incurs a monitoring cost κdt every period of length dt prior to investment. On the other hand, the shareholder can choose to hire a manager in which case he receives a payoff whose timing and level is determined by the manager.¹ The key differences between the manager and the shareholder are that the manager can exert cost-reducing effort and also monitor the project prior to investment at a lower cost than the shareholder.²

The form of the shareholder's hiring decision has significant implications for how this problem is modelled. We propose two ways to model this decision:

- **Static Hiring Decision**

¹The level is determined in the sense that the manager chooses at what level of Y he will invest.

²For simplicity we normalise the manager's cost of monitoring the project to zero.

At time zero the shareholder makes an irreversible decision about whether or not to hire a manager; or

- **Dynamic Hiring Decision**

The shareholder can delay the decision on whether or not to hire a manager, but if and when he does hire a manager, the hiring decision is irreversible.

These two models represent different ends of the spectrum of hiring flexibility. The first model represents no flexibility, that is, the shareholder is presented with the opportunity to hire a manager and if he delays the decision he loses the option to hire that or any other manager.³ The second model represents perfect flexibility in that any time prior to investment the shareholder can decide to hire a manager. Real world hiring decision will likely contain elements of both models and thus examining both ends of the spectrum is a useful way to analyse the problem. Note that in neither case can the manager be removed.

The reason for making the hiring decision irreversible is that reversibility would alter the manager's problem as we have framed it thus far. The impact of the potential to be dismissed on the manager's dynamic investment and effort decisions is an interesting question that we leave to future research.⁴ In addition it is important to note that unlike Chapter 6, we will treat the share of the firm given to the manager (α) as exogenous (this is analogous to there being a market "price" for a manager). Endogenizing α in a dynamic setting is another interesting question we leave to future work.

There are two inputs which we require to solve either problem. These are

³Note that one of the limitations of this set up is that there is only one potential manager. Therefore the shareholder does not have to incur any search costs or choose between multiple managers. This is an area left for future research.

⁴In a slightly different setting this question has been examined by Hugonnier and Morellec (2007a). In their model the manager's payoff was not directly linked to the profitability of the firm.

the market value of the project rights if it is managed by a shareholder (the opportunity cost of hiring a manager⁵) and the payoff from hiring a manager. Given we assume that the shareholder is well diversified and thus values the project rights using the CAPM, calculation of the market value of the project rights is a straightforward and is left to Appendix A. We now derive the payoff to the shareholder from hiring a manager.

8.3 Payoff from hiring a manager

8.3.1 Setup

Given the characterizations of the problem, the starting point is to calculate the shareholder's payoff given that he has already decided to hire a manager. In this situation the payoff to the shareholder depends on two stochastic variables, the level of the cashflow (Y_t) and the level of the manager's wealth (W_t). The manager's wealth determines the level of effort he will exert and thus the eventual investment cost. In this sense the investment cost is stochastic from the shareholder's point of view. Thus there are parallels between this and the real options analysis of projects where the cash flow and investment cost are stochastic.⁶

Once the manager has been hired the shareholder waits for the manager to invest. Given that shareholder's value function now depends on the manager's wealth, it can be derived using the same methodology used for the manager's problem. Following the same procedure used for the manager, Equation (3.8) becomes the following

$$S(W, Y, t) = e^{-\beta dt} E[S(W(t + dt), Y(t + dt), (t + dt))] \quad (8.1)$$

⁵When the shareholder hires a manager, he gives up the value of the project if he managed it himself.

⁶See, e.g. Dixit and Pindyck (1994), Section 6.5 pp207-211.

where this differs from Equation (3.10) in that the shareholder's consumption and/or investment do not affect his valuation of the cash flow and he is not maximising the function himself. The general value function (Equation (3.9)) thus becomes⁷

$$\beta S(W, Y, t) = E[S_t dt + S_w dW + S_y dY + \frac{1}{2} S_{ww} dW^2 + \frac{1}{2} S_{yy} dY^2 + S_{wy} dY dW] \quad (8.2)$$

The interesting feature here is that W_t is the manager's wealth and thus dW is defined in exactly the same way it was in the manager's problem. Therefore the manager's consumption and (portfolio) investment decisions impact the evolution of the W_t and thus the payoff to the shareholder. Substituting in dW and dY from Equations (3.4) and (3.7) and simplifying gives the differential equation that the shareholder's payoff from hiring a manager must satisfy:

$$\begin{aligned} \beta S(W, Y, t) = & S_t + S_w (rW + \pi(\mu_m - r) - C) + S_y \mu_y \\ & + \frac{1}{2} S_{ww} (\pi \sigma_m)^2 + \frac{1}{2} S_{yy} (\phi^2 + \rho^2 \sigma_y^2) + S_{wy} (\pi \rho \sigma_y \sigma_m) \end{aligned} \quad (8.3)$$

We already know from Equation (6.2) that the shareholder's payoff when the manager invests is

$$P^M[\alpha] = (1 - \alpha) \left(\frac{Y}{r} + \frac{\mu - \rho \sigma_y \Phi}{r^2} - I[\hat{e}] \right) \quad (8.4)$$

Therefore if the manager invests the shareholder receives the payoff defined by Equation (8.4), otherwise he receives the solution to Equation (8.3). This can be represented formally as

$$S^M(W, Y, t) = \begin{cases} P^M[\alpha] & \text{if } \alpha(G[Y] - I[\hat{e}]) \geq V^{RO}(W, Y, t) \\ \text{Solution to (8.3)} & \text{if } \alpha(G[Y] - I[\hat{e}]) < V^{RO}(W, Y, t) \end{cases}$$

⁷Note that this is not technically a Bellman equation since the shareholder does not have any choice variables.

We have already solved $P^M[\alpha]$, $G[Y]$, $\hat{e}$ and $V^{RO}(W, Y, t)$ when examining the manager's problem in Chapters 4, 5 and 7. Therefore all that is required to calculate $S^M(W, Y, t)$ is to solve the differential Equation (8.3).

8.3.2 Solution

Because Equation (8.3) does not involve any policy functions that must be solved, the problem can be solved using standard finite difference methods. Given that the technique is standard, it is not reproduced here. Before continuing it is worth repeating the parameters that we will use as the base case for this section.

Table 8.1: *Modified base case parameters for Chapter 8*

$\gamma = 1$	$r = 0.1$	$\mu_m = 0.15$	$\mu_y = 0.1$	$A = 100$	$\lambda = 1$	$W = 1$	$\kappa = 0.1$
$\beta = r$	$\sigma_m = 0.1$	$\sigma_y = 0.3$	$\rho = 0.0$	$B = 80$	$\theta = 0.1$	$Y = 10$	

Note that unlike the rest of Part II, we are assuming that $\sigma_m = 0.1$. This is because we are interested in examining the implications of the investment “spike”.

Figure 8.1 plots the shareholder's value function in the first date ($t = t_1 = 0$) as well as the date at which the investment option expires ($t = t_n = 10$). We can see from Figure 8.1 that the standard real options intuition still holds in that the shareholder's payoff is smoothed out when delay is possible. However, one has to keep in mind that the manager is making the decision for the shareholder. The obvious distinction between the two graphs is that when $t = t_n$ the graph is kinked as opposed to being *relatively* smooth. This occurs because there is no longer any opportunity to delay investment and thus the project value is simply the payoff from immediate investment.⁸ Conversely, at the first date, where the manager still has significant scope to

⁸Which could of course be zero if the manager does not invest.

delay investment, the shareholder receives value from the manager delaying the investment decision. This is evidenced by the smooth shape of the graph along the Y axis. The reasons behind the smoothness is that even when the project would have a negative immediate payoff, there is value in waiting to see if the project improves. The graphs also have the expected shape when one looks along the W axis in that the payoff increases as W decreases and then hits an asymptote. This is due to the facts that the manager's effort is higher for a lower W and that there is a limit on how much effort can reduce the investment cost.

Figure 8.1: Payoff from hiring a manager ($S^M(W, Y, t)$) at $t = t_1$ and $t = t_n$

Calculated using the base case parameters

It is important however to note that the undiversified manager is choosing an investment policy that maximises his utility. This will not necessarily coincide with the policy that would maximise the market value of the project rights. In fact, it turns out that in some situations the manager waits too long to invest from the shareholder's perspective - that is, the manager waits when investing immediately would give the shareholder a higher payoff. This is illustrated by comparing the shareholder's payoff from hiring a manager ($S^M(W, Y, t)$) and the payoff to the shareholder if the manager invested im-

mediately ($P^M[\alpha, Y]$). This comparison is shown in Figure 8.2

Figure 8.2: $S^M(W, Y, t)$ vs $P^M[\alpha, Y]$ and $P^{CAPM}[Y]$ vs $F^{CAPM}(Y, t)$

Calculated using the modified base case parameters and $W = 0$. The solid blue line represents the payoff from hiring a manager ($S^M(W, Y, t)$), the dashed green line represents the payoff to the shareholder if the manager invested immediately ($P^M[\alpha, Y]$), the solid red line is the CAPM valuation of the project rights ($F^{CAPM}(Y, t)$, i.e. option) and the dashed brown line is the CAPM valuation of the cash flow ($P^{CAPM}[Y]$).

In Figure 8.2 the solid blue curved line represents $S^M(W, Y, t)$ and the dashed green line represents $P^M[\alpha, Y]$. For comparative purposes we have also included the CAPM valuation of the cash flow ($P^{CAPM}[Y]$) as the dashed brown line and the CAPM valuation of the project rights⁹ ($F^{CAPM}(Y, t)$) as the solid red curved line. We can see that from just after $Y = 7$ the payoff from the manager investing is greater than the payoff to the shareholder from the manager delaying. Therefore in this situation the shareholder would be better off if the manager invested immediately. Even though he would prefer that the manager invested in this situation, he may still be better off than if he managed the project himself. This is illustrated by the fact over the range of Y we examine, $S^M(W, Y, t)$ exceeds $F^{CAPM}(Y, t)$.

The fact that the shareholder's payoff function ($S^M(W, Y, t)$) is actually below the payoff from immediate investment ($P^M[\alpha]$) for a wide range of Y , while

⁹i.e. the CAPM valuation of the project recognising the value of the delay option.

not unexpected, is interesting in and of itself. Figure 8.2 is plotted at $W = 0$ which under the base case parameters is in the area where the manager's investment "spike" occurs (see Chapter 7). Recall that in the "spike" region the manager delays investment beyond what a standard real options model predicts. Thus it makes sense that in this region the manager is delaying investment beyond what is optimal from the shareholder's perspective. To confirm this logic we re-plot Figure 8.2 for values of W outside of the "spike" region. Figure 8.3 does so for $W = -50$ and $W = 50$. The solid dark blue line is $S^M(-50, Y, t)$ while the solid brown line represents $S^M(50, Y, t)$. Given that the manager's optimal effort is a decreasing function of his wealth, it is not surprising that we find that $S^M(-50, Y, t) > S^M(50, Y, t)$. More importantly we find that in both cases the shareholder's value function from hiring a manager does not fall below the immediate payoff from investing. Thus for values outside of the "spike" region the manager doesn't appear to delay investment beyond what is optimal from the shareholder's perspective. It is also easy to show that as other parameters are changed in a way that reduces the spike¹⁰ we get a similar result. This occurs because the smaller the spike is, the less likely the manager is to delay investment beyond what is optimal from the shareholder's perspective.

It is useful to recall the results of Miao and Wang (2007) at this point. They found that relative to the risk-neutral model, investment would be delayed but that this is a second order effect. Our model is slightly different in that the point of comparison is a shareholder who uses the CAPM and incurs a monitoring cost. The assumption of the CAPM does not make a significant difference and thus (ignoring the monitoring cost) when W is very large or very small the effort decision is no longer a factor and thus we should obtain the same results as their model.

The only direct comparison that can be made between the shareholder's investment threshold absent a manager and the manager's investment thresh-

¹⁰For example setting $A = B$ or increasing σ_m . See Chapter 7.

Figure 8.3: $S^M(W, Y, t)$ vs $P^M[\alpha, Y]$ for $W = -50$ and $W = 50$

Calculated using the modified base case parameters. The dark blue line is $S^M(-50, Y, t)$ while the green line represents $S^M(50, Y, t)$.

old is when W is very large and thus the manager exerts no effort. This ensures that the investment cost is the same for the shareholder and the manager. For the base case parameters we find that the investment thresholds are the same for the manager and the shareholder. However, if the project's volatility (σ_y) is increased we get the second order effect of Miao and Wang (2007) in that the manager delays investment beyond what the shareholder would. In addition, if the monitoring cost (κ) is increased (which the shareholder bears but the manager does not), the manager invests later than the shareholder would in the absence of a manager. Intuitively, the monitoring cost makes the shareholder want to invest earlier to avoid the cost of monitoring, while idiosyncratic volatility makes the manager want to invest later for the reasons outlined in Miao and Wang (2007).

8.4 The Static Hiring Decision

8.4.1 Setup

As discussed in Section 8.2, we are going to model the shareholder's decision to hire a manager in two ways. The simplest way to model this is for the shareholder to examine a now-or-never (and irreversible) decision on whether or not to hire a manager at the initial date. In other words, if the shareholder delays they lose the opportunity to hire the manager. Given that the shareholder is making a now-or-never decision, the problem is static even though the valuation problem is dynamic.

In other words, the shareholder is choosing between $S^M(W, Y, t)$ (the value function from hiring a manager derived in Section 8.3.2) and $F^{CAPM}(Y, t)$ (derived in Appendix A). $F^{CAPM}(Y, t)$ is the market value of the project rights if the project is managed by the shareholder who incurs a monitoring cost of κdt while waiting to invest. Mathematically the static hiring decision is represented as

$$V^{\text{static}}(W_1, Y_1) = \max[S^M(W_1, Y_1, 0), F^{CAPM}(Y_1)] \quad (8.5)$$

As we have already calculated $S^M(W_1, Y_1, 0)$ and $F^{CAPM}(Y_1)$, calculating $V^{\text{static}}(W_1, Y_1)$ is straightforward.

8.4.2 Solution

The shareholder's value function, $V^{\text{static}}(W_1, Y_1)$, is plotted in Figure 8.4 using the modified base case parameter values.

This figure looks almost identical to Figure 8.1. Therefore it is useful to examine the difference between the payoff from hiring a manager ($S^M(W, Y, t)$) and the value function when the shareholder optimally chooses whether or

Figure 8.4: Value function ($V^{static}(W_1, Y_1)$) given optimal hiring decision

Calculated using the base case parameters

not to hire ($V^{static}(W_1, Y_1)$). This is shown in Figure 8.5.

Figure 8.5: $V^{static}(W_1, Y_1) - S^M(W_1, Y_1, 1)$

Calculated using the base case parameters

We can see from Figure 8.5 that the difference is in the region where W is large. This is because when W is large the manager exerts little to no effort. In other words if the shareholder hired a manager he would be receiving nothing in return. Therefore the shareholder is better off managing the project himself.

To more carefully examine the shareholder's decision we can use a contour plot to graphically show which decision is optimal for the shareholder in each region. For the base case this is presented in Figure 8.6.

Figure 8.6: *The Static Hiring decision: Base Case Parameters*

In this plot W is plotted on the x-axis and Y is plotted on the y-axis. The dark purple region is where a manager is hired who invests immediately, the blue region is where the shareholder invests immediately, the light blue region is where a manager is hired who then delays investing and the tan region is where the shareholder delays investing.

