Essays On Contracting For Experimentation

THE UNIVERSITY
of EDINBURGH

Aodi Tang

School of Economics

University of Edinburgh

A thesis submitted for the degree of

Doctor of Philosophy

2017

Declaration

I declare that this thesis has been composed solely by myself and that it has not been submitted, in whole or in part, in any previous application for a degree. Except where states otherwise by reference or acknowledgement, the work presented is entirely my own.

Aodi Tang

Acknowledgements

First and foremost, I would like to thank my supervisors Tim Worrall and Jonathan Thomas for their continuous support and guidance during the past four years PhD studies. I thank them for always being supportive and letting me explore the research questions that I truly interest in. I am very grateful that they are always willing to spare time for a meeting with me. This thesis can not be completed without their insightful guidance, challenging questions and inspiring ideas.

I thank Francesca Flamini and Ed Hopkins for being such kind and helpful Viva examiners. I thank School of Economics for providing such a nice and interactive research environment where all the professors are approachable and friendly. I thank Andrew Clausen, Ina Taneva, Jozsef Sakovics, Tatiana Kornienko, Andrew Snell and Michele Belot for their helpful comments and tips at different stages of PhD. I thank Ahmed Anwar for his guidance on tutoring. I benefit a lot from the first-year PhD training courses and plenty of presentation opportunities provided by the school. I thank Edinburgh University for the Principal's Career Development Scholarships and School of Economics for extra research funding. I thank Johannes Horner for introducing the frontier research area, strategic experimentation, to me at the PhD conference in York. I thank all my office mates in Room 4.12 and PhD classmates for their friendship and support.

Finally, I would like to thank my husband, Xin Wang, for supporting and encouraging me over the past eight years. Edinburgh has become my second home town because of his companion. I owe special thanks to my mother, Huiqin Zhong, and my father, Dong Tang. I have the courage to pursue my dream because of their unconditional support and love.

To my loving husband and parents

Abstract

This thesis is composed of four chapters and addresses the contracting issue under strategic experimentation.

The first chapter presents an overview of the thesis and introduces the strategic bandit model, which is commonly adopted in the other three chapters. The chapter also previews the main results and implications of the thesis.

The second chapter discusses the contracting issue between a principal and a team of agents where the actions of agents are unobservable to the principal. The main contribution of this chapter is to fill the gap of strategic experimentation literature by introducing the free-rider problem in teamwork. The chapter first deals with the optimal hiring choice of the principal under perfect information. Since the belief of the state being good decreases if no one succeeds over time, the paper shows that the principal tends to hire fewer agents in response to the downward-adjusted posterior belief. When the principal can neither monitor the agents' actions nor distinguish the agents who succeed, this chapter shows the optimal incentivising contract consists of an upfront payment from the agents to the principal, a bonus to every agent conditioning on success and a stopping time. Under this contract, the principal can implement first-best experimentation and incentivise all agents to work until the optimal stopping time.

The third and fourth chapters discuss the financial contracting issue in innovation where an innovator requires external funding from an investor. The third chapter adopts a "bad news" exponential bandit to study the financial contracting under adverse selection between the innovator and the investor. The innovator, owns the innovation project, is privately informed of either a high or low prior belief of the good state but seeks a large amount of external investment from the less-informed investor. Experimentation is conducted by the innovator using internal funding before the external investment. The posterior belief about the good state increases in the amount of internal funding if no bad news arrives during experimentation, but the project will be abandoned as long as bad news

arrives. The chapter shows that the amount of internal funding can be used by the investor to separate the agents with different priors. Under the unique least-costly separating equilibrium, the high-prior innovator spends even more than the low-prior first-best internal funding in order to deter the low-prior one from mimicking, and the low-prior one remains at his first-best. This chapter enriches the financial experimentation literature by proposing internal funding as a novel signalling tool and establishing a Pareto dominating separating equilibrium.

The fourth chapter studies a multi-stage innovation financing problem between an agent and an investor with asymmetric information on the progress of the project. The innovation is comprised of two stages where the agent needs to complete the first development stage in order to proceed to the second experiment stage. The model assumes that the completion of the first stage can be early or late following a binary distribution, and the arrival of success in the experimentation stage follows a "good news" exponential bandit. Each period, a fixed amount of investment is needed from the investor. However, the investor can not observe nor verify the project progress. The chapter shows that the optimal incentive-compatible contract consists of differential maximum funding periods in the event of early and late completion of the first stage respectively and subsequent bonuses to the investor conditioning on a success in the second stage. We prove that the first-best experimentation time is attainable as long as the bonus of the late completion exceeds that of the early completion, and the difference between the two bonuses should be confined within a certain range. In the extension, we consider the case when the first stage completion time is informative such that an early completion indicates a higher prior in the good state than the late completion. Under imperfect information, the agent has a stronger incentive to mimic the early completion if the first stage is completed late as a longer experimentation time will be granted according the first-best contract. The chapter proves that the first-best is still achievable under a similar bonus contract but the difference between the two bonuses becomes smaller. This chapter contributes to the experimentation financing literature including the information imperfectness on project progress and multi-stage spillover effects.

Contents

List of Figures

Chapter 1

Introduction

1.1 Overview

Uncertainties are prevalent in the process of innovating. An innovation project can be promising or not, which is usually unknown to innovators at the beginning. This thesis adopts an exponential bandit to model the experiment process in innovation and focuses on the agency problems where the principal and agents involved in experimentation have asymmetric information. Three different agency problems are addressed in the thesis respectively: the free-rider problem in team experimentation, the adverse selection problem in financing experimentation and the imperfect information problem under a multistage project financing. This thesis concentrates on the design of an optimal contract that eliminates the relevant agency problems and optimises the welfare of contracting parties.

The aim of this chapter is to provide an overview of the thesis, which is organised as follows: Section 1.2 reviews the general exponential bandit game; Section 1.3 and 1.4 introduce the motivation and main contributions of the thesis under the topic of team experimentation and innovation financing respectively.

1.2 The exponential bandit game

This section provides an introduction of the general strategic bandit model. A *multi-armed bandit* represents a sequential game where a player chooses between a number of arms with unknown returns. The selected arm will randomly generate a reward follows a probability distribution. A *two-armded bandit* often involves a safe arm, which generates a constant relative low return, and a risky arm, which randomly generates a relative high return. The player in the two-armed bandit game faces a trade-off between exploiting the safe arm and exploring the risky arm in order to maximise his expected payoff. The trade-off between exploitation and exploration is ubiquitous in real life. For example, people often face the choices of visiting the regular coffee shop or trying a newly-opened one; researchers can build work on an existing discovery or explore a novel frontier area.

This thesis adopts the two-armed exponential bandit by Keller et al. (2005) and Keller and Rady (2010) to model the learning and experimenting process in innovation. There are two states of the world, either good or bad, which would impact the expected payoff of the risky arm. If the state is good, the risky arm is good and will generate a high payoff following the exponential distribution; otherwise, the state is bad and the risky arm will generate nothing. The state is chosen by natural and fixed but is unknown the public. The mean of the lump-sum revenue generated from the risky arm is assumed to be strictly larger than it from the safe arm. Thus, the players have to strategically allocate resources on the two arms. Bolton and Harris (1999), Keller et al. (2005), Keller and Rady (2010) collectively build the seminal strategic experimentation model with the two-armed bandit. These papers mainly differ in the probability distributions of the risky arm where different equilibrium results are generated: in Bolton and Harris (1999), a success arrives following a Brownian process with an unknown drift, while in Keller et al. (2005) and Keller and Rady (2010), the arrival of either a success or failure follows an exponential distribution. We adopt the exponential bandit in the thesis as it is more tractable and has been widely applied in the innovation literature (Bergemann and Hege, 2005; Horner and Samuelson, 2013). Keller et al. (2005) and Keller and Rady (2010) both focus on the equilibrium analysis in a multi-player cooperative game where the free-rider problem exists. Keller et al. (2005) develops the so-called "good news" model where a good news is equivalent to a success or a breakthrough of the project. Chapter 2 and 4 both adopt the "good news" model. Keller and Rady (2010) builds the "bad news" model to describe the arrival of a failure or breakdown of the project, which is used in Chapter 3 of the thesis. We choose either the good news or the bad news arrival that fits best to the specific problem and context of each chapter, and we will give a thorough explanation later.

The main difference between the "good news" and "bad news" model is the evolution of posterior belief updating. Suppose the player is given a prior belief of the good state as p_0. Under the "good news" model, the player exerts effort and waits for a good news, which indicates a breakthrough or success of the innovation. Thus, the posterior belief of the good state decreases as long as no good news arrives. However, in the "bad news" model, an arrival of the bad news (breakdown/failure) shows the state is bad, which may result in an abandonment of the project, but the posterior belief of the good state evolves upward overtime as long as no bad news arrives. The opposite posterior belief updating gives us opportunities to model innovation under different contexts.

The "good news" model is adopted when the innovation project owner aims to achieve a breakthrough, but he needs to set up a stopping time after which the project generates non-positive profit. In Chapter 2, the principal owns the innovation project and employs a group of agents to explore the unknown state of the world. The optimal stopping time is determined after which the marginal expected benefit generated from the project is less than the marginal hiring cost, which is included in the optimal contract. In Chapter 4, the innovator owns the project and relies on the external funding from an investor to run experiment. In order to borrow from the investor, the innovators needs to write an incentive-compatible financial contract that makes the investor want to participate. As the innovator is short of funding, she could compensate the investor only if the project succeeds and a large amount of revenue is generated. Thus, Chapter 4 also uses the "good news" model to describe the innovator's experimentation in order to get a breakthrough. The "bad news" model is used to when the project holder conducts the experiment to check if the innovation involves any fault or error, which may result a breakdown or failure, before investment. Chapter 3 investigates an optimal financial contract for the innovator to get external investment under the adverse selection. Given a fixed period of experimentation, the innovator conducts experiment to guarantee that the project is worth investing as long as no bad news arrives. Thus, we apply the "bad news" model in this context. Note that in all these three chapters, pulling the safe arm is equivalent to take an outside option, which is assumed to generate zero payoff.

1.3 Chapter 2: Team experimentation

The free-rider problem is often the key concern of the principal especially when she cannot fully monitor the actions of team workers. This chapter considers the free-rider problem in innovation setup when team workers simultaneously conduct experiment to learn the unknown state of the world and the outcome of the project is uncertain. There are several unique features of the team experimentation model in this chapter, which are absent in the typical asymmetric information in teamwork research, such as Holmstrom (1982), Manso (2011): the outcomes of the innovation depends on the unknown state of the world, only when the state is good, the project can succeed; the principal can neither observe the actions of the agent nor distinguish the agents who make success; the principal can choose the number of agents to employ.

This work shows that the principal would equally distribute a reward to each agent when she cannot tell who contribute to the success under perfect information. In this case, the principal extracts the entire surplus and achieves the first-best experimentation. However, under imperfect information, the actions of the agents are unobservable to the principal. Assume the primary concern of the principal is to achieve the first-best level of experimentation, that is, to incentivise the agents to work until the first-best stopping time. We show that the target can be obtained as long as the principal rewards each agent with a relative large bonus whenever a success comes, and the value of the optimal bonus exceeds the revenue of the project. As the principal holds the entire bargaining power, she can charge each agent a large lump-sum upfront payment and promise the full-incentivising bonus to each participant conditional on a success, which makes each agent gain zero payoff in expectation and motivate the team to work until the first-best stopping time.

This chapter also provides insights on the optimal hiring choice of the principal in team experimentation. We endogenize the number of agents to employ and let the principal to decide the optimal hiring number and stopping time at the beginning. The result indicates that the optimal hiring choice is negatively correlated with the optimal experimentation time. Moreover, the time preference of the principal plays an important role in making the hiring choice: for a relative patient principal, fewer agents will be hired initially and the experiment will last longer. In the extension part, we allow the hiring choice to vary each period and consider the optimal hiring decision of a myopic principal. Since the belief of the project being good becomes pessimistic while no body succeeds, the model shows that the principal tends to hire fewer agents, which means that the termination of contract and firing may occur during the experimentation.

To summarise, this chapter provides an optimal incentivising contract to eliminate the free-rider problem in team experimentation. This work mainly contributes to the recent popular area of strategic experimentation by adding the study of the free-rider problem in teamwork. Moreover, the results of this chapter can be applied to optimise the hiring choice and the compensation scheme of innovation projects in real life.

1.4 Chapter 3 and 4: Innovation financing

The third and fourth chapters both focus on the financing of innovation, which involves an innovator (project owner/borrower) and an investor (financier/lender). It is motivated from the practical fact that it is very difficult for small innovation firms to get externally funded, and the main reason might be the information imperfection between investors and innovators (Hall and Lerner, 2010). These two chapters aim to design an optimal incentivising financial contract which eliminates the adverse selection problem on the types of the innovation project (Chapter 3) and the imperfect information on the project progress (Chapter 4), and hence improves the financing process.