Figure 8.6 is very similar to the graphs of Chapter 7 examining the manager's dynamic investment decision. To understand the shareholder's decision in Figure 8.6, it is best to think about how the manager's optimal investment and effort decisions are affected. The general breakdown of the regions is easily understood in this light. When Y is large both the shareholder and the manager will invest immediately given that there is little to be gained from waiting in terms of exploiting cash flow volatility. Conversely, when Y is small both the shareholder and the manager would wait. With respect to W , for a lower W the manager exerts more effort and vice-versa.¹¹ Therefore we can make the following general characterizations of Figure 8.6:

¹¹See Chapter 7.

- Low W , High Y : Hire a manager who invests immediately
- High W , High Y : The shareholder does not hire a manager and invests immediately
- High W , Medium Y : The shareholder does not hire a manager and delays investment
- High W , Low Y : The shareholder hires a manager who delays investment
- Low W , Low Y : The shareholder hires a manager who delays investment

Moving past generalizations, there are three features of Figure 8.6 that warrant further discussion: the shape of the boundaries between hiring (light blue) and not hiring (tan) in the waiting region, the observation that the shareholder hires a manager in the “spike” region (which was discussed extensively in Chapter 7) and the reason for hiring a manager in the “high W , low Y ” region. With respect to the “spike”, in this region if the manager’s wealth remains unchanged or increases, he will exert little to no effort. On the other hand, if his wealth is lower he will exert more effort. Therefore from the shareholder’s perspective the downside risk from hiring a manager is significantly mitigated.

The fact that the shareholder hires a manager in the “high W , low Y ” region is initially counter-intuitive. When W is high the manager is unlikely to exert any effort upon investment while a low Y means that investment is unlikely to occur in the first place. The intuition behind this result is explained by the difference in monitoring costs between the manager and the shareholder. In this region the prospect of investment occurring is quite low yet the shareholder must continue to incur the monitoring cost because we have not introduced an explicit abandonment option. Because the shareholder avoids the monitoring cost by hiring a manager, the project is effectively

abandoned in this region through hiring a manager. This raises the question of the manager's participation constraint, i.e. would he actually accept the job? Because we have assumed that the manager bears no monitoring cost this is not an issue. However, if the manager did incur a monitoring cost this would need to be considered. The introduction of an explicit abandonment option and monitoring cost for the manager is left to future work.

With respect to the boundaries, we denote the vertical boundary between the tan and light blue regions the "left boundary" and the horizontal boundary the "bottom boundary". We will discuss what happens "along" each boundary, but the reader must keep in mind that the current analysis is static and thus the shareholder is not actually waiting to make a decision along these boundaries. Nonetheless it is a convenient way to describe what is happening.

The interesting feature of the left boundary is that as Y is reduced it slopes. This implies that for a lower Y the shareholder is more likely to hire a manager. Along this boundary, the shareholder is effectively making a decision as to whether it is more likely W and Y will end up in the region where it will be *ex post* optimal to hire (the purple region) or to invest himself (the blue region). The reason this boundary slopes is that as Y is decreased, the increase in Y required to end up in the blue region is greater than the fall in W required to end up in the purple region. Put simply, when Y is small, the probability that it will be *ex post* optimal to hire a manager increases and thus it is *ex ante* optimal to hire now. As we show in Section 8.4.3 the slope of the left boundary depends on the relative volatility of W and Y as this determines which decision (hiring/not hiring) is more likely to be optimal *ex post*.

8.4.3 Comparative statics

Having examined the base case, it is useful to examine how the shareholder's decision changes as we make (*ceteris parabus*) changes to the parameters of the model. In this section we conduct comparative statics for the same variables as we considered for the manager's dynamic investment decision (Chapter 7). To illustrate the effect of the different variables we simply redraw the decision plot presented in Figure 8.6 for various parameters. The graphs in this section can thus be interpreted in the same way as Figure 8.6.

We begin by examining the effect of the monitoring cost the shareholder incurs while waiting to invest (κ). Figure 8.7 graphically represents the shareholder's hiring decision for different levels of κ in $\{W, Y\}$ space. We can see from these graphs that as κ increases, the tan area in the contour plots becomes smaller and thus the shareholder is more likely to hire a manager prior to investment. This occurs because the manager can costlessly monitor the project and thus the shareholder avoids the monitoring cost when he hires a manager. In fact we can see that when κ is quite large and so is W , the shareholder will hire a manager even though the manager will exert no effort upon investing (and thus if Y was larger the shareholder would invest himself without a manager). This occurs because when the monitoring cost is high, the costs saved by avoiding the monitoring cost can outweigh the costs incurred from hiring a manager (i.e. giving them αY).

We now examine the effect of the manager's skill/the natural scope for cost savings by examining the effect of B (the investment cost floor) on the static hiring decision. Figure 8.8 presents the hiring decision for various levels of B . Recall that the potential efficiencies of the project are represented by $(A - B)$ and that in the base case $A = 100$. Therefore as B decreases the potential cost savings increase. The results of Figure 8.8 are not surprising. When $A = B$ there are no potential cost savings and thus a manager is only hired when Y is very low in order to avoid the monitoring cost (effectively abandoning the project). Conversely, when B is quite low (and thus the

Figure 8.7: κ and the static hiring decision

potential cost savings are high), a manager is hired over a much larger range of W and Y . This is shown by the dark purple and light blue areas in Figure 8.8 getting larger as B falls.

Figure 8.9 plots the shareholder's static hiring decision for various levels of the volatility of the market asset (σ_m). To understand Figure 8.9, we must first review the effect that σ_m has on the manager's dynamic investment decision. In Chapter 7 we found that as σ_m increases, the "spike" gets smaller.¹² The reason this occurs is that by making *ceteris parabis* changes to σ_m , we are changing the "Sharpe Ratio" of the market asset, the end result being that an increase in σ_m actually makes the manager's wealth less volatile because he

¹²See Figure 7.4.

Figure 8.8: B and the static hiring decision

decreases his holdings of the market asset. In other words, as σ_m increases the volatility of W decreases. We are now in a position to understand Figure 8.9.

As σ_m increases, the spike in Figure 8.9 dissipates because it is dissipating for the manager too. The more interesting feature of Figure 8.9 is the change in the slope of the boundary between the light blue (hire a manager who waits) and tan (do not hire a manager and wait) areas as σ_m increases. As discussed above, the reason that the “left boundary” is sloped is because as Y decreases, it becomes less likely the shareholder will invest and more likely that the manager will invest (because the manager can invest at lower cost). However, as σ_m is increased, the W becomes relatively less volatile than Y . Therefore it is less likely W will end up in a region where the manager

will exert effort. Thus at the margin the shareholder is less likely to hire a manager. As outlined previously, as $\frac{\sigma_W}{\sigma_y}$ falls, the “left boundary” becomes steeper.

Figure 8.9: σ_m and the static hiring decision

Figure 8.10 shows the effect of the volatility of the cashflow (σ_y) on the hiring decision. While increasing σ_y decreases the value of the cash flow, Figure 8.10 shows that an increase in σ_y also results in the shareholder not hiring a manager over a much greater range of Y . Using our previous terminology, the “left boundary” becomes steeper and the “bottom boundary” shifts down. The “bottom boundary” shifts down because when W is large, a high σ_y means the shareholder is much more likely to end up in the region where he invests himself (high W , high Y) than the region where the manager would invest and exert a high level of effort. The “left boundary” becomes steeper

for the same reason as when σ_m increases, $\frac{\sigma_W}{\sigma_Y}$ falls. While in the waiting region, the decision of whether or not to hire at the margin is determined by the relative volatility of W and Y . The more volatile Y is relative to W , the more likely it is that the shareholder will invest himself and thus he is less likely to hire a manager. On the other hand, when σ_Y is low, W is relatively more volatile than Y so at the margin the shareholder is more likely to hire a manager, since he is more likely to end up in the region where the manager exerts effort.

Figure 8.10: σ_Y and the static hiring decision

Figure 8.11 plots the static hiring decision for various levels of the manager's risk aversion (γ). In Chapter 7 we found that the "spike" gets "wider" in that it spans a greater range of W as γ gets smaller. This occurs because the curvature of the manager's utility function is less when γ is small. The

result of this is that the manager exerts effort over a greater range of W . It is therefore not surprising that the shareholder is more likely to hire a manager when γ is smaller. There is however another factor at play. Figure 8.11 also shows that the “left boundary” becomes steeper as γ increases. As we have discussed previously, the slope of this boundary is determined by the relative volatility of W and Y . While γ has no effect on σ_y , it does have an effect on the volatility of W through the manager's portfolio decision. As γ increases the manager is becoming more risk averse and thus all other things being equal he will alter his asset allocation away from risky assets. The result of this is that the manager's holdings of the market asset decrease and thus the volatility of his wealth decreases. That is, $\frac{\sigma_W}{\sigma_y}$ falls. Therefore, for the same reasons as discussed above, the “left boundary” is steeper as γ increases.

Figure 8.11: γ and the static hiring decision

In this chapter we are treating the manager's share of the firm (α) as exogenous and thus we also consider the effect of this on the hiring decision. Figure 8.12 plots the static hiring decision for various levels of α . Figure 8.12 shows two main effects of α on the hiring decision. The first and easiest to understand is that the "bottom boundary" in the waiting region shifts down as α increases. Recall that at this boundary the trade off being made is saving the monitoring cost (κ) and giving the manager a proportion α of the firm. As α increases the cost of hiring a manager increases and thus it is unsurprising that the shareholder is less likely to hire a manager.

The second effect is that the boundary between hiring and not hiring in the investment region (i.e. high Y) becomes sloped as α increases. In other words, the shareholder is less likely to hire a manager. That the shareholder is less likely to hire a manager as the cost of hiring increases is easy to understand. The slope however changes because the benefit of hiring a manager also decreases when α is large. In Chapter 5 we examined the manager's optimal effort decision and found that as Y increases the manager's optimal effort decreases.¹³ This is very similar to a wealth effect in that the utility benefit from exerting effort is smaller when Y is high. Because α scales Y , it also makes this "wealth" effect stronger. That is when α is high, the manager's effort decreases more quickly when Y is increased than when α is small. Therefore the benefit to the shareholder from hiring is decreasing at a faster rate when Y increases when α is large.

¹³See the discussion about Figure 5.3 in particular.

Figure 8.12: α and the static hiring decision (a) $\alpha = 0.05$ (b) $\alpha = 0.1$ (c) $\alpha = 0.2$ (d) $\alpha = 0.25$

8.5 The Dynamic Hiring Decision

8.5.1 Setup

We now turn to the general model where the shareholder can delay the decision on whether or not to hire a manager. The model of Section 8.4 is a specific case of this model where the shareholder must make an upfront irreversible decision on whether or not to hire a manager. In this situation the manager has the following possible actions to choose from prior to investment/hiring:

- Hire a manager and receive the payoff $S^M(W, Y, t)$ calculated in Section 8.3.2.
- Invest himself and receive the market value of the cashflow ($PCAPM$) denoted by Equation (6.1).
- Delay making a decision on whether or not to hire the potential manager/invest and incur the monitoring cost κdt .

In this situation the decision to wait encompasses both a decision to defer hiring a manager and a decision to delay investment. Thus neither of the shareholder's options are extinguished by waiting.

We assume that if the shareholder has not hired a manager, he incurs a monitoring cost of κdt while waiting. This represents the cost/effort the shareholder must incur to monitor the project when he does not have manager. Without this assumption the shareholder would never hire a manager prior to investment which would appear to be an unrealistic result.¹⁴

With respect to the payoff from the different decisions the shareholder can make, we have already calculated $PCAPM$ and $S^M(W, Y, t)$. Therefore the

¹⁴This stems from the fact that hiring a manager prior to investment would destroy the shareholder's option to invest himself without providing the shareholder any other benefit.

only additional calculation we need to make is the payoff from waiting. In this situation, because the shareholder incurs a cost of κdt while he waits, the shareholder's general value function becomes

$$S^D(W, Y, t) = e^{-\beta dt} E[S^D(W(t+dt), Y(t+dt), t+dt)] - \kappa dt \quad (8.6)$$

which can be expanded out to

$$\beta S^D(W, Y, t) = E[S_t^D + S_w^D dW + S_y dY + \frac{1}{2} S_{ww}^D dW^2 + \frac{1}{2} S_{yy}^D dY^2 + S_{wy}^D dY dW] - \kappa dt \quad (8.7)$$

which is then simplified to

$$\begin{aligned} \beta S^D(W, Y, t) = & S_t^D + S_w^D (rW + \pi(\mu_m - r) - C) + S_y^D \mu_y + \frac{1}{2} S_{ww}^D (\pi\sigma_m)^2 \\ & + \frac{1}{2} S_{yy}^D (\phi^2 + \rho^2 \sigma_y^2) + S_{wy}^D (\pi\rho\sigma_y\sigma_m) - \kappa \end{aligned} \quad (8.8)$$

We can thus formally represent the shareholder's value function as

$$S^D(W, Y, t) = \begin{cases} P^M[\alpha] & \text{if the shareholder invests} \\ S^M(W, Y, t) & \text{if the shareholder hires a manager} \\ \text{Solution to (8.8)} & \text{if the shareholder waits} \end{cases}$$

8.5.2 Solution

Similarly to Section 8.3, $S^D(W, Y, t)$ is solved using standard finite difference techniques which are not reproduced here. A good starting point for a discussion of the results from the dynamic model is a comparison to the static model. This allows us to make a *ceteris paribus* evaluation of the effect of flexibility in the hiring decision. Figure 8.13 shows the contour plots for the dynamic and static hiring decision when investment can be delayed.

Three differences are immediately apparent between the dynamic and static hiring decisions; the “left boundary” has shifted left, the “bottom boundary” has shifted down and the “spike” extends much further up creating a “ridge”

Figure 8.13: *The dynamic vs static hiring decision: the effect of flexibility*

between the regions where it is optimal for the shareholder to invest now or hire a manager who invests now.

This “ridge” is particularly interesting as it represents a saddle point of sorts. This is because along the ridge the shareholder is waiting but will either invest himself or hire a manager if there is a shock to W . Intuitively this occurs because if the shareholder hired a manager, the manager would invest now and exert no effort. Thus in the absence of the ability to wait it would be optimal for the shareholder to invest himself. On the other hand, when delay is available the shareholder can wait and see if there is a shock to W that would induce the manager to exert effort upon investment. Thus the shareholder waits in the hope of a larger future payoff even though in a now-or-never setting it would be optimal to hire a manager or invest now himself.

The movement of the left boundary is also interesting. Notice that this results in the shareholder delaying hiring in the region below the “spike”. In this region the shareholder delays hiring because the manager still has a “choice” over his effort level. This choice occurs because W is not low enough that the

manager always exerts “maximum” effort,¹⁵ nor is it high enough that the manager always exerts no effort. Therefore in this region there is uncertainty over the effort level the manager will exert upon investment if hired. Thus it makes intuitive sense that the shareholder delays the decision of whether or not to hire.

To understand the shift in the bottom boundary, it is useful to recall that this boundary represents the effective abandonment region - the shareholder hires a manager to avoid the monitoring cost, not because he expects investment to occur. With that in mind, the downward shift in the lower boundary means that the shareholder is less likely to abandon the project when waiting is allowed. Intuitively this is easy to understand, the shareholder does not want to hire a manager (for abandonment purposes) and find out subsequently not only that investment is worthwhile, but that it would have been optimal *ex post* to not hire a manager.

There is a common theme running through each of these explanations, the shareholder makes an *ex ante* decision to delay hiring a manager because of uncertainty over whether it will have been optimal to hire a manager *ex post*. In other words the “bad news” principle discussed in Chapter 7 can be appealed to in order to explain the shareholder’s behaviour. It is thus worth repeating yet again the description of the “bad news” principle given by Dixit and Pindyck (1994):

“...it is the ability to avoid the consequences of “bad news” that leads us to wait.”