In Chapter 3, the innovator is privately informed of a high or low prior belief of the risky project being good and has to signal this information in order to get funded by the less-informed investor. As the innovator normally covers the experiment cost, the choice of internal funding spent on the experimentation can be adopted to as a signal of the confidence in the project. Under the "bad news" setup, learning speed increases in the amount of internal financing allocated, and the posterior belief of the good state increases as long as no bad news arrives during experimentation. External investment from the investor will be triggered if no bad news arrives until an endogenously fixed deadline.

This chapter first presents the innovator's optimal options of the internal financing under perfect information: in the first-best, the high-type innovator tends to allocate less internal financing based on the higher prior belief compared with the low-type innovator. However, under imperfect information, the low-type would have incentive to mimic to be a high-type by allocating less amount of internal financing, which results in a welfare loss of the investor. We then consider an incentive-compatible financial contract that separates the high-type from the low-type and makes the investor break-even in expectation. By including the amount of internal financing in the financial contract, the innovator could send useful information upfront to the investor, which indicates his confidence in the prospect of the project. Thus, we forms up a high-type optimisation programme by choosing the amount of internal financing and a revenue share in the event of success subject to a participation constraint of the investor and the incentive compatibility constraints of both types. This chapter is able to establish an unique least costly separating equilibrium under imperfect information which indicates that the high-type at the optimum would commit to a higher level of internal financing, which exceeds the amount of the low-type first-best internal financing, and the low-type remains at the low-type first-best. In other words, the investor can distinguish between different types of innovators via the signal of internal funding. There also exist multiple pooling equilibria where cross-subsidisation from the high to the low is commonly involved, all of which are prove to be Pareto dominated by the least costly separating equilibrium.

As it is technically challenging to determine a Pareto dominating equilibrium in the signalling game, this chapter adopts the informed-principal optimisation approach by Maskin and Tirole (1992) where the innovator has the entire bargaining power and proposes the financial contract to a less-informed investor. The investor is endowed with a prior belief of the type of the innovator that with probability of α, the innovator is a high-type; with probability of $(1-\alpha)$, the innovator is a low-type. At the beginning, the innovator proposes a menu of two contracts, including the high-type and low-type optimum, and executes one of the contracts once the financial contract is signed. We first consider the so-called low-information intensity equilibrium where each type at the minimum can get regardless of the prior belief of the investor. The equilibrium allocation shows that to prevent the low-type from deviating, the low-type is guaranteed with the low-type first-best payoff, but the high-type sacrifices a lot. However, the high-type can do better than the low-information intensity equilibrium by letting the investor break-even in expectation. In other words, the investor can make a loss on the low-type contract but is still willing to participate as long as he earns a non-negative payoff in expectation given a menu of two contracts. Thus, we form the high-type optimisation problem, which includes the average participation constraint of the investor, the incentive compatibility constrains of both types and the constraint that guarantees the low-type at least the low-type-first-best payoff. Thus, the unique least costly separating equilibrium is realised from this optimisation program.

Chapter 4 considers another imperfect information problem, which is ubiquitous in a multistage innovation financing, that is the progress of the innovation is unobservable to the investor. This chapter aims to provide an optimal financial contract that motivates the innovator to reveal the innovation progress truthfully and enables a non-negative payoff of the investor. The chapter models a two-stage innovation where the first-stage involves no experimentation and can be completed with certainty, but the second-stage requires the innovator to learn the unknown state of the world. Only in the good state, the innovation project can succeed and generate a large amount of revenue. The "good news" bandit is adopted to show the arrival of success in the second experimentation stage. The innovator needs a constant amount of funding to run the project from the investor until the first-best stopping time. This chapter shows that the first-best financial contract should include a bonus transfer from the innovator to the investor at the breakthrough and maximum funding times (the first-best stopping times) conditioning on the completion times of the first stage, under this contract the innovator captures the entire surplus. However, if the information is imperfect, the investor can neither observe nor verify the completion of the first stage. The agent may hide an early completion or pretend to finish the first stage early in order to get a higher profit. To prevent such deviation, this chapter establishes that with some variance on the bonus, the above financial contract can still obtain the first-best and motive a truthful revelation of the innovation progress. In the extension, this chapter consider a spillover effect across two stages when an early completion of the first stage indicates a higher prior belief of the good state. In this case, the progress of the innovation contains information about the good state, which would affect the optimal stopping times of the second experimentation stage. The agent would have a higher incentive to mimic an early arrival due to a higher prior belief and an extended stopping time. This chapter offers an optimal solution to eliminate such potential deviation and achieve the first-best, which involves the similar bonus contract but with some dispersion on the bonus.

Chapter 4 mainly contributes to the multistage innovation financing under imperfect information, which is a relatively new area. Most of the papers in this area focus on the moral hazard problem where the optimal contract of the investor is to incentivise the innovators to work until an optimal stopping time (Green and Taylor, 2016; Moroni, 2016; Wolf, 2017). However, this chapter focuses on the problem of imperfect information on the progress of the innovation project and the design of an innovator-optimum contract that enables a truthful information disclosure. The results of this chapter shed lights on the design of compensation scheme especially for the long-term financial contract and help the competent innovators to get funded more easily under information imperfection.

Chapter 2

Optimal contracts for team experimentation

2.1 Introduction

2.1.1 Motivation

Nowadays, uncertainty involved in innovation has drawn much attention and interest of economic theorists as it may result in the under-provision of research funding and the early abandonment of research, which are arising from asymmetric information between the worker and project owner (Hall and Lerner, 2010). In real life, uncertainty is quite prevalent in various industries. For instance, a pharmaceutical company wants to engage in a brand-new drug field, which requires substantive research and technological development. However, the future of this area can be promising or not as it relates to various unknown factors, such as economic background and prospects, intensity of future competition and so on. Uncertainty also prevails among the discovery of an oil field when workers explore resources by testing and drilling wells. The field may be abundant or barren, which is determined by nature but unknown at the outset of the project. In this paper, we use the unknown state of the world to describe the above-mentioned uncertainty in innovation. In the good state, a breakthrough will happen with higher probability; otherwise, in the bad state, the project cannot succeed. As the owner of the innovation project, the principal would like to explore the unknown state of the world and optimise his expected payoff. The agents recruited can gradually learn the state of world through experimentation and acquire the information for the principal, but they may be less motivated under imperfect monitoring. Due to imperfect information, it is difficult for the principal to discover whether the failure of the project is because of agents' shirking or the bad state of the world. Thus, a number of interesting contracting problems arise in team experimentation, such as how many agents should be employed each period, how will the principal motivate agents and realise the optimal payoff?