¹⁵Recall that “maximum” in this context refers to exerting enough effort to drive the investment cost very close to B .

8.5.3 Comparative statics

Given the similarity between this model and the model where hiring cannot be delayed, many of the comparative statics are qualitatively the same. The key differences introduced by allowing hiring to be delayed are the “ridge” and the shift in the left boundary.¹⁶ The comparative statics analysis will thus focus on understanding these areas in greater detail.

In this context the volatility of the market asset (σ_m) has a significant effect on both areas. Figure 8.14 plots the shareholder's investment/hiring decision for various levels of σ_m . From this graph we can see that as σ_m increases the “ridge” disappears and the left boundary shifts right. With respect to the “ridge”, recall that in the absence of the ability to delay hiring, in this region it would be optimal to either invest now or hire a manager who would invest now. Thus the shareholder is waiting along the ridge in case there is a negative shock to W which will induce the manager to exert effort. In this context the effect of σ_m is quite simple. As we discussed in Chapter 7, increasing σ_m actually decreases the volatility of W .¹⁷ Therefore as σ_m increases, shocks to W become smaller. This means that the benefit of delaying investment is reduced since any shock will be unlikely to induce the manager to exert a large amount of effort. Therefore along the ridge we have a situation similar to that in Chapter 7 where an increase in σ_m spurs action rather than delays it.

The intuition behind the shift in the left boundary is similar to that for the “ridge”. The rightward shift of the boundary means that prior to investment the shareholder is more likely to commit to hiring a manager. This occurs because as σ_m increases, the magnitude of the “bad news” the shareholder can receive reduces. In other words, because the manager's wealth is not very

¹⁶The comparative statics analysis of the bottom boundary are little different from when hiring cannot be delayed.

¹⁷All other things being equal, an increase in σ_m decreases the Sharpe ratio of the market asset and thus the manager reduces his holdings of the market asset.

volatile, if his wealth is low prior to investment, it will also likely remain so at the time investment occurs.

Figure 8.14: σ_m and the dynamic hiring decision

The next parameter that sheds light on the “ridge” and the left boundary is the investment cost floor B . Figure 8.15 plots the shareholder’s hiring/investment decision for various levels of B . Beginning with the “ridge”, we can see from these graphs that as B decreases (and thus the benefit of effort increases) the ridge becomes wider. Along the “ridge” the shareholder is delaying investment in case W falls and thus a higher payoff can be obtained by hiring a manager. As B is decreased the benefit of a manager exerting effort increases. Therefore it is intuitively sensible that for a lower B the shareholder is willing to defer hiring over a wider range.

Figure 8.15: B and the dynamic hiring decision

Perhaps more interesting is the behaviour around the left boundary as B decreases. Notice that the “bottom boundary” stays a fixed distance from the threshold where the shareholder invests himself (the boundary between the tan and dark blue areas). This occurs because the bottom boundary is the abandonment threshold when W is large (and thus the manager would exert no effort) and thus is unrelated to B . On the other hand, for a lower B the threshold where the manager invests when W is low (and thus the manager exerts “maximum” effort) shifts downwards. This results in an “overhang” of sorts whereby the shareholder waits for intermediate values of W but will hire a manager if W rises (to abandon) or if W falls (in anticipation of the manager exerting effort and investing). To confirm this logic we can examine what happens to the “overhang” when ρ changes as

this determines the relationship between changes in W and Y . Figure 8.16 plots the shareholder's hiring decision for $\rho = \{0.5, -0.5\}$ when $B = 70$.

Figure 8.16: ρ and the dynamic hiring decision when $B = 70$

Figure 8.16 shows us that when ρ is positive the “overhang” extends into the low W region whereas when ρ is negative the overhang disappears. The intuition behind this is quite simple: when ρ is positive, if Y increases W is also likely to increase. Thus in the region where investment is being delayed, the shareholder is more likely to end up in the region where it is optimal to invest himself (i.e. high Y , high W) and thus he is less likely to hire a manager prior to investment.

8.6 Impact on shareholder wealth

The results of this chapter so far have focused on the shareholder's *decision* (i.e. hire a manager, delay hiring/investing or invest himself) as certain parameters are changed. This however does not answer the question of whether or not the shareholder is better off if a parameter takes one value over another. In particular, while we do not directly model the decision in this

thesis, analysing the shareholder's *payoff* allows us to partially answer the question "what type of manager would a shareholder want?". In addition, if a managerial parameter has little effect on the shareholder's payoff, if there are search costs the shareholder will be relatively indifferent to the level of this parameter that their manager has.

As it turns out the effect of the different parameters is qualitatively quite similar whether or not the shareholder is able to delay the hiring decision. We will thus examine the effect of managerial parameters on the shareholder's payoff jointly for both models.

Given the complexity of our numerical solutions, the way we will get at this question is to simply calculate the shareholder's payoff for a "high" and "low" value of each parameter of interest and then look at the difference between these payoffs in $\{W, Y\}$ space.

The first parameter we will examine is the investment cost floor (B). Recall that this parameter determines the achievable cost savings and thus is a measure of not only the manager's "skill", but also the natural scope for cost savings in the project. Figure 8.17 plots the difference in the shareholder's payoff for $B = 70$ and $B = 90$. In this graph the payoff when $B = 90$ is subtracted from the payoff when $B = 70$ and thus positive values indicate that the shareholder is better off when $B = 70$.

The immediate observation from Figure 8.17 is that when W is relatively high there is no difference in the shareholder's payoff. This occurs because the manager would not exert any effort in this region and thus the amount by which the manager could reduce costs does not impact the shareholder's payoff. Similarly in the "low W , high Y " area the difference is flat and equals 18. In this region the manager is exerting enough effort to drive the investment cost very close to B and thus the difference in the payoff is simply $1 - \alpha$ (0.9) multiplied by the change in B (20).¹⁸

¹⁸For different parameter values or if Y is sufficiently large the difference begins to

Figure 8.17: *Difference in shareholder's payoff: $\{B=70, B=90\}$*

These graphs calculate the shareholder's payoff when $B=70$ and then subtracts from that the shareholder's payoff when $B=90$. Thus a positive value indicates that the shareholder is better off when $B=70$ and vice versa.

When Y is very low there is no difference because it is unlikely that investment will occur. To explain why the difference increases with Y when W is low, we need to remember that when W is low the manager will generally exert “maximum” effort. Therefore as Y increases the difference becomes positive because the *investment* option is much more valuable when $B = 70$ than when $B = 90$. That is, the expected payoff from the waiting to invest is higher when the investment cost is lower.

The “bump” in the middle is the unusual feature of this graph. This bump is related to the *effort*-related delay option. In the region where the manager has a “choice” over his level of effort, the shareholder's option value of waiting (to see if the manager exerts more effort) is much greater when the potential cost reductions are greater (i.e. B is lower).

The next parameter we consider is λ , another measure of the manager's “skill”. This parameter determines how much effort the manager must exert to reduce to zero since effort is a decreasing function of Y as shown in Chapter 5.

to drive the investment cost close to B (where a higher value of λ results in less effort being required to reach B). Recall that this parameter has the exact opposite effect to θ (the utility cost of exerting effort) and thus we only need to examine one to infer the effect of the other. Figure 8.18 plots the difference in the shareholder's payoff between $\lambda=1$ and $\lambda=0.5$.

Figure 8.18: *Difference in shareholder's payoff: $\{\lambda=1, \lambda=0.5\}$*

These graphs calculate the shareholder's payoff when $\lambda=1$ and then subtracts from that the shareholder's payoff when $\lambda=0.5$. Thus a positive value indicates that the shareholder is better off when $\lambda=1$ and vice versa.

To understand Figure 8.18 we need to remember the effect that λ has on the manager's timing/effort decision in Chapter 7. The key result in this context is that increasing λ simply moves the investment "spike" to the right, without changing its shape (i.e. the manager exerts "maximum" effort over a wider range of W). This rightward shift in the "spike" is effectively what we are witnessing here. The area where the "ridge" occurs in Figure 8.19 is an area where investment would occur immediately and the manager would not exert "maximum" effort when $\lambda = 0.5$ yet does when $\lambda = 1$. The ridge disappears to the left because in this area maximum effort is exerted for both parameter values while to the right no effort is exerted for both parameter values. The "spike" that occurs for lower values of W occurs because the region where

the effort option has value has shifted right in W space.

The next managerial parameter of interest is the manager's level of risk aversion (γ). Figure 8.19 plots the difference in the shareholder's payoff between $\gamma = 0.5$ and $\gamma = 2$, where higher levels of γ represent greater risk aversion.

Figure 8.19: *Difference in shareholder's payoff: $\{\gamma=0.5, \gamma=2\}$*

These graphs calculate the shareholder's payoff when $\gamma=0.5$ and then subtracts from that the shareholder's payoff when $\gamma=2$. Thus a positive value indicates that the shareholder is better off when $\gamma=0.5$ and vice versa.

There are two things occurring in this graph, a positive “ridge” in the region when Y is large and a negative “dip” when W and Y are relatively lower. The ridge occurs because γ has a negative effect on effort and thus a lower value of γ results in the manager exerting effort (and thus being hired) for a wider range of W . To understand the “dip” it is worth reexamining Figure 8.11 which is reproduced in Figure 8.20 with only the figures for $\gamma = \{0.5, 2\}$.

Figure 8.20 illustrates that the negative dip (i.e. the shareholder getting a higher payoff from a more risk averse manager) occurs because the transition from no effort to “maximum” effort is much less sudden as W changes when γ is low. This results in an area (the “dip”) where when γ is large the manager exerts “maximum” effort and invests whereas with a lower γ he is waiting to

Figure 8.20: γ and the static hiring decision

invest or exerting a lower level of effort and investing.

Unlike the model of Chapter 6, in this model we have treated the proportion of the firm owned by the manager (α) as exogenous. The interpretation of α is therefore the market price of hiring a manager. In this context it is therefore worthwhile considering whether or not the shareholder is better off if the manager desires a high or low share of the firm (holding all other parameters constant). Figure 8.21 plots the difference in the shareholder's payoff between $\alpha=0.2$ and $\alpha=0.05$.

Figure 8.21 shows that if W is sufficiently large then the shareholder is indifferent between a high and low α . This occurs because when W is relatively large the manager will exert no effort and thus the shareholder does not hire a manager. The two interesting features of Figure 8.21 are the positive "ridge" around where the investment spike occurs for the manager and that the difference becomes increasingly negative as Y increases when W is relatively low. The explanation for the increasing negative difference is relatively simple. We know from Chapter 5 that as Y increases the manager's optimal level of effort decreases. Given that α effectively scales Y , an increase in α

Figure 8.21: *Difference in shareholder's payoff: $\{\alpha=0.2, \alpha=0.05\}$*

These graphs calculate the shareholder's payoff when $\alpha=0.2$ and then subtracts from that the shareholder's payoff when $\alpha=0.05$. Thus a positive value indicates that the shareholder is better off when $\alpha=0.2$ and vice versa.

amplifies this effect.¹⁹ Therefore, at a sufficiently low W , when α is higher an increase in Y causes a greater reduction in effort relative to when α is small, leaving the shareholder worse off.

The “ridge” is slightly more complicated. To understand why this occurs we will reexamine the manager's investment decision from Chapter 7. Figure 8.22 plots the manager's investment decision for $\alpha=0.2$ and $\alpha=0.05$.

Recall that the dark purple area is where the manager invests and the light blue area is where the manager waits but would invest in a “now or never” world. Figure 8.22 therefore shows that when $\alpha=0.2$ the area where investment occurs is now larger (i.e. the light blue waiting region has shrunk). This is the reason behind the “ridge” in Figure 8.21, it represents the region in Figure 8.22 where investment occurs when $\alpha=0.2$ but the manager waits when $\alpha=0.05$.

¹⁹This is the same reason for the curved boundary in Figure 8.12.

Figure 8.22: *The manager's investment decision: $\{\alpha=0.2, \alpha=0.05\}$*

The final managerial parameter we consider is the shareholder's monitoring cost (κ). While this may not directly be a manager-specific parameter, it does represent the cost the shareholder avoids by hiring a manager. Put another way, it is the manager's relative monitoring cost advantage. Figure 8.23 plots the difference in the shareholder's payoff between $\kappa=0.1$ and $\kappa=0.01$.

Figure 8.23 is simple relative to other graphs we have examined in this section. This occurs because κ only affects the shareholder's payoff in regions where he is waiting to make a decision. Therefore if the shareholder invests himself or hires a manager, κ has no effect on the shareholder's payoff. Unsurprisingly in the areas where the shareholder is waiting,²⁰ his payoff is higher when κ is lower.

²⁰Which of course change as κ changes.

Figure 8.23: *Difference in shareholder's payoff: $\{\kappa=0.1, \kappa=0.01\}$*

These graphs calculate the shareholder's payoff when $\kappa=0.1$ and then subtracts from that the shareholder's payoff when $\kappa=0.01$. Thus a positive value indicates that the shareholder is better off when $\kappa=0.1$ and vice versa.

8.7 Conclusion

In this chapter we examined the shareholder's decision whether or not to hire a manager who has the ability to exert effort to reduce the cost of a project whose timing is flexible. The hiring decision was examined in two different contexts:

- the choice of whether or not to hire is a “now or never” decision; and
- the hiring decision can be delayed.

In our model the shareholder incurs a monitoring cost while waiting to invest. Because the shareholder avoids the monitoring cost (but incurs the cost αY) when a manager is hired, as the monitoring cost increases the shareholder hires a manager over a wider range of parameters. In the limit, if the monitoring cost is large enough the shareholder always hires a manager prior to investment.

The key result in the “now-or-never” model is that prior to investment the relativity between the volatility of W and Y is crucial as this affects what will be *ex post* optimal from the shareholder’s perspective. At the margin on the “left boundary”, the more volatile W is relative to Y , the more likely it is that it will be optimal to hire a manager. This is because it is more likely that we will end up in the area where the manager exerts effort.

The decision at the margin along the “bottom boundary” is driven by the implicit abandonment option available to the shareholder. Because the shareholder incurs a monitoring cost, he can effectively abandon the project by hiring a manager.²¹ Because an increase in the volatility of Y makes it more likely that the project will eventually be profitable, the bottom boundary shifts downwards as σ_y increases.

Turning to the model where hiring can be delayed, the key difference is that in regions where the manager has a “choice” over how much effort he exerts, the ability to delay hiring will often cause the shareholder to delay hiring relative to a “now or never” world. This can be explained using a re-expression of the famous “bad news” principle. Essentially, the shareholder will delay hiring now to avoid hiring a manager who subsequently exerts no effort. The other interesting distinction between these two models is that when Y is very large and immediate investment is optimal in the absence of the ability to delay hiring, we find a “ridge” where the shareholder will delay hiring. The shareholder does this in the hopes of an effort-inducing negative shock to the manager’s wealth. In simple terms the shareholder is willing to wait another day in the hope of higher payoff.

Examining the shareholder’s *decision* under different parameters only tells one side of the story. To understand what type of manager the shareholder desires we must examine the payoff to the shareholder for different levels

²¹Note that the lack of explicit abandonment option is a weakness of this framework. However, in this region the shareholder loses nothing by giving the project to the manager since he would have abandoned the project anyway.

of the relevant managerial parameters. The general results of this exercise are unsurprising, the shareholder will generally prefer a manager who is less risk averse, more skilled (low B and high λ) and who requires a lower level of firm ownership (α). However, when one moves past generalisations, there are situations where the opposite is true and also situations where the shareholder is indifferent.

Part III

Direct wealth effects

Chapter 9

CRRA Utility with GBM Cash Flow

9.1 Introduction

In Chapter 3 we discussed two possible utility functions that could be used for our analysis: CARA and CRRA. CARA was used for the “clean path” because the manager’s valuation of the cashflow does not depend on his financial wealth and thus the rest of the analysis was tractable. The purpose of this chapter is to go back to the first step of the clean path and see what difference *direct* wealth effects¹ have on the initial valuation problem. This chapter is thus a parallel analysis to Chapter 4 using CRRA utility instead of CARA.²

Because the manager’s valuation will depend directly on his financial wealth, the dimensionality of the problem has increased (i.e. the valuation now depends on W as well as Y). This changes the analysis of Chapter 4 in two key areas:

¹This is in contrast to the *indirect* wealth effects we witness through the manager’s effort decision in the “clean path”.

²As will become clear shortly, we also use a GBM cashflow instead of SBM to be consistent with the literature on CRRA valuations.

- there is no closed form solution to the manager's valuation problem and thus the problem must be solved numerically; and
- The implicit valuation (IV) concept used in Chapter 4 is not equal to the utility indifference (UI) valuation and thus there are two measures of value.

In simple terms, the IV valuation is the marginal valuation and the UI valuation is the average valuation. The distinction between these two concepts is discussed in more detail in section 9.3. A lack of a closed form solution and two measures of value complicates the analysis significantly relative to that carried out in Chapter 4. In Chapter 4 there was a closed form solution and only one measure of value needed to be considered since the IV and UI valuations were equivalent.

The numerical solution also turns out to be relatively complex and thus, unlike Chapter 4, a significant portion of this chapter is devoted to solving the model. Part of this complexity lies in the fact that the UI and IV valuations are not equivalent meaning there are multiple choices for numerical boundary conditions.

The only paper to have presented numerical solutions to the CRRA problem is Munk (2000). Munk examined the valuation problem of a CRRA investor who has a non-traded GBM cash flow that can be partially hedged by the market asset. However, the numerical method Munk implements has some shortcomings which we seek to address. The use of a numerical (as opposed to theoretical grounded) boundary condition is the weakness of Munk's analysis. We therefore extend the work of Munk (2000) by using a numerical finite difference method which uses theoretically grounded conditions at the upper boundary as opposed to the numerical condition imposed in Munk (2000). The analytical results of Munk (2000), Duffie, Fleming, Soner, and Zariphopoulou (1997) and Koo (1998) provide a menu of alternative conditions that can be imposed at the model's upper boundary. We compare the

accuracy of these results and contrast them with Munk's results in order to select the most appropriate boundary condition. The use of theoretically grounded boundary conditions significantly alleviates the problems encountered by Munk at the upper boundary. While the selection of a small value for the upper boundary (i.e. a smaller computational grid) introduces more error into the computations, we find that selection of an appropriate boundary condition can significantly mitigate this error. This allows the use of smaller, less refined grids in computation.

Munk's model also considers an agent who owns the entire cash flow, whereas we are interested in the situation where the manager is constrained to own an exogenous portion of the cash flow (α). We can thus think of Munk's model as a special case of the partial ownership model when $\alpha = 1$. As it turns out, because of the way the problem reduces, the solution method is actually no different when $\alpha < 1$. We thus focus on improving the solution method for Munk's model first. Then, with the solution method and preferred boundary condition in place, we examine the case where the manager only owns part of the firm. This allows us to examine the impact that changes in the manager's ownership level have on his valuation of the firm in a wealth dependent setting.

This chapter differs from the rest of the thesis in that the majority of it is devoted to the "journey" of solving the model. The complexity of the solution method and the differences between this model and that of Chapter 4 do however mean that this journey is interesting in its own right.

The rest of this chapter is laid out as follows: Section 9.2 sets up the model, Section 9.3 discusses the different measures of value relevant to our analysis, Section 9.4 draws on the literature to discuss the limiting behaviour of the value function (which is relevant for determining boundary conditions), Section 9.5 outlines our numerical solution method, Section 9.6 discusses our results and contrasts them to those obtained by Munk (2000), Section 9.7 analyses the implication of our results for the discount rate used by a man-

ager who owns part of the firm, Section 9.8 introduces effort to the model and Section 9.9 summarises the results of this chapter.

9.2 The Model

By substituting the CRRA utility function $U(C) = C^\gamma$, and the process for the GBM cashflow (Equation (3.5)) into Equation (3.9) we get a highly non-linear second order partial differential equation for the manager's HJB. The choice of CRRA utility means that the solutions will be wealth dependent and thus standard numerical solution procedures are difficult to implement³. We follow the approach of Duffie, Fleming, Soner, and Zariphopoulou (1997) and show that the problem can be reduced to a single state variable $Z = \frac{W}{\alpha Y}$. We also follow Duffie, Fleming, Soner, and Zariphopoulou (1997) and impose the restriction that $Z > 0$.⁴ To reduce the problem we recognize that we can write the value function as $J(W, Y) \equiv (\alpha Y)^\gamma F[Z]$.⁵ Substituting this in gives the following reduced form for the HJB

$$\begin{aligned} \beta F[Z] = & \max_{\bar{C}} [\bar{C}^\gamma (\alpha Y)^{-\gamma} - \frac{\bar{C} F'(Z)}{\alpha Y}] + \frac{1}{2} Z^2 F''(Z) \sigma_y^2 + F(Z) \left(\frac{1}{2} (\gamma - 1) \gamma \sigma_y^2 + \gamma \mu \right) \\ & + (Z (r - \mu - (\gamma - 1) \sigma_y^2) + 1) F'(Z) \\ & + \max_{\bar{\pi}} \left[\bar{\pi} \frac{1}{2(\alpha Y)^2} ((\bar{\pi} \sigma_m^2 - 2W \rho \sigma_y \sigma_m) F''(Z) - 2\alpha Y (r - \mu_m - (\gamma - 1) \rho \sigma_m \sigma_y) F'(Z)) \right] \end{aligned}$$

³This is because there are two state variables; W_t and Y_t .

⁴Note that no restrictions were placed on the manager's financial wealth in the CARA models of the clean path given that the variable Z was not necessary.

⁵For details of this transformation see Duffie, Fleming, Soner, and Zariphopoulou (1997).

Following Munk (2000) we define the following policy functions

$$\xi = \frac{\bar{C}}{\alpha Y} - 1$$

$$\psi = \frac{\bar{\pi}}{\alpha Y}$$

Substituting these into the HJB equation allows us to write it as a second order non-linear ordinary differential equation with a single state variable and two policy functions

$$\begin{aligned} \beta F[Z] = & \max_{\xi} [(\xi + 1)^{\gamma} - (\xi + 1)F'(Z)] + \frac{1}{2}Z^2F''(Z)\sigma_y^2 + F(Z) \left(\frac{1}{2}(\gamma - 1)\gamma\sigma_y^2 + \gamma\mu \right) \\ & + (1 + Z(r - \mu - (\gamma - 1)\sigma_y^2))F'(Z) \\ & + \max_{\psi} \left[\psi(\mu_m - r - (1 - \gamma)\rho\sigma_m\sigma_y)F'(Z) + \left(\frac{1}{2}\psi^2\sigma_m^2 - \rho\sigma_m\sigma_yZ\psi \right)F''(Z) \right] \end{aligned}$$

Again following Duffie, Fleming, Soner, and Zariphopoulou (1997) we make use of the following notation to simplify the HJB

$$\begin{aligned} \hat{\beta} &= \beta - \mu_m\gamma + \frac{1}{2}\sigma_y\gamma(1 - \gamma) \\ k_1 &= \mu_m - r - (1 - \gamma)\sigma_m\sigma_y\rho \\ k_2 &= r - \mu + \sigma_y^2(1 - \gamma) \end{aligned}$$

which gives the following reduced form for the HJB and policy functions

$$\begin{aligned} \hat{\beta}F[Z] = & \max_{\xi} [(\xi + 1)^{\gamma} - (\xi + 1)F'(Z)] + \frac{1}{2}F''(Z)\sigma_y^2Z^2 + F'(Z)(1 + k_2Z) \\ & + \max_{\psi} \left[\left(\frac{1}{2}\psi^2\sigma_m^2 - \rho\sigma_m\sigma_yZ\psi \right)F''(Z) + \psi k_1F'(Z) \right], \\ \bar{\xi}[Z] = & \left(\frac{F'[Z]}{\gamma} \right)^{\frac{1}{\gamma-1}} - 1, \\ \bar{\psi}[Z] = & \frac{\sigma_y\rho Z}{\sigma_m} - \frac{k_1F'[Z]}{\sigma_m^2F''[Z]}. \end{aligned} \tag{9.1}$$

Although we have defined Z differently, this is the exact system of equations of Munk's model. Therefore once we improve the numerical solution to his model, we can take advantage of the duality of the problems to analyse our partial ownership case.

9.3 Measures of Value

In Chapter 4 (and the rest of the clean path) we only had to consider one measure of value. As noted by Miao and Wang (2007) in the context of the CARA framework used in Part II of this thesis:

The two interpretations of $G(x)$ - the certainty-equivalent wealth and the implied option value - are the same in our setup. This is due to the absence of the wealth effect under CARA utility. We will thus use certainty-equivalent wealth (from the consumption literature perspective) and implied option value (from the investment literature perspective) interchangeably throughout the remainder of the paper.

Given that CRRA has wealth effects, this is no longer the case and we thus have to consider two different measures of value (the IV and UI values). As mentioned previously, the IV valuation is effectively the manager's marginal valuation of the cashflow while the UI valuation is his average valuation.

The reason it matters that the UI and IV valuations are not equivalent is that the model must be solved numerically, and the UI and IV valuations are candidates for theoretically grounded boundary conditions. In addition, once the model is solved, we want to analyse the manager's valuation of the cash flow. With that in mind, it is worth examining in detail the definition of each valuation and how they relate to each other.

The UI concept is what Miao and Wang (2007) refer to as the "certainty-equivalent" valuation. It is defined as the least increase in initial wealth the manager would require to forgo the entire income stream. It therefore gives a natural measure of the value of the *entire* income stream. The IV valuation is slightly different in that it is the value of the cashflow implied by the manager's consumption and asset allocation decisions.

The implicit value is derived by Koo (1998). This was done by showing that the value function and optimal policies can be written as

$$\begin{aligned} J(W, Y) &= A(Z)^{\gamma-1} (W + B(Z)Y)^\gamma \\ C(W, Y) &= A(Z)(W + B(Z)Y) \\ \pi(W, Y) &= \frac{\sigma_m - r}{\sigma_m^2} \frac{W + B(Z)Y}{1 - \gamma + B'(Z)} + \frac{\sigma_y \rho}{\sigma_m} \left(W - \frac{1 - \gamma}{1 - \gamma + B'(Z)} (W + B(Z)Y) \right) \end{aligned}$$

Where $B(Z)$ and $A(Z)$ are defined as

$$B(Z) = \frac{J_Y(W, Y)}{J_W(W, Y)} \quad A(Z) = \left(\frac{J(W, Y)}{(W + B(Z)Y)^\gamma} \right)^{\frac{1}{\gamma-1}}$$

Based upon the form of the value function, Koo interprets the implicit valuation as

$$V^I = B(Z)Y$$

Thus he describes the manager's "Accounting Total Wealth" as $W + V^I$. Given that V^I is by definition the marginal value of the income stream it is sensible that this is the valuation that drives the manager's consumption/savings decision.⁶ However, as Munk (2000) points out, this is an unnatural measure of value, especially when one is looking at decisions to buy or sell an income stream as one would want to examine the total value of the stream as opposed to the marginal values.

Munk notes that the UI valuation provides a more natural valuation of the entire cashflow and thus prefers it to the IV value. He shows that the UI valuation can be derived as

$$\begin{aligned} V^{UI} &= B^*(Z)Y, \\ B^*(Z) &= A^{\frac{1-\gamma}{\gamma}} F(Z)^{\frac{1}{\gamma}} - Z \end{aligned} \tag{9.2}$$

For the case of CRRA, these two valuations do not in general coincide. We therefore use the complete markets valuation as benchmark to compare the

⁶Given that optimality is based upon the equalization of the marginal utility from consumption with the marginal utility of deferring consumption.

two value. This is the valuation the agent would reach in the absence of liquidity constraints and if markets were complete. Koo (1998) derived this measure by solving for the manager's Certainty Equivalent Present Value (CEPV) of the income stream in the unconstrained, complete markets case.

$$V^C = \frac{1}{\lambda} Y$$

$$\lambda \equiv r - \mu_y + \frac{(\mu_m - r)\sigma_y}{\sigma_m} > 0 \quad (9.3)$$

As we are interested in asset pricing and investment decisions it is natural that we focus on the utility indifference value as this determines the amount the manager is willing to pay for the income stream. However, in the limit as Z becomes large, the distinction does not matter as Munk (2000) has proved that the two valuations converge. He does this by showing that the implicit income multiplier can be written as

$$B(Z) = A(Z)^{\frac{1-\gamma}{\gamma}} F(Z)^{\frac{1}{\gamma}} - Z \quad (9.4)$$

Using the fact that Koo (1998) proved that the following limits hold

$$\lim_{Z \rightarrow \infty} A(Z) = A \quad (9.5)$$

$$\lim_{Z \rightarrow \infty} B(Z) = \frac{1}{\lambda}, \quad (9.6)$$

Munk showed that the following holds

$$\lim_{Z \rightarrow \infty} B(Z) = B^*(Z).$$

Therefore it must be the case that

$$\lim_{Z \rightarrow \infty} B^*(Z) = \frac{1}{\lambda}. \quad (9.7)$$

Thus we can see that in the limit both measures of value converge to the complete markets valuation.

9.4 Limiting Behavior of the Value Function

Given the model is going to be solved numerically and that we aspire to have theoretical relationships governing the behavior at the boundaries, it is essential that we examine the behavior of the model in the limit as we approach the imposed boundaries.

As we have imposed the restriction that the investor's financial wealth must be non-negative at all times,⁷ this is equivalent to requiring that $Z \geq 0$ at all points in time given that $Z \equiv \frac{W}{Y}$. It is therefore necessary to examine the behavior of the value function as $Z \rightarrow 0$. This has been done by Duffie, Fleming, Soner, and Zariphopoulou (1997) who showed that the following limit holds

$$\lim_{Z \rightarrow 0} ZF''[Z] = 0$$

The other limit we need to consider is when Z gets very large. We therefore must consider the behavior of the model as $Z \rightarrow \infty$. This is equivalent to looking at the behavior as $Y \rightarrow 0$ and/or $W \rightarrow \infty$. It is important to note that because we have assumed that Y_t follows GBM it must be the case that once the income process hits zero, the investor behaves as if he is in a Merton (1969, 1971) world since the income process is zero thereafter. Therefore it must be the case that

$$J(W, 0) = J^M(W)$$

For the general case when $Y \neq 0$, as $Z \rightarrow \infty$ Duffie, Fleming, Soner, and Zariphopoulou (1997) show that the optimal policies and value function become asymptotically equivalent to the Merton no-income case. Thus their

⁷See Section 9.2.

Theorem 6 states that as $\frac{W}{Y} \rightarrow \infty$

$$\begin{aligned} J(W, Y) &\rightarrow J^M(W) \\ C(W, Y) &\rightarrow \bar{C}^M(W) \\ \pi(W, Y) &\rightarrow \bar{\pi}^M(W) \end{aligned}$$

which essentially states that as the ratio of wealth to income becomes very large, the investor will behave as if he is not receiving an income stream.