This paper builds on the typical principal-agent framework to study team experimentation when the actions of agents cannot be fully monitored by the principal. In the model, there are a group of agents who work together to learn the state of the world. The arrival of success is random and depends on the state of world as well as how much effort in total is exerted such that the more effort is allocated, more likely is success. The model adopts a "good news" exponential bandit model, so that in a good state success occurs with positive probability, but the posterior belief of the good project will decrease over time if no success arrives (Keller et al., 2005). The moral hazard problem arises when the principal cannot observe the actions of the workers such that some of them may shirk and free ride on the others, which has the potential to create asymmetric beliefs about the good state between the two parties. Additionally, due to the imperfect monitoring, the principal cannot even distinguish the agents who make the success, which makes contracting more complicated. In the seminal teamwork paper by Holmstrom (1982), the principal creates incentives by generously rewarding each agent involved in which case the budget is unbalanced. This paper follows Holmstrom's motivation scheme to resolve the free-rider issue, that is, giving agents a substantial reward whenever there is a success. Thus, we apply such a bonus contract to tackle the two layers of moral hazard problem in our model.

This work starts with a relatively straightforward case when the number of workers is exogenously set. One might relate it to the practical scenario when there is a fixed amount of machines available for the production which exogenously determines the number of workers. The first-best stopping time is defined as the socially efficient stopping time before which the n-agent-teamwork constantly generates a positive expected surplus. When the principal cannot identify the agents who make the success, to motivate each agent work until the first-best stopping time, our analysis indicates that the principal should equally distribute the revenue to each agent in the event of success. As there is no credit constraint imposed on the principal nor agents, our model allows for an upfront transfer from the agent to the principal. In this case, the principal can propose a bonus contract to each agent with an upfront transfer from the agents, which guarantees a non-negative payoff and the participation of the agents, and the bonus that equally shares the revenue among the team in the event of success. Under this bonus contract, the principal is able to extract the entire surplus and implement the first best stopping time.

However, under imperfect information, the agents have incentive to free-ride on the others and shirk. This chapter shows that providing a large reward conditioning on success is still an effective way to eliminate the free-rider issue. Section 2.4 derives the minimum bonus required to incentivise agents for the principal. The approach is to characterise the symmetric agent's optimal stopping time assuming other agents are exerting effort. Then the fundamental issue is to define an optimal bonus which keeps individual agent working until the first-best stopping time. The result implies that the optimal bonus required should be greater than the revenue generated from the project if more than one agent is employed, which matches the statement by Holmstrom (1982) that the principal acts as a "budget breaker" in teamwork so as to provide incentives. Moreover, the optimal bonus is increasing in the size of the team, which means that the larger the team size, the severer the incentivise issue, hence larger bonus should be imposed. In this case, the principal can optimise his payoff and extract the entire surplus by charging each agent a larger upfront payment, which balances the participation constraint of each agent. Intuitively, by signing the bonus contract, the agent pays a deposit to the principal and will be rewarded with a large bonus as long as a success arrives. This bonus contract design also appears in Halac et al. (2016a). Their model focuses on the strategic experimentation problem under imperfect information between the principal and a single agent where the optimal contract is characterised by an optimal stopping time and a front-loading transfer. However, as a major departure from Halac et al. (2016a), this paper explores the optimal contracts under multi-agent teamwork where the identity of the agents who make success is unknown to the principal. In the team production, the principal would incur some other contracting problems, involving how to determine the optimal number of workers and identify the worker that contributes to the final success. The rest of the paper will tackle each problem respectively.

The analysis becomes more complicated when the hiring choice is endogenously determined by the principal. There exists a trade-off between accelerated experimentation and increased cost. The discrete-time model setting allows us to analyse the principal's hiring choice, that is, a profit-maximising principal would choose an optimal number of agents, n, at the beginning based on the total expected discounted payoff, which also determines the efficient stopping time. For simplicity, we first assume that once the optimal n is chosen, the principal should keep these n-agents as a team until the end. Under perfect information, Proposition 2 shows that the optimal experimentation time decreases in the optimal number of agents. Intuitively, teamwork shortens the experimenting time and increases the total success probability. Furthermore, the comparative static study indicates that principal's time preference would play an important role in his hiring decision: for an impatient principal, he would choose larger amount of agents to experiment for shorter periods since future success is less valuable to him; in contrast, as a sufficiently patient principal he would rather save some effort costs by employing fewer agents and having a longer experiment.

As an extension, we explore the case where a myopic social planner optimally chooses his hiring choice over time. In this case, the current numbers employed by the principal will just affect the success probability this period as the planner does not take future payoff and belief updating into consideration. As the posterior belief of the good state decreases over time as long as no good news comes, the planner hires fewer agents accordingly. In the simulation, we are able to capture a clear decreasing hiring pattern based on the downward-adjusted belief. Moreover, there is a steep decrease of the hiring choice if the revenue of the project is very high. Intuitively, the planner would hire more agents in the first period if the revenue is higher, whereas if it fails, the posterior belief will drop by a lot, which results in much fewer agents employed in the second period.

This chapter belongs to the literature of moral hazard in team and proposes an optimal contract eliminating the free-rider problem. This work fills the gap of this literature by considering a special imperfect information case where the principal cannot distinguish who make the success. Moreover, this chapter explores the team work under strategic experimentation where there exists uncertainty on the outcome of the project. The results of the principal's optimal hiring choice also provide some implications on the contracting and recruiting strategies in a team innovation.

2.1.2 Related literature

The paper first contributes to the literature on strategic experimentation. The research in this area mainly concentrates on the equilibrium analysis where each player independently allocates resource between the safe arm and risky arm as a best response to the actions of the others. Bolton and Harris (1999), Keller et al. (2005), Keller and Rady (2010) built the standard models based on the so-called two-armed bandit model. These papers mainly differ in the distributions of the success arrival. The "good news" model by Keller et al. (2005) is the closet to our experimentation setting. In their model, the breakthrough from the risky arm arrives randomly and follows an exponential distribution, and the average payoff generated from the risky arm is strictly larger than that of the safe arm, so the players would play the risky arm and gradually become pessimistic about the good state as long as no success arrives. However, there is no moral hazard problem exists in the above-mentioned paper since either the actions of players or the outcomes of project are public observable. Under the similar bandit model setup, this paper focuses on the design of an optimal contract that prevents team members from free riding.

This chapter adds the uncertainty of the innovation into the area of moral-hazard-in-teams. Since the seminal paper by Holmstrom (1982), researchers have proposed various incentive contracts to mitigate the moral hazard problem in teamwork, such as relational contracts based on mutual trust among team members suggested by Baker et al. (2002), Levin (2003) and Rayo (2007), using implicit contracts where the subjective performance evaluation scheme is implemented (MacLeod and Malcomson, 1989; Bentley MacLeod, 2003). There are several papers dealing with the moral hazard issue in the team learning environment. Bonatti and Horner (2011) claims that endogenously forming a deadline is a feasible way to get rid of free-riders among a team. However, the principal, who holds no bargaining power and acts as an outsider, plays no role in their model. Moreover, the multi-stage innovation in teamwork has been modelled by Moroni (2014). With a limited liability constraint and varying bonus setup, their result indicates that the principal tends to employ fewer agents in the early experiment stage as it gets difficult to suceed later. The employment choice is adjusted upwards in accordance with the increasing posterior belief based on the "bad news" bandit model. As a contrast, in our model, the decreasing hiring is driven by the more pessimistic beliefs of the good state under the "good news" bandit. Additionally, there is just one stage in our model, which means the project stops as long as a breakthrough arrives. This setup allows us to discuss the imperfect information about the team members who make the breakthrough in a straightforward paradigm.

Our paper also relates to Halac et al. (2016b), which discusses team experimentation under a contest design. The principal with a fixed budget maximises the probability of success by appropriately distributing the prize amongst the agents. In their model, the principal has full information about contestants' performance and only values the one makes the first success. However, in our model, how to mitigate the free-rider issue is the major focus as the principal has imperfect monitoring over agents' actions. In addition, the number of agents is exogenous in their model where neither entry nor exit is allowed.

The next section describes the set-up and environment of the model. Section 2.3 defines the first-best in the team experimentation model. Sections 2.4 discusses the free-rider problem and derives optimal contracts. Section 2.5 is the extension and conclusion.

2.2 Environment

The principal in charge of the risky project wants to hire a group of homogeneous agents to explore the unknown state of world. Adopting the terminology from the bandit game literature, exploring the risky project is similar as pulling the risky arm: under the good state, the risky project is supposed to succeed with some probability and generates relatively high lump-sum h randomly; in the bad state the project will never succeed no matter how much effort the agents allocate, and the principal will get nothing. The agents can pull the safe arm, which generates zero payoff as an outside option. To avoid complexity, this model uses discrete-time such that the success can only be observed at the end of each period. As a group of homogeneous agents work together on the project, individual contribution to the success each period is denoted as λ where $\lambda \in (0, 1)$. Assume that the arrival of the success follows an exponential distribution in the good state. Let function f be a total success probability function in the good state, which depends on the number of agents working per period and the success arrival rate of individual agent. Define $f = 1 - e^{-n\lambda}$ with the total success arrival rate λn if the number of n agents exerting effort, and f ranges within $(0, 1)$.

Once hired, the agents choose whether to work each period, that is, we set agent i's binary actions as: $x_i \in \{0, 1\}$. He could either pull the risky arm and explore the project by working where $x_i = 1$, or take the safe arm by shirking where $x_i = 0$. Let $x_{i,t}$ be agent i's action at period t. Working on the project each period incurs a fixed cost c. Then we need to define the aggregate success probability in the n-agent-teamwork problem given the state is good. As long as one of the agents makes a breakthrough, the project succeeds, and the experiment stops. Thus, the aggregate success probability at any time t depends on each agent's effort at that period. Let f_t denotes per-period aggregate success probability when efforts are publicly known. In particular, if the principal keeps the same amount of agents and everyone works all the time, aggregate success probability any period will be constant. If n is treated as endogenous, which means that the principal chooses how many agents to hire, $f(n)$ as defined is a concave function in n, i.e $f'(n) > 0$, $f''(n) < 0$, and the per-period aggregate success probability will be the same so long as the principal keeps these n-agents till the end.

If the actions of the agents cannot be monitored, it is hard for the principal tell who contributes the most to success. In other words, it is difficult to separate individual's contribution to the final breakthrough in teamwork. For instance, a group of workers are employed to drill a water well. If there is a discovery of water, they might be equally rewarded by the manager due to asymmetric information on workers' effort.

2.2.1 Timing

Both parties share a common prior belief about the good state $\beta_0 \in (0, 1)$. Period 0 is defined as a negotiation period as no experimentation occurs, thus $\beta_0 = \beta_1$. The project will be terminated when the expected payoff of the principal over the next period becomes non-positive, and the optimal stopping time is defined as t^*. Therefore, $\{0, 1, ..., t^*\}$ is the time line of the project. All players discount future at the same rate $\delta \in [0, 1]$.

When working on the project, agents gradually learn the state of world or the quality of project as long as there is no success. They will update beliefs independently based on the available information. Intuitively, they have two ways of updating their beliefs: they could learn by experimenting individually, or by just observing whether the project succeeds in that period. Given that only the outcomes of the project are public observable, a typical free-rider problem arises where agents have incentives to take advantages of the information imperfectness and shirk every period. Therefore, the principal needs take actions to motivate agents and avoid losing profit.

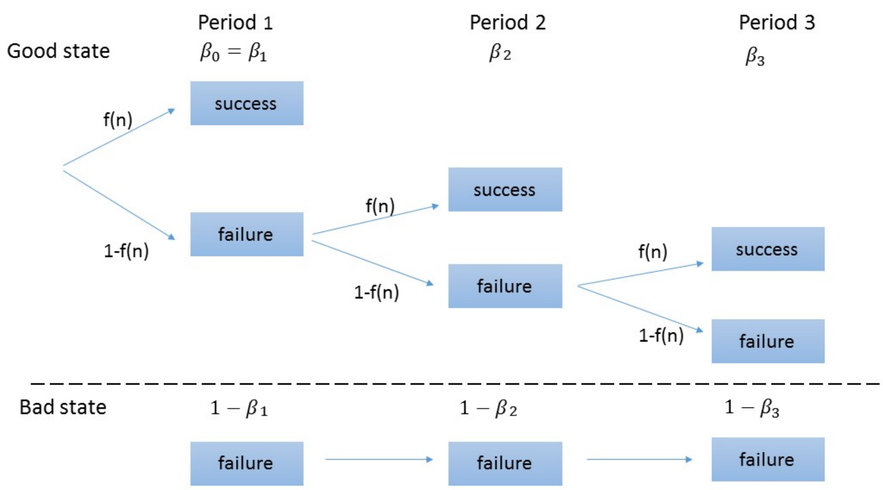

Figure 2.1: Timing

Figure 2.1 illustrates the timing of the problem where there are n agents exert effort each period. β_0 is the prior belief of the good state at the first period. In the good state, with f probability that the project will succeed if n agents choose to work; with $(1-f)$ probability it will fail in the good state. The second period of experiment starts conditional on the failure of the first trial, however, the belief about the good state will be updated downward accordingly to β_2. If the project succeeds, it generates the surplus $h-nc$ in that period. However, the project fails every time under the bad state as shown in the bottom of the time-line, and the cost nc is paid every period.