Koo (1998) puts forward a slightly different proposition. As discussed previously, he shows that the implicit income multiplier ($B(Z)$) converges to the complete markets income multiplier (Equation (9.6)) and that the consumption multiplier ($A(Z)$) converges to the Merton (1969) consumption multiplier (Equation (9.5)). Koo shows that the complete markets value function can be written as

$$J^C(W, Y) = A^{\gamma-1} \left(W + \frac{1}{\lambda} Y \right)^\gamma$$

Given the form of the value function, this shows that the value function converges to the complete markets case

$$\lim_{Z \rightarrow \infty} A(Z)^{\gamma-1} (W + B(Z)Y)^\gamma = A^{\gamma-1} \left(W + \frac{1}{\lambda} Y \right)^\gamma \quad (9.8)$$

It is interesting to note that while very similar, this is slightly different from the proposition put forward by Duffie, Fleming, Soner, and Zariphopoulou (1997). We must therefore check to see whether these two propositions are consistent. Consistency requires that the following holds

$$\lim_{Z \rightarrow \infty} A^{\gamma-1} W^\gamma = A^{\gamma-1} \left(W + \frac{1}{\lambda} Y \right)^\gamma$$

To show that this holds we rewrite the proposition of Koo as

$$\lim_{Z \rightarrow \infty} J(W, Y) = A^{\gamma-1} W^\gamma \left(1 + \frac{1}{\lambda Z} \right)^\gamma$$

Which we can see implies that

$$\lim_{Z \rightarrow \infty} J(W, Y) = A^{\gamma-1} W^\gamma$$

Therefore the two propositions are consistent with each other.

It is interesting to note that this consistency holds for any finite value of the complete markets multiplier $\frac{1}{\lambda}$ and thus the proposition of Duffie *et al* says nothing about the investor's valuation of the income stream, just that income becomes insignificant relative to financial wealth.

9.5 Numerical Solution Method

Because the HJB is a non-linear second order differential equation, we cannot use standard finite difference methods. Reduction of the problem to one state variable reduces the dimensionality of the problem and thus the computational complexity of the problem, but does not eliminate the non-linearity of the differential equation. We therefore implement the “Policy Iteration” Algorithm⁸ on the system (9.1). This is essentially a two step variation of the standard finite difference algorithm where non-linearity is circumvented by alternating between calculating the policy functions and the value function implied by those policies.

9.5.1 Description of Algorithm

We begin by defining a grid in Z space with an imposed upper (Z_{max}) and lower (Z_{min}) boundary. Selecting a number of steps (N) for the grid allows us to define the grid co-ordinates as

$$dZ = \frac{Z_{max} - Z_{min}}{N}$$

$$Z_i = Z_{min} + dZ(i - 1)$$

Once values are chosen for Z_{max}, Z_{min} and N the following steps are used to implement the “Policy Iteration” algorithm

⁸See Judd (1998) for an overview of this method.

1. The initial values for the policy functions are specified as the optimal policies for the Merton (1969) no-income case

$$\begin{aligned}\xi[Z_i] &= AZ_i \quad \forall \quad i \\ \psi[Z_i] &= \frac{\mu_m - r}{(1 - \gamma)\sigma_m^2} Z_i \quad \forall \quad i\end{aligned}$$

2. Boundary conditions are imposed at the upper and lower boundaries for Z_i and the HJB is solved for $F[Z_i]$ using the current specification of the policy functions;
3. “Policy Improvements” are estimated by calculating the optimal value for the policy functions implied by the current estimate of $F[Z_i]$; and
4. Steps (2)-(3) are repeated until convergence is reached.

9.5.2 Finite Difference Approximations

To estimate the function $F[Z]$ and the policy functions $(\xi[Z], \psi[Z])$ we take central difference approximations of the first and second derivatives of $F[Z]$

$$\begin{aligned}F'[Z_i] &\approx \frac{F_{i+1} - F_{i-1}}{2dZ} \\ F''[Z_i] &\approx \frac{F_{i+1} - 2F_i + F_{i-1}}{dZ^2}\end{aligned}$$

where the shorthand $F[Z_i] \equiv F_i$ has been used. Substituting the derivative approximations into HJB gives our finite difference approximation for $F[Z_i]$

$$\begin{aligned}\hat{\beta}F_i &\approx (1 + \xi[Z_i])^\gamma + \frac{(F_{i+1} - F_{i-1})(k_2 Z_i - \xi[Z_i] + k_1 \psi[Z_i])}{2dZ} \\ &\quad + \frac{(F_{i+1} - 2F_i + F_{i-1})(\sigma_y Z_i^2 - 2\sigma_y \rho \sigma_m Z_i \psi[Z_i] + \sigma_m^2 \psi[Z_i]^2)}{2dZ^2}\end{aligned}$$

Similarly the finite difference approximations for the policy functions are

$$\begin{aligned}\xi[Z_i] &\approx 2^{-\frac{1}{\gamma-1}} \left(\frac{F_{i+1} - F_{i-1}}{dZ\gamma} \right)^{\frac{1}{\gamma-1}} - 1 \\ \psi[Z_i] &\approx \frac{-4\sigma_y \rho \sigma_m F_i Z_i + F_{i-1} (dZ k_1 + 2\sigma_y \rho \sigma_m Z_i) - F_{i+1} (dZ k_1 + 2\sigma_y \rho \sigma_m Z_i)}{2\sigma_m^2 (F_{i-1} - 2F_i + F_{i+1})}\end{aligned}$$

9.5.3 Boundary Conditions

In this study we consider three different boundary conditions for large values of Z . As general rule, a straight numerical boundary condition should be used as a last resort, but in many applications the solution is unaffected by the boundary condition and thus numerical boundaries perform quite well. In Munk (2000), as well as previous work on numerical solutions to Merton's no income problem⁹, the solutions for the optimal controls were found to perform very poorly near the upper boundary.¹⁰

In this situation we have several theoretical boundary conditions available and thus we will compare the performance of these against the numerical condition.

Numerical

A very simple numerical boundary condition that can be imposed is to assume that $F''[Z] = 0$ at the boundary. This has the effect of making the value at the boundary a projection of the previous values. We can express this condition in finite differences as

$$F_{N+1} = 2F_N - F_{N-1} \quad (9.9)$$

This is a simple and often effective boundary condition for numerical algorithms when the exact behavior at the boundary is not known. However, as we will see in the next section, a numerical boundary condition of this sort performs quite poorly in the context of an HJB style problem

⁹Munk (2003), Munk (1997a), Munk (1997b).

¹⁰Munk (1997b), pg 196 states that "The numerically computed controls can therefore not be trusted for values of z larger than approximately 70% of $\hat{z}$ ".

Convergence to the Merton (1969) Value Function

Duffie, Fleming, Soner, and Zariphopoulou (1997) proved that as the ratio of wealth to income tends towards infinity, the value function becomes asymptotically equivalent to the Merton (1969) value function. We can use this result to directly calculate the value function at the upper boundary Z_{N+1} . We do so by exploiting the following limit

$$\lim_{Z \rightarrow \infty} J(W, Y) = A^{\gamma-1} W^\gamma$$

Using the fact that $J(W, Y) \equiv Y^\gamma F[Z]$, we can rewrite this as

$$\lim_{Z \rightarrow \infty} Y^\gamma F[Z] = A^{\gamma-1} W^\gamma$$

Evaluating the above expression at Z_{N+1} gives our boundary condition

$$F_{N+1} = A^{\gamma-1} Z_{N+1}^\gamma \quad (9.10)$$

Quasi-Merton

While the simple Merton (1969) boundary condition allows for direct computation of the *level* of the value function, this may not be appropriate if the upper boundary Z_{max} is too small. This is because if Z_{max} is too small relative to infinity, where the Merton (1969) solution holds, imposing a value for the level might be quite inaccurate. It thus may be more appropriate to look at the *shape* of the function. To do this we build a boundary condition based on the derivative of the Merton value function. As shown above as Z tends towards infinity the following holds

$$F[Z] = A^{\gamma-1} Z^\gamma$$

Differentiating this with respect to Z yields

$$F'[Z] = \gamma A^{\gamma-1} Z^{\gamma-1}$$

From this it is easy to show that the following equality holds

$$ZF'[Z] = \gamma F[Z]$$

Evaluating this equality at Z_n using finite differences gives the Quasi-Merton boundary condition

$$Z_N \left(\frac{F_{N+1} - F_{N-1}}{2dZ} \right) = \gamma F_N \quad (9.11)$$

Convergence of Implicit Income Multiplier to the Complete Market Income Multiplier

As previously discussed, Koo (1998) proved that the implicit income multiplier converges to the complete markets income multiplier as the ratio of wealth to income goes to infinity. Using the fact that Munk (2000) showed that we can express the implicit income multiplier as

$$B(Z) = \frac{\gamma F[Z]}{F'[Z]} - Z$$

we can rewrite (9.6) as

$$\frac{\gamma F[Z]}{F'[Z]} - Z = \frac{1}{\lambda}$$

Simplifying this expression and converting to finite differences gives our condition for the upper boundary of Z

$$\gamma F_N = \frac{(F_{N+1} - F_{N-1})(Z_N + \frac{1}{\lambda})}{2dZ} \quad (9.12)$$

Convergence of Utility Indifference Income Multiplier to the Complete Market Income Multiplier

We also know that the utility indifference multiplier converges to the complete markets income multiplier. We can therefore combine (9.2) and (9.7)

and evaluate the resulting expression at Z_{N+1} allowing us to directly calculate the value of F_{N+1}

$$F_{N+1} = \left(A^{\frac{\gamma-1}{\gamma}} \left(\frac{1}{\lambda} + Z_{N+1} \right) \right)^{\gamma} \quad (9.13)$$

Hybrid

We can also construct a hybrid boundary condition that combines both of the income multiplier boundary conditions. This is done by evaluating (9.13) at Z_N instead of Z_{N+1} . This then gives an expression for F_N which can be substituted into (9.12) giving the following boundary condition

$$\gamma \left(A^{\frac{\gamma-1}{\gamma}} \left(\frac{1}{\lambda} + Z_N \right) \right)^{\gamma} = \frac{(F_{N+1} - F_{N-1})(Z_N + \frac{1}{\lambda})}{2dZ} \quad (9.14)$$

9.6 Numerical Results

9.6.1 Review of Munk's (2000) Numerical Results

The base case parameters for Munk's analysis are as follows

$$\begin{array}{llll} \gamma = 0.5 & r = 0.1 & \mu_m = 0.15 & \mu_y = 0.05 \\ \beta = 0.2 & \sigma_m = 0.3 & \sigma_y = 0.1 & \rho = 0.0 \end{array}$$

Thus Munk is examining the case where the the risk of the cash flow is completely idiosyncratic. Munk solves (9.1) using the “Policy Iteration” variant of the Markov Approximation Method. This method uses a tri-nomial tree structure for the state variable Z . Thus for each value of Z , there is an associated transition probability for Z evolving up, down or staying the same. Since a finite grid is used for the state variable a condition must be imposed at the artificial upper boundary for Z to approximate the behavior as $Z \rightarrow \infty$. Munk does this by assuming that at Z_{max} the transition probability

for Z increasing is 0. This condition is not based upon theory and is roughly equivalent to our boundary condition where it is assumed that the second derivative of the value function is constant at the upper boundary.

In addition to the qualitative results concerning the value function and optimal controls described in Section 2.4.2, Munk conducts comparative statics analysis to determine the effect of the correlation parameter (ρ), income volatility (σ_m), income drift (μ_m) and the time preference rate (β) on the agent's relative optimal consumption, relative optimal investment in the risky asset¹¹ and utility indifference income multiplier.

For ρ the results are intuitively straightforward and appealing. For low levels of Z the income stream is relatively important and thus the agent has a greater desire to hedge against fluctuations in the income stream. This results in a negative relationship between the proportion of wealth invested in the risky asset and ρ due to the diminished hedging ability of the risky asset as it's correlation with the income stream increases.¹² The flip side of this reduction in risky investment is that the agent consumes more as ρ increases. For high levels of Z , the income stream is relatively unimportant to the investor which results in consumption and investment being roughly constant across ρ . This is because the income stream is relatively insignificant and thus hedging the income stream does not factor into the agent's optimal policies. Using the same intuition about the ability to hedge the income stream, it is no surprise that the income multiplier is a strictly decreasing function of ρ and that the more significant the income stream is to the agent's total wealth (low Z), the steeper the function is.

¹¹The fraction of financial wealth (W) that is optimally consumed and invested in the risky asset.

¹²This is because the investor cannot short sell. If the investor were able to shortsell the risky asset, we would see a convex relationship centered around $\rho = 0$ since the investor could perfectly hedge the income stream at $\rho = \pm 1$, whereas the income stream would be entirely idiosyncratic at $\rho = 0$.

The effect of σ_y on consumption is consistent with the standard precautionary saving motive whereby the agent reduces consumption and increases savings in the face of uncertainty over the future. Again this effect is most prominent for low values of Z where the income stream is relatively important. The effect of σ_y on investment in the risky asset is initially less clear. At high levels of Z when the income stream is relatively less important the fraction of wealth invested in the risky asset decreases with σ_y , whereas for low levels of Z when income is relatively important the proportion of wealth invested into the risky asset increases. Munk attributes the negative relationship at high levels of Z to the agent substituting away from the risky asset since the income stream is becoming a closer substitute to it as σ_y increases. Munk however makes no argument for the increasing positive relationship at low levels of Z . The positive relationship can be explained by acknowledging that there is another force affecting the agent's investment saving decision. When an investor engages in precautionary saving he is reducing current consumption and thus by definition saving more. Some of these savings will be put into the risky asset. The fact that investment in the risky asset increases with σ_y at low levels is not surprising when one considers that consumption decreases substantially as σ_y increases at low levels of Z . Thus at low levels of Z it is simply the case that the desire to shift away from the risky asset is offset by the larger proportion of wealth that is being saved instead of consumed.

Munk next goes on to show that relative consumption increases with the drift of the income process (μ_y), and that the effect of μ_y on consumption is greater when the income stream is relatively significant. As Munk recognizes, the effect of μ_m on investment in the risky asset is harder to understand. Munk notes that when income is relatively significant, the optimal investment in the risky asset is decreasing in μ_m , but when the income stream is relatively less important, investment in the risky asset is marginally increasing in μ_m . A simple explanation can be offered for why optimal investment decreases with μ_m when income is important if one recognizes that the risky asset and

the income stream are in some sense substitutes for each other. Therefore as the growth rate of the income stream increases the investor can obtain the same portfolio return with a lower holding of the risky asset, something that would benefit any risk averse investor. The positive relationship between μ_m and investment in the risky asset when income is insignificant can be explained by recognizing that when income is insignificant the investor will not be as averse to holding the risky asset since a significant portion of his wealth is not already made up by a holding in the risky non-traded income stream. Therefore just as a higher consumption rate can be achieved with a higher μ_m , more wealth can be invested in the risky asset. The distinction between the case of high and low Z values comes from the fact that when income is very substantial the investor is over exposed to risky assets and thus where possible will want to shift away from them, whereas this is less of a problem when income is insignificant.

There are however some weaknesses in the numerical procedure employed by Munk (2000). As noted in Munk (1997b), the computed optimal policies cannot be trusted for values of Z larger than approximately 70% of the upper boundary $\hat{Z}$. The primary cause of this is the error introduced by the boundary condition Munk selects and the fact that he generally uses relatively small values of $\hat{Z}$ which amplify the propagation of the error to the interior solutions¹³.

9.6.2 Our Results

Using Munk's base case parameters we use the policy iteration algorithm described in Section 9.5 to solve the system set out in Equation 9.1. For all the boundary conditions we examined we found the same general qualitative

¹³This stems from the fact that behavior around $\hat{Z}$ is supposed to approximate behavior as $Z \rightarrow \infty$ and thus it is expected that too small a value for $\hat{Z}$ would introduce error into the computations.

results for the value function and optimal policies. Figure 9.1 plots the manager's value function ($J(W, Y)$), optimal consumption ($C(W, Y)$) and optimal investment in the risky asset ($\pi(W, Y)$) using the hybrid boundary condition.

Figure 9.1: *The Value Function, Optimal Consumption and Optimal Investment*

Calculated using the Hybrid boundary condition and $Z_{max} = 1000$

Because the qualitative results are consistent with Munk regardless of the boundary condition used, the focus of our analysis will be on the accuracy of the different boundary conditions. As the key problem with Munk's method is that the optimal controls are only reliable up to 70% of Z_{max} ,¹⁴ a useful starting point in our analysis is to compare the optimal controls for the

¹⁴Munk (1997b).

different boundary conditions we are examining.

Figure 9.2 shows the Value Function, Implicit Income Multiplier, The Utility Indifference Income Multiplier, consumption policy function ($\xi(Z)$) and the risky investment policy function ($\psi(Z)$). Examination of the graphs for $F(Z)$ reveals that all of the boundary conditions do a good job of calculating the value function. In fact we get nearly identical results for the value function across all of the boundary conditions. What stands out about Figure 9.2 is that although we get nearly identical results for the value function, the calculated optimal policies and income multipliers vary drastically across different boundary conditions.

If we look at the optimal policies for the numerical condition we see that we are experiencing similar problems near the upper boundary to those documented in Munk (1997b), in that at approximately 75% of Z_{max} the solutions begin to fall apart. We also witness a similar collapse for the Merton boundary condition which is surprising given that the condition is theoretically based. The Quasi-Merton boundary condition performs much better as we can see that the solution is accurate for roughly 80-90% of Z_{max} after which point the solution deviates, but still remains stable. This is expected as the Quasi-Merton condition is imposing behavior around the shape of the curve, rather than the level and thus is likely to be more accurate at approximating behavior as $Z \rightarrow \infty$. The final three boundary conditions we examine all produce similarly accurate and stable solutions for the optimal policies over the whole range of Z_{max} and thus so far seem to significantly outperform the other boundary conditions as well as the methodology of Munk (2000).

The most striking feature of Figure 9.2 is the income multipliers. Despite the fact that the value function is accurate over the whole range of Z_{max} and the policy functions are accurate over 75% of Z_{max} , we see that for the Numerical, Merton and Quasi-Merton boundary conditions, the solutions for the income multipliers fall apart spectacularly after approximately 30% of Z_{max} . For the Numerical and Merton boundary conditions the same behavior is observed

Figure 9.2: *Value Function, Policy Functions and Income Multipliers for $Z_{max} = 1000$*

(a) Numerical

(b) Merton

(c) Quasi- Merton

(d) Implicit Valuation Multiplier

(e) Utility Indifference Multiplier

(f) Hybrid

In the first column we present the graphs for $F(Z)$ where $J(W, Y) \equiv Y^\gamma F(Z)$. In the last column the horizontal line represents the Complete Markets Income Multiplier as defined by Equation (9.3), the dashed line represents the Implicit Income Multiplier and the solid curved line is the Utility Indifference Income Multiplier

whereby the Implicit Value (IV) Multiplier ascends towards infinity and the Utility Indifference (UI) Multiplier decreases to 0 as $Z \rightarrow \infty$. We already know that this does not make sense because (9.6) and (9.7) show that both multipliers converge to the Complete Markets (CM) multiplier as $Z \rightarrow \infty$. The fact that the income multipliers degenerate significantly sooner than the policy functions suggests that it is the instability of the solution for the income multipliers that is driving the instability of the computed policy functions, rather than the other way round. For the Quasi-Merton boundary condition the behavior is the opposite but is equally bad. Here we observe the IV Multiplier decreasing towards zero while the UI multiplier rises towards infinity as $Z \rightarrow \infty$. It is interesting that we still get behavior for the income multipliers that cannot be reconciled with theory, despite the fact we have a smooth solution to the value function and relatively smooth solution to the optimal policies. For the IV and Hybrid boundary conditions we see that the solution remains stable for approximately 90% of Z_{max} , after which point the IV Multiplier is dragged up towards the CM Multiplier and actually crosses the UI Multiplier. Given that we are forcing the IV Multiplier to equal the CM Multiplier at Z_{max} it is unsurprising that we see this happen. The UI boundary condition exhibits the smoothest solution of all the boundary conditions as we don't see either curve getting artificially "dragged" in any direction.

Given that we have five theoretical boundary conditions, why do some perform so poorly relative to the others? The answer to this question is directly related to the discussion in Section 9.4 on whether or not the propositions of Koo (1998) and Duffie, Fleming, Soner, and Zariphopoulou (1997) concerning the limiting behavior of the value function are consistent with each other. Recall that Duffie, Fleming, Soner, and Zariphopoulou (1997) show that the value function converges to the Merton (1969) value function which is equivalent to acting as if you are not receiving any income. This is contrasted with the result of Koo (1998) who shows that the value function converges to the complete markets case where income is valued using the complete markets

income multiplier $\frac{1}{\lambda}$. As discussed earlier these two propositions are consistent with each other, but the Duffie *et al* proposition states only that income becomes so insignificant that it does not effect the agent's decisions. Since it does not make any statement about how the income stream is valued, it is actually not surprising that when we use the proposition of Duffie *et al* to formulate a boundary condition, it gives nonsensical results for the income multipliers and as a result inaccurate computations for optimal consumption and investment.

Having established that the Numerical, Merton and Quasi-Merton boundary conditions provide unstable solutions, it is now necessary to determine which boundary condition is the most accurate. From Figure 9.2 one might conclude that because the UI boundary condition provides the smoothest solution that it is also the most accurate. This is not necessarily the case and thus we will compute the value function, Policy Functions and Income Multipliers at $Z = 100$ and $Z = 500$ for a range of different values for Z_{max} in order to investigate the error that is propagated to the interior solutions as they get closer to the artificially imposed upper boundary.

In Figures 9.3 and 9.4 we can see that as long as the point at which any of the functions is being calculated is far enough from the upper boundary, the choice of boundary condition does not have any effect as we get identical results across all of the boundary conditions. If we accept that the value calculated when quite far away from the upper boundary represents a fair approximation of the true value, then the deviation away from that value represents a good measure of how accurate each boundary condition is. Applying this criteria to Figures 9.3 and 9.4 allows us to make a judgement on the accuracy of each boundary condition. When we look back at Figure 9.2 it is unsurprising that the Numerical, Merton and Quasi-Merton boundary conditions are all incredibly inaccurate and deviate substantially from the true level as $Z \rightarrow Z_{max}$. On the other hand we can see that the other boundary conditions all give relatively accurate computations as $Z_{max} \rightarrow Z$, with the

Figure 9.3: *Key Variables at $Z = 500$*

Figure 9.4: Key Variables at $Z = 100$

most accurate results being provided by the Hybrid Boundary condition. The Hybrid Boundary condition is closely followed by the IV boundary condition which is only slightly less accurate. Interestingly the UI boundary condition, while still being substantially more accurate than the numerical and Merton based boundary conditions, is much less accurate than the Hybrid and IV conditions. This is despite the fact that it provides a much smoother solution as evidenced by Figure 9.2. It is also interesting to note that the deviation across all boundary conditions as $Z_{max} \rightarrow Z$ is much larger for $Z = 100$ than for $Z = 500$. This is actually to be expected as we are imposing the boundary conditions for much lower values of Z_{max} and thus we cannot expect behavior that holds as $Z \rightarrow \infty$ to hold when Z_{max} takes a relatively small value.

To illustrate the dangers of using a small value for Z_{max} with a numerical boundary condition, it is useful to recreate Figure 4 from Munk (2003). In this graph Munk plots both income multipliers and the complete markets multiplier to illustrate how slowly they converge. The multipliers are calculated using $Z_{max} = 200$ yet he only plots the graph out to $Z = 100$. In Figure 9.6(a), we recreate Munk's plot with inclusion of the same result using the Hybrid Boundary Condition, in addition we also show the full plot out to $Z = 200$ for comparison.

Just like in Munk's Figure 4, panel (a) of Figure 9.5 shows that when we plot the graph out to $Z = 100$ we see that the lower dashed line representing the implicit multiplier calculated using the numerical boundary condition starts to veer upwards as $Z \rightarrow 100$, while the corresponding line for the Hybrid boundary condition continues to change at the same rate. Given that these plots are calculated using $Z_{max} = 200$ this is suspicious. While panel (b) of Figure 9.5 is not presented in Munk (2003), the fact that our plots in panel (a) for the numerical boundary condition are identical suggest that Munk must have experienced something similar. The results presented in panel (b) illustrate the failing of the numerical boundary condition by providing a stark contrast with the Hybrid boundary condition. We can see that soon after Z

Figure 9.5: *Comparison of Hybrid and Numerical Boundary Conditions Income Multipliers for $Z_{max} = 200$ and $Z_N = 2000$*

Where the dashed lines correspond to the numerical boundary condition and the solid lines are the Hybrid condition. The horizontal solid line at 15 is again the complete markets multiplier.

becomes greater than 100, the two income multipliers veer off sharply, the implicit towards infinity, the utility indifference towards zero and actually cross. This makes no sense theoretically and thus is purely a manifestation of the method used to solve the problem. While the multipliers from the Hybrid condition do veer off and cross each other, the change is nowhere near as drastic and in fact does not occur until after $Z = 150$. This suggests that while any boundary condition is going to have problems if a small grid is used (i.e small Z_{max}), the use of an *appropriate* theoretical boundary condition can help mitigate this significantly.¹⁵

¹⁵Where the word appropriate is used because even the Merton and Quasi-Merton boundary conditions performed poorly.

9.7 Manager's subjective discount rate

To gain more insight into the significance of the base case parameters it is useful to calculate the complete markets income multiplier

$$\frac{1}{\lambda} = \frac{1}{r - \mu_y + \frac{\sigma_y}{\sigma_m}(\mu_m - r)}$$

Calculating this for the base case parameters gives an income multiplier of 15. Since this is essentially a complete markets, continuous time version of the Gordon growth model, λ can be interpreted as the discount rate applied to the project's cashflows. Thus we can say that the complete markets discount rate for this project is $\lambda = \frac{1}{15} = 6.7\%$. This transformation is useful because the discount rates are a more natural measure to use when discussing investment decisions than multipliers. Applying the same transformation to the UI multiplier¹⁶ allows us to analyse the discount rate a constrained agent would use for a non-traded cashflow. Thus if the UI income multiplier is $B^*(Z)$, then this implies the discount rate is $\frac{1}{B^*(Z)}$.

If we recall that we can define $Z \equiv \frac{W}{\alpha Y}$, we can analyse the decision of a manager who is constrained to own the fraction α of the firm's cashflow. To analyse this decision and its implications for shareholders, we can fix the ratio $\frac{W}{Y}$ at various levels and examine the impact that changing α has on the discount rate used by the manager to evaluate the project.

Figure 9.6 plots the discount rate used by the manager depending on how much of the firm he owns for various levels of the ratio $\frac{W}{Y}$. Three things are apparent from this graph. First, unless they are independently wealthy ("rich") and they do not own "too much" of the firm, managers adopt a discount rate that exceeds the complete markets rate. The reason is simple: the inability to trade their stake in the firm makes managerial wealth dependent on the fortunes of the employing firm, and so managers adopt

¹⁶Given we are interested in investment decisions, the UI value is a more natural measure since it describes what the manager would give up to obtain the cash flow.

Figure 9.6: *UI Discount Rate vs α*

Calculated for $\frac{W}{Y} = 0.02, 0.1, 1, 5, 20, 100, 200, 500, 1000$, where the lower dashed lines correspond to the high values of $\frac{W}{Y}$

a safety-first approach that screens out projects of marginal profitability (but which would nevertheless add to shareholder wealth). Thus there is an ‘under-investment’ problem: managers generally invest in fewer projects than shareholders would like. Second, “poor” managers (those who are not independently wealthy) choose a higher discount rate than “rich” managers because firm ownership has a greater impact on the diversification of their portfolio. That is, they care more about the firm’s specific risk. Third, for similar reasons, the discount rate chosen by “poor” managers is much more sensitive to their ownership share.

The most important thing to take from Figure 9.6 is that the impact of under-diversification on investment decision making can be severe: even a tiny amount of firm ownership can approximately double the hurdle rate adopted by a manager with little wealth.

9.8 Adding Effort to the CRRA Model

The analysis of this chapter has thus far demonstrated that wealth effects the manager's investment decision through its impact on the valuation of the cash flow. For a higher W , the manager's subjective discount rate decreases which makes investment *more* likely. This contrasts to the CARA model of Chapter 4 where the manager's valuation of the cash flow is independent of his financial wealth. The introduction of effort in Chapter 5 changed this and wealth affected the manager's investment decision through the impact of wealth on optimal effort. Increasing W reduces the marginal utility of wealth and thus the manager has less incentive to exert cost reducing effort. Therefore investment is *less* likely as W increases in the CARA/SBM model with effort.

The purpose of this section is to introduce effort into the CRRA/GBM model and determine if a similar result to Chapter 5 (wealth reducing optimal effort) still holds.

9.8.1 Solution method

The setup for this problem is essentially the same as for Chapter 5. The manager's value function and the investment cost function are defined as

$$J^e(W, Y) = J^2(W - \alpha I[e], Y) - \theta e$$

$$I[e] = \exp(-\lambda e)A + (1 - \exp(-\lambda e))B \quad e \geq 0$$

As in the previous section of this chapter, we make use of the transformation $J^2(W, Y) = (\alpha Y)^\gamma F[\frac{W}{\alpha Y}]$, where $F[\frac{W}{\alpha Y}]$ is the interpolated solution for the manager's value function solved in the previous part of this chapter. The manager's value function can thus be expressed as

$$J^e(W, Y) = (\alpha Y)^\gamma F\left[\frac{W - \alpha I[e]}{\alpha Y}\right] - \theta e \quad (9.15)$$

Because we are interested in solving for the level of effort (e^*) that maximises Equation (9.15), and e appears inside the interpolated function F , the problem must be maximised numerically. This is done by specifying a grid for W and Y and then using the NMaximise command in Mathematica to find the level of e that maximises Equation (9.15) subject to the following constraints:

1. $e \geq 0$
2. $W \geq \alpha A$

The first constraint is the same “non-shirking” constraint as imposed in Chapter 5, while the second constraint is introduced because the GBM/CRRA model of this chapter constrains wealth to be positive. In the absence of this second constraint the manager could have negative wealth at the time of investment which results in the interpolated function F being undefined. This could also be dealt with by constraining the manager to exert enough effort to make wealth non-negative if he invests, but this introduces another dimension to the manager’s effort decision which limits comparability with the results of Chapter 5.

9.8.2 Results

To solve for e^* we use the same base case parameters as Section 9.6 to solve for the interpolated function F and the base case effort related parameters from Chapter 5, which for the sake of completeness are

Table 9.1: *Effort related parameters*

$$A = 100 \quad B = 80 \quad \theta = 0.1 \quad \lambda = 1$$

The grid over which we will solve for e^* is specified in Table 9.2 below.