2.2.2 Contracting

The model adopts the typical take-it-or-leave-it contract where the agents have no bargaining power, and the principal proposes the contract and extracts the entire surplus under perfect information. The contract is signed upfront and both parties have to fully commit to the contractual terms. The contract that each agent gets is denoted by C, which consists of an upfront payment W_0, a fixed bonus b conditional on the success of the project and a specific termination time t, i.e $C = (W_0, b, t)$. Note that W_0 can be negative such that an upfront transfer from the agents to the principal is allowed. In other words, the principal could sell the project to n-agent at the price of W_0 as there is no limited liability constraint. In practise, the negative transfer can be viewed as a type of sunk cost that the agents paid to participate the project: workers usually need to spare some effort and get trained in order to be qualified to operate the machine or conduct the project. A bonus is a commonly used way of incentivising agents. The size of the bonus directly affects how long the agent would experiment. If the bonus is sufficiently large, the agents are willing to work a longer time although the belief about the good state is decreasing overtime when no success comes. However, the bonus is negatively correlated with the total payoffs that the principal could get. Thus, the principal faces a trade-off between providing the bonus and motivating agents.

Let U_0 be the total discounted expected payoff of the agent according to the contract above as follows:

$$U_0 = \sum_{t=1}^{t^*} \delta^t \{(1-f(n))^{t-1}\beta_0(f(n)b-c) - (1-\beta_0)c\} + W_0, \tag{2.1}$$

where the term in the big bracket is the expected payoff each period. $(1-f(n))^{t-1}$ denotes the probability that no success arrives before time t. Conditioning on this failure probability, $\beta_0(f(n)b-c)$ is the expected payoff at time t when the state is good, while $(1-\beta_0)c$ is the bad-state payoff.

We use Π_0 denotes the discounted expected total payoff of the principal under the bonus contract:

$$\Pi_0 = \sum_{t=1}^{t^*} \delta^t (1-f(n))^{t-1} \beta_0 f(n)(h-nb) - nW_0. \tag{2.2}$$

The principal receives h amount of revenue and rewards the bonus b to everyone in the team when the project succeeds.

2.2.3 Learning and belief updating

In the n-agent-teamwork, each agent's belief about the good state is updated via his own effort as well as the actions of the others if the information is perfect. Consider a perfect information scenario where the actions and outcomes of the project are public information. Suppose every one works in the first period, but no success arrives. A rational agent i adjusts down his belief about the good state given the others' working and the outcome. We use Bayes rule to formally characterise the belief updating process at any date t. Let agent i's posterior belief about good state at the beginning of time t conditional on previous failures be denoted by $\overline{\beta}_{i,t}$, which can be calculated as the following:

$$\overline{\beta}_{i,t}(\mathbf{x_{i,t-1}}, \mathbf{x_{-i,t-1}}) = \frac{\beta_0 \Pi_{s=1}^{t-1}(1-f(x_{i,s}, x_{-i,s}))}{\beta_0 \Pi_{s=1}^{t-1}(1-f(x_{i,s}, x_{-i,s})) + (1-\beta_0)},$$

where $\mathbf{x_{i,t-1}}$ denotes agent i's set of actions before period t such that $\mathbf{x_{i,t-1}} = [x_{i,1}, x_{i,2}, ..., x_{i,t-1}]$, and $\mathbf{x_{-i,t-1}}$ denotes the other $n-1$ agents' actions till time t, which is a $(t-1)$ by $(n-1)$ matrix. The numerator on the right hand side is the probability that the success never happens before period t conditional on the state being good, and the denominator is the unconditional probability of failure, which consists of two events: the failure conditional on the state being good, which is the same as the numerator; the failure conditional on the bad state, which happens with probability 1. If the project succeeds at period $t-1$, agent i's belief about the good state $\overline{\beta}_{i,t}$ will jump to 1 at the beginning of time t.

Under perfect information, if all agents are motivated to work every period, the belief about the good state will be updated symmetrically, and omitting the subscript i, the updated belief at period t will be :

$$\overline{\beta}_t = \frac{\beta_0(1-f)^{t-1}}{\beta_0(1-f)^{t-1} + (1-\beta_0)},$$

where f is the per-period aggregate success probability of n-agent's working conditional on the good state: $f = 1 - e^{-\lambda n}$.

If the principal could endogenously determine n at the beginning, per-period updated belief varies with n, the intensity of experimentation, which takes the form:

$$\overline{\beta}_t(n) = \frac{\beta_0(1-f(n))^{t-1}}{\beta_0(1-f(n))^{t-1} + (1-\beta_0)},$$

where $f(n)$ is per-period aggregate success probability expressed as a concave function of n: $f(n) = 1 - e^{-\lambda n}$. As n increases, $f(n)$ increases, which triggers a faster decrease of $\overline{\beta}_t(n)$. Intuitively, more agents induce an accelerated speed of learning, given the same times of failure, the belief about the good state would decrease more quickly.

The next section will explore the efficient stopping time on the equilibrium path and the first-best policy.

2.3 First-best policy

2.3.1 Socially efficient stopping time

In this section, we consider a benevolent social planner whose goal is to optimise the surplus of the risky project. Assume n is exogenously fixed and remains constant. Given n, the socially optimal level of experiment requires all these n agents work till some point as long as the marginal expected surplus from experimenting dominates that from giving up. In other words, the experiment ceases as soon as the marginal surplus becomes negative or a success arrives, so the social planner specifies an optimum stopping time to terminate the experiment. Note that we apply the same methodology as Halac et al. (2016a) in order to solve the first-best stopping time. Define this efficient stopping time as t^* before which experimenting will always be profitable.