Table 9.2: *grid parameters*

$$\begin{array}{lll}
W_{min} = \alpha A & W_{max} = W_{min} + 10 & W_N = 20 \\
Y_{min} = 5 & Y_{max} = 15 & Y_N = 20
\end{array}$$

Using this grid and the effort related parameters specified in Table 9.1, the solution for e^* is shown in Figure 9.7. This graph demonstrates that optimal effort is a decreasing function of both W and Y . This is the same result as we found in Chapter 5, thus the inclusion of wealth effects does not qualitatively change the manager's optimal *effort* decision.

Figure 9.7: e^* as a function of W and Y

However, the manager's *investment* decision is different. Figure 9.8 shows a contour plot of the manager's threshold. The light area represents those areas where the manager would invest while the dark shaded area represents those areas where he would not. This graph shows that as W increases, the investment threshold decreases (i.e. investment is *more* likely).

Figure 9.8: *Investment threshold in the CRRA/GBM model with effort*

This contrasts to the results of Chapter 5 where the investment threshold increases with W and thus investment is *less* likely. This occurs because an increase in W increases the manager's valuation of the cash flow in the wealth dependent CRRA model. Therefore, despite effort decreasing as W increases, the impact of this on the investment decision is more than offset by the increased value placed on the cashflow.

9.9 Summary

This chapter set out to consider the implications of CRRA on the manager's valuation problem by conducting a parallel analysis to that of Chapter 4. By using CRRA utility and a GBM cash flow, the “first step” of the analysis this thesis seeks to carry out requires a complex numerical solution. This

chapter therefore examined these numerical issues in isolation given they made “continuing down the path” intractable. The main methodological lesson learned from this chapter is that the choice of boundary condition can be critical when conducting numerical analysis. In particular, theoretically grounded boundary conditions perform much better than numerical conditions, although some theoretical boundary conditions perform much better than others and some are even as bad as the numerical condition. The key for this model appears to be selecting a boundary condition that says something about the manager’s valuation of the cashflow, rather than general conditions surrounding the value function.

Once the “journey” of solving the model was complete, the CRRA model was used to examine how wealth dependence affects the manager’s valuation problem. It is commonly argued that remunerating managers with stock and options grants will make them think like shareholders (since they are entitled to a share of profits), and hence run the company in a manner desired by shareholders. However, this overlooks the fact that such grants paradoxically create a conflict-of-interest problem: managers use a higher discount rate than is optimal for shareholders and hence pass up investment projects that would enhance shareholder wealth.

In designing managerial remuneration policy, shareholders must therefore trade off the benefits of incentive alignment with the conflict-of-interest cost caused by over exposing the manager to the firm. In this chapter we have shown that if managers have CRRA utility functions, then it is not just the level of their firm ownership that impacts how much idiosyncratic risk they take into account when making investment decisions, but that the hurdle rate they select depends crucially on their levels of wealth. In fact, the results of this chapter suggest that it is a “rich” manager that is most likely to think like a shareholder, but interestingly he will generally do this even if he owns only a small portion of the firm. On the other hand, the more firm ownership a “poor” manager is given the less he thinks like a shareholder as his personal

wealth becomes overly dependent on the fortunes of the firm and he will thus make overly cautious investment decisions as evidenced by the high hurdle rate he selects.

Introducing effort doesn't appear to change this result either. In the wealth dependent CRRA setting, the impact of wealth on the manager's valuation of the cash flow more than offsets his lowered incentive to exert effort as wealth increases.

Part IV

Conclusion and Appendices

Chapter 10

Summary and Conclusion

This thesis began by asking two questions:

1. if managers own too much stock, how does this affect the investment decision they make for the firm? and
2. given the answer to (1), how does this affect the shareholder's decision to hire a manager?

The fact that the answer to the first question is required to answer the second had implications for the structure of this thesis. We initially considered two utility functions for the manager, Constant Relative Risk Aversion (CRRA) and Constant Absolute Risk Aversion (CARA). However, the problem is significantly complicated by the fact that financial wealth affects the manager's valuation of the cashflow when he has a CRRA utility function and is constrained to own part of the firm.

Therefore, with the goal of having a “clean path” to answering the second question, Part II focused on the CARA model and completed the various steps required to analyse both questions.

Chapters 4 and 5 analysed the first question in a now-or-never setting. The key result from Chapter 4 is that constraining the manager to hold “too much” of the firm causes him to value the firm’s project less than a well diversified shareholder would. Chapter 5 attempted to analyse the “other side of the coin” by introducing effort into the model. The key result here was the manager’s optimal level of effort is highly non-linear in the proportion of the firm he owns and depends heavily on how risk averse the manager is. In other words, more is not always better when it comes to incentivising managers.

Chapter 6 analysed the second question in a static setting by using the model of Chapter 5. The shareholder is effectively trading off three things when determining how much of the firm to give to the manager:

1. making the manager less diversified;
2. incentivising the manager to exert more effort; and
3. diluting his own share of the firm.

The factors that have the greatest impact on the optimal level of managerial ownership are, perhaps unsurprisingly, those specific to the manager. More interesting is that market and project specific parameters have smaller effects on the optimal level of managerial ownership. This is interesting as it suggests that the type of firm a manager works for is not that important when determining his compensation.

Chapter 7 addressed the first question when the manager is able to delay investment. The interesting finding from this chapter is that because effort can be used to hedge the market,¹ managers will delay investment beyond the standard predictions. This is in order to preserve the ability hedge. Put another way, they do not want to invest now, exert no effort and subsequently receive “bad news” that will make them wish they had worked harder.

¹I.e. you work harder when your portfolio investments perform poorly.

To finish Part II, Chapter 8 addressed the second question when investment can be delayed both when the shareholder must make a now-or-never decision to hire and when the decision to hire can be delayed. In both the now-or-never and dynamic hiring models, the more likely it is the manager will exert effort upon investment, the more likely it is that a manager will be hired now. When the ability to delay hiring is introduced, the shareholder will delay his decision in some regions where he would have hired a manager in a now-or-never setting. This region corresponds to the area where the manager has “choice” over his level of effort.² Intuitively, the shareholder delays hiring in this region to avoid hiring a manager who subsequently exerts no effort upon investment.

Chapter 8 also attempted to address the question of what type of manager a shareholder would prefer, keeping in mind that in the model there is only one manager and no search. The general results of this exercise are largely as expected - the shareholder is better off with a manager who is less risk averse, more skilled and requires a smaller share of the firm as payment. However, the non-linearities in the model mean that there are situations where the shareholder is indifferent and in fact where the opposite occurs.

In Part III of the thesis we returned to the CRRA model. Because this utility function causes the manager’s valuation to depend on his financial wealth, the numerical solution method becomes complicated. However, the complications to the numerical solution method are particularly interesting and thus Chapter 9 analysed the unique numerical issues introduced by wealth dependent valuations. The key methodological insight from this chapter is that the choice of boundary condition can be of crucial importance. While not to the same extent as in Part II, this model also gives some insight into how the introduction of wealth effects alters the answers to the first question addressed above. In particular, a “rich” manager is likely to make investment

²The “choice” in this context stems from the fact that the manager is not so rich that he never exerts effort and not so poor that he never exerts effort.

decisions more in line with shareholder interests than a poor manager. This is because being undiversified is relatively less important for a rich manager.

This thesis has also identified many areas for future research. In Part II of the thesis this would involve introducing the following features into the models (either individually or collectively):

- make the manager's level of effort unobservable to the shareholder
- make the manager's characteristics initially unobservable to the shareholder³
- allow the manager to be dismissed
- introduce multiple managers that the shareholder must choose from and more importantly *search* for
- endogenise α in the model of Chapter 9
- introduce an explicit abandonment option
- examine the interplay between the executive compensation and the level of investment flexibility⁴

The key extension that could be made to Part III of this thesis would be to introduce effort into the CRRA/GBM model. This would allow examination of how wealth dependant valuations affect the effort decision. That is, the interaction of indirect and direct wealth effects could be analysed.

³This would be done in the style of a learning over time model.

⁴Note that the period for which investment can be delayed has effectively been fixed throughout this thesis.

Appendix A

ROA Valuation Using the CAPM

A.1 Setup

In this appendix we discuss the solution method used value the project rights when an investor is compensated for risk according to the capital asset pricing model (CAPM). As discussed in Section 2.2, the key assumption of the CAPM is that investors are not compensated for idiosyncratic risk.

The setup for this model is the same as for Part II of this thesis (the “clean path”) in that the shareholder can pay a lump sum I to receive a cashflow (Y) which follows simple brownian motion and the investment decision can be delayed. The difference between the shareholder and the manager, besides using the CAPM to value the project rights, is that the shareholder incurs a monitoring cost of κ while waiting to invest.

Given that we are assuming the CAPM holds, solving the shareholder’s valuation of the project rights is relatively straightforward. Equation 3.27 of Trigeorgis (1996) is the differential equation that any contingent claim with a single state variable (V) must satisfy, subject to a terminal condition and

to a lower and an upper boundary condition.

$$\frac{1}{2}\sigma^2V^2F_{VV} + (\alpha - \lambda\sigma)VF_V - F_\tau + d = 0 \quad (\text{A.1})$$

As noted by Trigeorgis (1996,p97), if the CAPM holds then $\lambda = (\mu_m - r)\frac{\rho}{\sigma_m}$. Therefore, Equation 3.27 of Trigeorgis (1996) can be re-written using the notation of this thesis for the CAPM valuation of the project rights ($F(Y, t)$).

$$\frac{1}{2}\sigma_y^2Y^2F_{YY} + (\mu_y - (\mu_m - r)\frac{\rho}{\sigma_m}\sigma_y)YF_Y - F_t - \kappa = 0 \quad (\text{A.2})$$

The terminal and boundary conditions are discussed in the next section in the context of the numerical solution method.

A.2 Numerical Solution

To solve this model we use what is effectively the explicit finite difference method. The steps of the algorithm are as follows:

1. The value function at the last possible date (t_{N+1}) is calculated using the terminal condition
2. boundary conditions are imposed along the upper (Y_{N+1}) and lower boundaries (Y_1)
3. the algorithm progressively steps backwards in time, using the solutions for the value function at t_i to calculate the value function at t_{i-1}

To implement this algorithm we must define the numerical grid, choose finite difference approximations for the derivatives of the value function and specify the terminal, upper boundary and lower boundary conditions.

Given we only have the single state variable (Y) and a time variable (t), the grid is defined as follows

$$dY = \frac{Y_{max} - Y_{min}}{Y_N}$$

$$\begin{aligned}
Y_i &= Y_{min} + dY(i - 1) \\
dt &= \frac{t_{max} - t_{min}}{t_N} \\
t_i &= t_{min} + dt(i - 1)
\end{aligned}$$

where (Y_{max}, t_{max}) are the maximum values for the grid, (Y_{min}, t_{min}) are the minimum values and (Y_N, t_N) are the number of steps for each variable.

To obtain a finite difference approximation of equation (A.2), the following approximations are used for the derivatives of $F(Y_i, t_i)$

$$\begin{aligned}
F_t &\approx \frac{F[W_i, Y_i, t_i] - F[W_i, Y_i, t_{i-1}]}{dt} \\
F_Y &\approx \frac{F[W_i, Y_{i+1}, t_i] - F[W_i, Y_{i-1}, t_i]}{2dY} \\
F_{YY} &\approx \frac{F[W_i, Y_{i+1}, t_i] - 2F[W_i, Y_i, t_i] + F[W_i, Y_{i-1}, t_i]}{dY^2}
\end{aligned}$$

That is, for derivatives with respect to Y central difference approximations are used, while for the time derivative a backward difference is used. The result of this is that it is possible to solve for the value function at date t_{i-1} ($F[Y_i, t_{i-1}]$) as a function of the date t_i value function. This is why the algorithm can solve the value function by stepping backwards in time starting from the terminal values.

For the terminal condition, we assume the investor must make a now-or-never decision of whether or not to invest, and that the cashflow is valued using the CAPM. Using the CAPM valuation of an SBM cashflow set out in Section 2.2, the terminal condition is therefore

$$F(Y_i, t_{n+1}) = \max \left[\frac{Y_i}{r} + \frac{\mu_y - \rho\sigma_y\Phi}{r^2} - I, 0 \right] \quad (\text{A.3})$$

For the boundary conditions, we assume that when Y is very large, investment occurs immediately and that when Y is very small $F_{YY} = 0$.^{1,2} Therefore the upper and lower boundaries can be expressed as

$$F[Y_{n+1}, t_{i-1}] = \frac{Y_{n+1}}{r} + \frac{\mu_y - \rho\sigma_y\Phi}{r^2} - I \quad (\text{A.4})$$

$$F[Y_1, t_{i-1}] = 2F[Y_2, t_{i-1}] - F[Y_3, t_{i-1}] \quad (\text{A.5})$$

Note that the lower boundary is not bounded at 0. This is because no explicit abandonment option has been included in the model, though as discussed in Chapter 8 the shareholder has an implicit abandonment option through giving the project to the manager.

¹This is the numerical boundary condition discussed in Section 9.9 of this thesis. In effect it is assumed that the value at the boundary is a linear projection of the interior values.

²Note that to obtain a boundary condition for Y_1 , the finite difference approximation for F_{YY} is evaluated at Y_3 and then solved for $F(Y_1, t_{i-1})$.

Appendix B

Additional Comparative statics: Chapter 5

B.0.1 Investment Cost Ceiling (A)

For low values of A the manager doesn't exert enough effort to warrant hiring him so the shareholder sets $\alpha^* = 0$ which means the manager's utility is $J^M(W)$. As A increases the manager starts to exert effort and thus the hiring constraint stops binding and we get a positive value for α^* . The interesting feature of Figure B.1 is that α^* is insensitive to changes in A . The fact that the ceiling for the investment cost doesn't affect the choice of α^* other than satisfying the participation and hiring constraints is interesting. Intuitively this makes sense as it implies that it is the lower limit of the investment cost that matters, not the upper limit. Given that the whole point of hiring a manager is to reduce the investment cost, it makes sense that it is the *level* the investment cost can be reduced to that determines α^* .

Figure B.1: *Comparative Statics For A*

(a) Solid Line: $\text{Payoff}^{manager}[\alpha^*]$, Dashed Line: Payoff^{CAPM}

(b) Solid Line: $J^2(W - \alpha^* I[\hat{e}], Y)$, Dashed Line: $J^M(W)$

B.0.2 Aversion to effort (θ)

The effect of θ on α^* is essentially the opposite of λ which is unsurprising given that λ is essentially the benefit of effort while θ is the cost. As the cost of effort increases, the manager's optimal level of effort decreases and thus α^* increases to offset this. Because optimal effort is falling, the payoff to the shareholder is falling and thus if θ is large enough the hiring constraint starts to bind and α^* drops to zero.

Figure B.2: *Comparative Statics For θ*

(a) Solid Line: $\text{Payoff}^{manager}[\alpha^*]$, Dashed Line: Payoff^{CAPM}

(b) Solid Line: $J^2(W - \alpha^* I[\hat{e}], Y)$, Dashed Line: $J^M(W)$

B.0.3 Correlation Co-efficient (ρ)

The effect of ρ is quite simple. An increase in ρ increases the systematic risk of the cashflow and thus lowers the payoff to the shareholder. To offset this the shareholder increases α^* to induce the manager to exert more effort. The result of this is that $\text{Payoff}^{manager}[\alpha^*]$ decreases at a slightly lower rate than Payoff^{CAPM} . However this difference isn't that pronounced because the change in α^* is very small due to the fact that ρ doesn't have a very significant impact on the CAPM valuation of an SBM cashflow.