As all n agents are supposed to work before this efficient stopping time, the per-period aggregate success probability conditional on the good state will be identical. At any time period t the belief of the good state will be updated from previous $t-1$ times' failure, defined as $\overline{\beta}_t$. Let s_t be the expected surplus from team experimentation at time t given that all agents work, expressed as: $s_t = \overline{\beta}_t[f(h-nc)-(1-f)nc]-(1-\overline{\beta}_t)nc = \overline{\beta}_t hf - nc$. On the other hand, if none of the agents work, i.e $\forall i \leq n, x_{i,t} = 0$, $f = 0$, so $s_t = 0$. Therefore, the social planner wants to set up an efficient stopping termination time before which the marginal surplus of team experimentation is non-negative:

$$\begin{aligned} &t^* \in max_t \, \{t : t \in \mathbb{Z} \mid \overline{\beta}_t hf - nc \geq 0\} \\ &where\ \overline{\beta}_t = \frac{\beta_0(1-f)^{t-1}}{\beta_0(1-f)^{t-1}+(1-\beta_0)}. \end{aligned}$$

Solving for the efficient stopping time, we get:

$$t^* = 1 + \lfloor \frac{log(\frac{nc}{fh-nc}\frac{1-\beta_0}{\beta_0})}{log(1-f)} \rfloor, \tag{2.3}$$

where f is the per-period aggregate success probability of n agents' working conditional on the good state, and we apply the floor function in order to be consistent with the discrete-time setting. As indicated in Equation 2.3, t^* decreases in the effort cost c and increases in the revenue h. Intuitively, if it is more costly to hire an agent, the planner would save the cost by shortening the experimentation; if the revenue generated from the success is higher, the planner will extend the experimentation for few more periods.

2.3.2 Optimal contract under perfect information

As a group of agents work cooperatively on the project, they share the outcome of the project but their individual contribution cannot be separated. The breakthrough of the project should be associated with the collective effort of all the team workers. Thus, the planner has to reward everyone in the case of success. Suppose the planner could perfectly monitor agent's actions, so the agents have no chance of shirking. This section determines an optimal reward that should be large enough to motivate everyone to work till the first-best stopping time is reached.

The planner's problem is to achieve the socially efficient surplus, which is equivalent of searching for a cut-off posterior belief below which the planner wants to terminate the project. So the cut-off $\overline{\beta}_t^*$ is formally calibrated as the following:

$$\overline{\beta}_t f(n)h - nc \geq 0 \rightarrow \overline{\beta}_t^* = \frac{nc}{f(n)h}$$

Under perfect information, the actions of the team workers can be fully monitored by the principal. If the agent shirks, he will get the outside option zero payoff as there is no punishment specified in the contract[1]. Under perfect information, their belief about the good state will be updated symmetrically if all workers are rational and choose to work. They are willing to work jointly on this project as long as the expected payoff is greater than the outside option zero. Define their cutoff belief as β_t^* below which agents will shirk and choose the outside option. Therefore, each agent's problem is solved as follows:

$$\beta_t f(n)b - c \geq 0 \rightarrow \beta_t^* = \frac{c}{f(n)b}.$$

Thus, as long as $\overline{\beta}_t^* \geq \beta_t^*$ holds, agents will be fully incentivised to work until the cutoff belief is reached. Solve the inequality equation $\overline{\beta}_t^* \geq \beta_t^*$, we can determine a minimum bonus needed to motivate the agents, which is $\frac{h}{n}$. Thus, each agent gets a share of the revenue in the event of success.

Here we formally write down the optimisation programme of the social planner subject to the participation constraint of the agent:

$$\begin{aligned} \underset{b}{\text{minimise}} \quad & \overline{\beta}_t^* \geq \beta_t^*(b) \\ \text{subject to} \quad & U_0 \geq 0, \end{aligned}$$

where U_0 is the total expected payoff of the agent expressed in Equation 2.1. Solve the problem, we just let $U_0 = 0$ and choose a minimum reward.

Proposition 1 *Under perfect monitoring, with a given number of workers n, the planner would propose an equal sharing contract to implement the first-best when the contribution of individual agent cannot be fully separated such that* $C = (W_0, \frac{h}{n}, t^*)$, *where* $t^* = 1 + \lfloor \frac{log(\frac{nc}{fh-nc}\frac{1-\beta_0}{\beta_0})}{log(1-f)} \rfloor$, *and* $W_0 = -\sum_{t=1}^{t^*} \delta^t \{(1 - f(n))^{t-1}\beta_0(f(n)\frac{h}{n} - c) - (1-\beta_0)c\}$.

[1]Assume the agents can always walk away and choose their safe arm (the outside option). Thus, we cannot enforce a punishment in the contract.

The above proposition indicates that to achieve the first-best stopping time the social planner would equally distribute a share of the revenue to each agent whenever the project succeeds, which amounts to h/n. According to the participation constraint of the agent, W_0 is negative, which means there is an upfront payment from each agent to the planner. In this case, the planner can extract all the surplus, and the first-best experimentation time can be achieved.

2.3.3 Optimal hiring choice

The social planner cares not only about the efficient stopping time but also the size of the total surplus. If the planner is allowed to determine the optimal n and acts like the principal, what would be the optimal hiring choice? The answer is not transparent since there exists a trade-off when making the hiring decision. Recruiting more agents shortens experimentation time and increases the probability of success; however, it costs more.

The following determines the optimal hiring choice via the discounted expected total surplus $S(n)$ at the initial period. Let $s_t(n)$ be the expected surplus if no success arrives before t when hiring n agents, where n is constant over time. $f(n)$ is defined as the per-period success probability when all these n-agent are supposed to work the entire time, so $f(n)$ is same across time. Then we express $s_t(n)$ as follows:

$$s_t(n) = \delta^t[\beta_0(1 - f(n))^{t-1} + 1 - \beta_0](\overline{\beta}_t hf - nc),$$

which consists of two parts: the probability that the project does not succeed before period t, referring to the first term, and the marginal surplus based on the updated belief as the second term. To further interpret the surplus we substitute $\overline{\beta}_t$ according to its definition and make some rearrangement:

$$s_t(n) = \delta^t\beta_0(1 - f(n))^{t-1}(fh - nc) - \delta^t(1 - \beta_0)nc,$$

where the first term could be interpreted as the expected payoff conditional on the previous $t-1$ times' failure in the good state , and the second is the expected payoff in the bad state. At period 0, taking account of the total expected discounted surplus over the entire experiment, denoted as $S_0(n)$:

$$
\begin{aligned}
S_0(n) &= \sum_{t=1}^{t} \delta^t (1-f(n))^{t-1}\beta_0(f(n)h - nc) - \delta^t(1-\beta_0)nc \\
&= \sum_{t=1}^{t} \delta^t (1-f(n))^{t} \frac{\beta_0(f(n)h - nc)}{1-f(n)} - \delta^t(1-\beta_0)nc \\
&= \frac{\delta(1-f(n))[1-\delta^t(1-f(n))^t]}{1-\delta(1-f(n))}\frac{\beta_0(f(n)h-nc)}{1-f(n)} - \frac{\delta(1-\delta^t)}{1-\delta}(1-\beta_0)nc.
\end{aligned}
$$

The social planner could get the optimum size of surplus by maximising the total expected discounted surplus; formally define his optimisation problem as follows:

$$
\begin{aligned}
n^* = \arg\max_n \{S_0(n) = & \frac{\delta(1-f(n))[1-\delta^{t^*}(1-f(n))^{t^*}]}{1-\delta(1-f(n))}\frac{\beta_0(f(n)h-nc)}{1-f(n)} \\
& - \frac{\delta(1-\delta^{t^*})}{1-\delta}(1-\beta_0)nc\},
\end{aligned}
$$

where $t^*(n^*) = 1 + \lfloor \frac{log(\frac{nc}{f(n)h-nc}\frac{1-\beta_0}{\beta_0})}{log(1-f(n))} \rfloor$. In this case, the optimal stopping time is a function of the optimal hiring choice n^*.