Figure B.3: *Comparative Statics For ρ*

(a) Solid Line: $\text{Payoff}^{manager}[\alpha^*]$, Dashed Line: Payoff^{CAPM}

(b) Solid Line: $J^2(W - \alpha^* I[\bar{e}], Y)$, Dashed Line: $J^M(W)$

B.0.4 Cash Flow Growth Rate (μ)

We already know from Chapter 5 that all other things being equal, the manager's optimal effort is decreasing in μ which would lead one to think that α^* would increase in μ . As Figure B.4 shows this isn't the case and α^* decreases in μ . To understand why this occurs it is useful to examine how the manager's level of effort and thus the investment cost change as μ changes when α^* is endogenised and thus depends on μ .

Figure B.5 shows that when α^* is endogenous, an increase in μ decreases the

Figure B.4: *Comparative Statics For μ*

(a) Solid Line: $\text{Payoff}^{\text{manager}}[\alpha^*]$, Dashed Line: $\text{Payoff}^{\text{CAPM}}$

(b) Solid Line: $J^2(W - \alpha^* I[\hat{e}], Y)$, Dashed Line: $J^M(W)$

manager's optimal effort which in turn increases the investment cost. Given that the investment cost is increasing, for this policy to be the optimal for the shareholder, it must be the case that the gain from diluting the manager's claim on the firm (decreasing α^*) is greater than the loss from the increased investment cost. Given that the payoff to the shareholder from employing the manager remains approximately parallel to the payoff from not employing a manager, this trade off is satisfied.

Another way to think about why this is happening is in terms of the marginal

Figure B.5: *Optimal Effort ($\hat{e}$) and the Investment Cost ($I(\hat{e})$) when α^* is endogenous*

costs and benefits from a shareholder's perspective. As μ increases, all other things being equal the manager will exert less effort. Therefore to cause the manager to exert the same level of effort, the manager require a higher level of α . The result of this is that the marginal cost of reducing the investment cost rises as μ increases, and thus the shareholder's optimization results in a lower α^* .

Appendix C

Additional Comparative Statics: Chapter 9

Aside from the “ridge” in the investment region and the shift of the “left boundary”, the comparative statics for the shareholder’s dynamic hiring decision (i.e. when hiring can be delayed), differ little from when hiring cannot be delayed (Section 8.4). For the interested reader, the graphs that were shown in Section 8.4, but not in Section 8.5 because they did not shed light on the “ridge” or the left are produced here.

Figure C.1: κ and the dynamic hiring decision

Figure C.2: σ_y and the dynamic hiring decision

Figure C.3: γ and the dynamic hiring decision

(a) $\gamma = 0.5$

(b) $\gamma = 1$

(c) $\gamma = 2$

(d) $\gamma = 3$

Figure C.4: α and the dynamic hiring decision

(a) $\alpha = 0.05$

(b) $\alpha = 0.1$

(c) $\alpha = 0.2$

(d) $\alpha = 0.25$

Bibliography

- Abowd, John M., and David S. Kaplan, 1999, Executive Compensation: Six Questions That Need Answering, *Journal of Economic Perspectives* 13, 145–168. 6
- Akdoğan, Evrim, and Peter MacKay, 2008, Investment and Competition, *Journal of Financial and Quantitative Analysis* 43, 299–330. 25
- Amihud, Yakov, and Baruch Lev, 1981, Risk reduction as a managerial motive for conglomerate mergers, *The Bell Journal of Economics* 12, 605–617. 7
- Berger, Philip G., and Eli Ofek, 1995, Diversification's Effect on Firm Value, *Journal of Financial Economics* 37, 39–65. 8
- Bettis, J. Carr, Jeffery L. Coles, and Michael L. Lemmon, 2000, Corporate policies restricting trading by insiders, *Journal of Financial Economics* 57, 191–220. 3
- Bliss, Robert R., and Nikolaos Panigirtzoglou, 2004, Option-implied risk aversion estimates, *Journal Of Finance* LIX, 407–446. 50

- Bodie, Zvi, Robert C. Merton, and William F. Samuelson, 1992, Labor supply flexibility and portfolio choice in a life cycle model, *Journal of Economic Dynamics & Control* 16, 427–449. 119
- Boyer, Marcel, Éric Gravel, and Pierre Lasserre, 2004, Real options and strategic competition, *CIRANO Working Paper*. 25
- Boyle, Glenn, and Graeme Guthrie, 2003, Investment, uncertainty and liquidity, *Journal of Finance* 58, 2143–2166. 26
- , 2006, Hedging the value of waiting, *Journal of Banking and Finance* 30, 1245–1267. 26
- , and Richard Meade, 2006, Real options and transmission investment: the new zealand grid investment test, *available at www.iscr.org.nz/documents/git-iscr.pdf*. 27
- Brealey, Richard A., Stewart C. Myers, and Franklin Allen, 2006, *Principles of Corporate Finance* (McGraw-Hill Irwin) 8 edn. 18
- Bulan, Laarni, Christopher Mayer, and C. Tsurriel Sommerville, 2009, Ir-reversible investment, real options, and competition: Evidence from real estate development, *Journal of Urban Economics* 65, 237–251. 24
- Campbell, John Y., and Luis M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors* (Oxford University Press). 30, 49, 50
- Childs, Paul D., Steven H. Ott, and Timothy J. Riddiough, 2002, Optimal valuation of claims on noisy real assets: Theory and an application, *Real Estate Economics* 30, 415–433. 25
- Chirinko, Robert S., and Huntley Schaller, 2009, The irreversibility premium, *Journal of Monetary Economics* 56, 390–408. 24

- Comment, Robert, and Gregg A. Jarrell, 1995, Corporate Focus and Stock Returns, *Journal of Financial Economics* 37, 67–87. 8
- Copeland, Thomas, J. Fred Weston, and Kuldeep Shastri, 2005, *Financial Theory and Corporate Policy* (Pearson Addison Wesley) 4 edn. 18, 21
- Denis, David J., Diane K. Denis, and Atulya Sarin, 1997, Agency problems, equity ownership and corporate diversification, *The Journal of Finance* LII, 135–160.
- Dixit, Avinash K., and Robert S. Pindyck, 1994, *Investment Under Uncertainty* (Princeton University Press). 8, 12, 23, 28, 115, 117, 126, 131, 138, 161
- Duffie, Darrell, Wendell Fleming, H. Mete Soner, and Thaleia Zariphopoulou, 1997, Hedging in incomplete markets with HARA utility, *Journal of Economic Dynamics & Control* 21, 753–782. 32, 33, 34, 179, 181, 182, 186, 187, 191, 200
- El Karoui, Nicole, and Monique Jeanblanc-Picqué, 1998, Optimization of consumption with labor income, *Finance and Stochastics* 2, 409–440. 32
- Fama, Eugene, and Kenneth French, 1992, On the Cross-Section of Expected Returns, *Journal of Finance* 47, 427–466. 43
- , 1995, Size and Book-to-Market Factors in Earnings and Returns, *Journal of Finance* 50, 131–155. 43
- , 1996, Multifactor Explanations of Asset Pricing Anomalies, *Journal of Finance* 51, 55–84. 43
- Garen, John E., 1994, Executive Compensation and Principal-Agent Theory, *The Journal of Political Economy* 102, 1175–1199.
- Grenadier, Steven R., and Neng Wang, 2005, Investment Timing, Agency and Information, *NBER Working Paper 11148*. 43

- Guthrie, Graeme, 2007, Missed opportunities: optimal investment timing when information is costly, *Journal of Financial and Quantitative Analysis*. 28
- , 2010, House prices, development costs, and the value of waiting, *Journal of Urban Economics* 68, 56–71. 25
- Hall, Brian J., and Jeffrey B. Liebman, 1998, Are CEOs Really Paid Like Bureaucrats?, *The Quarterly Journal of Economics* CXIII, 653–691. 6
- Hall, Brian J., and Kevin J. Murphy, 2002, Stock options for undiversified executives, *Journal of Accounting and Economics* 33, 3–42. 31
- He, Hua, and Henri F. Pages, 1992, Labor income, borrowing constraints, and equilibrium asset prices, *Economic Theory* 3, 663–696. 30, 31, 32
- Henderson, Vicky, 2007, Valuing the Option to Invest in an Incomplete Market, *Mathematics and Financial Economics* 1, 103–128. 11, 34, 35, 38, 39, 40, 50
- , 2009, Utility indifference pricing - an overview, in R. Carmona, ed.: *Volume On Indifference Pricing* (Princeton University Press). 33, 116
- Holmstrom, Bengt, 1979, Moral hazard and observability, *The Bell Journal of Economics* 10, 74–91. 11, 73, 90
- Hugonnier, Julin, and Erwan Morellec, 2007a, Corporate control and real investment in incomplete markets, *Journal of Economic Dynamics and Control* 31, 1781–1800. 43, 137
- Hugonnier, Julien, and Erwan Morellec, 2007b, Real Options and Risk Aversion, *Simon Business School Working Paper No. FR 03-19*. 35, 40
- Ingersoll, Jonathan E., 2006, The Subjective and Objective Evaluation of Incentive Stock Options, *The Journal of Business* 79, 453–487. 31

- Jensen, Michael C., 1986, Agency Costs of Free Cash Flow: Corporate Finance and Takeovers, *American Economic Review* 76, 323–329.
- , 2001, How Stock Options Reward Managers for Destroying Value and What To Do About It, *Harvard Business School Negotiations and Markets (NOM) Working Paper*. 6
- , and William H. Meckling, 1976, Theory of the Firm: Managerial Behaviour, Agency Costs and Ownership Structure, *Journal of Financial Economics* 3, 305–360. 5, 6
- Jensen, Michael C., and Kevin J. Murphy, 1990a, CEO Incentives-It's Not How Much You Pay, But How, *Harvard Business Review* pp. 139–153. 6
- Jensen, Michael C., and Kevin J. Murphy, 1990b, Performance Pay and Top-Management Incentives, *The Journal of Political Economy* 98, 225–264. 6
- Jensen, Michael C., Kevin J. Murphy, and Eric G. Wruck, 2004, Remuneration: Where we've been, how we got to here, what are the problems, and how to fix them, *ECGI Finance Working Paper No. 44/2004*. 3, 6, 7
- Jin, Li, 2002, CEO compensation, diversification, and incentives, *Journal of Financial Economics* 66, 29–63. 42, 105
- Judd, Kenneth L., 1998, *Numerical Methods in Economics* (MIT Press, Cambridge, MA). 38, 110, 188
- Kahl, Matthias, Jun Liu, and Francis A. Longstaff, 2003, Paper millionaires: how valuable is stock to a stockholder who is restricted from selling it?, *Journal of Financial Economics* 67, 385–410. 31
- Kole, S., 1997, The complexity of compensation contracts, *Journal Of Financial Economics* 43, 79–104. 3

- Koo, Hyeng Keun, 1998, Consumption and Portfolio Selection with Labour Income: A Continuous Time Approach, *Mathematical Finance* 8, 49–65. 32, 33, 179, 184, 185, 187, 192, 200
- Lambert, Richard A., William N. Lanen, and David F. Larcker, 1989, Executive stock option plans and corporate dividend policy, *The Journal of Financial and Quantitative Analysis* 24, 409–425. 6
- Lamont, Owen A., and Christopher Polk, 2002, Does diversification destroy value? Evidence from industry shocks, *Journal of Financial Economics* 63, 51–77. 8
- Lins, Karl, and Henri Servaes, 1999, International Evidence on the Value of Corporate Diversification, *The Journal of Finance* LIV, 2215–2239. 8
- Markowitz, Harry, 1952, Portfolio selection, *The Journal Of Finance* 7, 77–91.
- May, Don O., 1995, Do Managerial Motives Influence Firm Risk Reduction Strategies?, *The Journal of Finance* L, 1291–1308. 8
- McDonald, Robert, and Daniel Siegel, 1986, The Value of Waiting to Invest, *Quarterly Journal of Economics* 101, 707–726. 23, 36, 37, 40
- Meier, Iwan, and Vefa Tarhain, 2007, Corporate investment decision practices and the hurdle rate premium puzzle, *HEC Montreal Working Paper*. 24
- Merton, Robert C., 1969, Lifetime portfolio selection under uncertainty: The continuous-time case, *The Review of Economics and Statistics* 51, 247–257. 31, 46, 54, 55, 56, 61, 67, 79, 83, 112, 113, 116, 187, 189, 191, 200
- , 1971, Optimum Consumption and Portfolio Rules in a Continuous-Time Model, *Journal of Economic Theory* 3, 373–413. 31
- Meulbroek, Lisa K., 2001, The Efficiency of Equity Linked Compensation: Understanding the Full Cost of Awarding Executive Stock Options, *Financial Management* 30, 5–44. 31

- Miao, Jianjun, and Neng Wang, 2007, Investment, Consumption and Hedging under Incomplete Markets, *Journal of Financial Economics* 86, 608–642. 10, 11, 34, 35, 38, 40, 50, 58, 60, 61, 62, 133, 143, 144, 183
- Milne, Alistair, and A Elizabeth Whalley, 2000, Time to build, option value and investment decisions: a comment, *Journal of Financial Economics* 56, 325–332. 27
- Munk, Claus, 1997a, Numerical Methods for Continuous Time, Continuous State Stochastic Control Problems, *University of Southern Denmark Working Paper*. 190
- , 1997b, Optimal consumption/portfolio policies and contingent claims pricing and hedging in incomplete markets, Ph.D. thesis School of Business and Economics, Odense University. 190, 196, 197, 198
- , 2000, Optimal consumption/investment policies with undiversifiable income risk and liquidity constraints, *Journal of Economic Dynamics & Control* 24, 1315–1343. 10, 33, 62, 179, 180, 182, 184, 185, 190, 192, 196, 198
- , 2003, The Markov chain approximation approach for numerical solution of stochastic control problems: experiences from Merton’s problem, *Applied Mathematics and Computing* 136, 47–77. 190, 204
- Murphy, Kevin, 1999, *Handbook of Labor Economics*. chap. Executive Compensation, pp. 2485–2563 (North-Holland, Amsterdam). 41
- Novy-Marx, Robert, 2007, An equilibrium model of investment under uncertainty, *Review of Financial Studies* 20, 1461–1502. 25
- Pindyck, Robert, 1993, Investments of Uncertain Cost, *Journal of Financial Economics* 34, 53–76. 27

- Poterba, James M., and Lawrence H. Summers, 1995, A CEO Survey of US Companies' Time Horizons and Hurdle Rates, *Sloan Management Review*. 24
- Ross, Steven A, 1976, The Arbitrage Theory of Capital Asset Pricing, *Journal of Economic Theory* pp. 343–362. 18
- SCHWARTZ, EDUARDO S., AND LENOS TRIGEORGIS (eds.), 2004, *Real Options and Investment under Uncertainty: Classical Readings and Recent Contributions* (The MIT Press). 24
- Svensson, Lars E.O., and Ingrid M. Werner, 1993, Nontraded assets in incomplete markets: Pricing and portfolio choice, *European Economic Review* 37, 1149–1168. 10
- Tobin, James, 1958, Liquidity preference as behavior towards risk, *The Review of Economic Studies* 25, 65–86. 30
- Triantis, Alexander, 2005, Realizing the Potential of Real Options: Does Theory Meet Practice?, *Journal of Applied Corporate Finance* 17, 8–16. 24
- Trigeorgis, Lenos, 1996, *Real Options: Managerial Flexibility and Strategy in Resource Allocation* (The MIT Press). 23, 220, 221
- Wall Street Journal, and Mercer Consulting, 2004, CEO Compensation survey and Trends, www.mercerhr.com. 4
- Wonder, Nicholas, 2006, Contracting on real option payoffs, *Journal of Economics and Business* 58, 20–35. 42