Proposition 2 *The optimal stopping time decreases in the optimal number of agent hired.*

Proof. See Appendix. ■

Intuitively, if the planner hires more agents, the total success probability increases, which gives rise to a faster learning speed. Thus, the experimentation will last for a shorter period of time.

2.3.4 A simple example: one versus two

Here we present a simple example to show the changes of the surplus corresponding to employing an additional agent. Thus, we compare the total surplus generated by having one agent working for two periods versus two agents working for just one period. Assume that the stopping time is determined by the first-best stopping rule in both cases such that $t^*(1) = 2$, $t^*(2) = 1$.

Focusing on the success probability function solely, one might think that hiring n^* agents to work t^* periods is equivalent to hiring a sufficient large number of agents to finish the project within just one period as the same success probability might hold. In the case of one versus two, hiring just one agent to work two periods gives the same success probability as hiring two agents to work one period, i.e. $f(1) + [1 - f(1)]f(1) = f(2)$. However, these two choices are not equivalent in terms of the total surplus received when taking accounting of the belief updating and discounting.

Let $S_0(1)$ and $S_0(2)$ denote the total expected discounted surplus in the above two cases accordingly:

$$\begin{aligned} S_0(1) &= \delta(\beta_0 f(1)h - c) + \delta^2\{(1 - f(1))\beta_0(f(1)h - c) - (1 - \beta_0)c\}, \\ S_0(2) &= \delta(\beta_0 f(2)h - 2c). \end{aligned}$$

Then we take difference of $S_0(2)$ and $S_0(1)$:

$$\begin{aligned} S_0(2) - S_0(1) = \delta[\beta_0(f(2) - f(1))h - c] - \delta^2\{[1 - f(1)]\beta_0(f(1)h - c) \\ - (1 - \beta_0)c\}, \end{aligned}$$

where the first term denotes the discounted marginal benefit of hiring one more agent, and the second denotes the discounted marginal benefit of extending the experiment for one more period. To make the hiring decision, principal basically do the revenue-cost-comparison. If the increased expected revenue dominates the cost increase, the principal will hire two agents. Otherwise, hiring one agent is preferable.

Particularly, when $\delta = 1$, $S_0(2) - S_0(1) = -\beta_0 f(1)c < 0$, having one agent is preferred. An extremely patient principal does not mind delaying the experimentation and prefers hiring just one agent in order to save some effort cost. Since $\beta_0 f(1)$ is the chance of success with one agent in the first period, so $-\beta_0 f(1)c$ could be interpreted as the amount of cost saved if the project succeeds in the first period with just one agent. However, when δ approaches closely to 0, the principal is extremely impatient, the difference between two choices gets close to 1. Intuitively, when the future payoff values almost nothing to the planner, he wants the success as soon as possible, thus the principal may employ an extra agent to end the experiment within just one period. We can conclude in the following result from this simple exercise.

Result 1 *The social planner's hiring decision can largely be altered by his time preference. For an extremely patient planner, he tends to hire as fewer agents as possible, presumably, just one if he is patient enough and delay the experimentation. On the other hand, an extremely impatient principal would in general hire more agents in order to accelerate the experimentation and achieve the success as soon as possible.*

2.3.5 Simulation

Figure 2.2 shows the evaluation result: $n^* = 60$ and $t^* = 1$ when $\beta_0 = 0.5$ and $\delta = 0.05$[2] . Given $\beta_0 = 0.5$ the principal believes that this risky project has the prior belief of 0.5 to be in the good state, so it's worth experimenting[3]. After the first-period failure, the principal would become very pessimistic about the state of the world. It is no longer worth experimenting given updated posterior belief. The discount factor is set to be 0.05 such that the principal is very impatient and prefers to complete the project as quickly as possible.

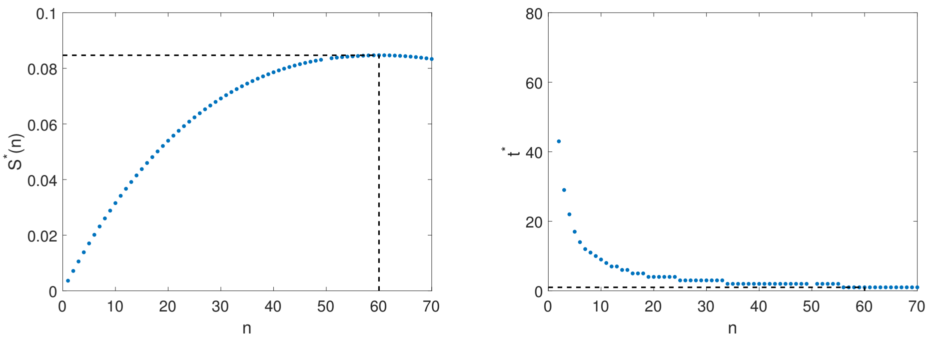

Figure 2.2: Optimal hiring choice of an impatient principal ($\delta = 0.05$)

Figure 2.3 evaluates the case when the principal holds a sufficient large discount factor, set $\delta = 0.95$, then $n^* = 17$ and $t^* = 5$, controlled for other parameter values. As the principal values future almost as much as present, he would hire fewer agents to experiment for more periods.

[2]Note that the optimal results of n and t are the local optimum, whereas, the global optimum can be obtained by solving the total surplus function, which is left for future research.

[3]Parameters values are chosen to insure that the principal has the incentive to start the project, which means $\beta_0 f(n)h - nc > 0$ holds. Thus, when $n \in \mathbb{Z}^+$, $\beta_0 = 0.5$, $\delta = 0.05$, $c = 0.03$, $\lambda = 0.02$, $h = 10$.

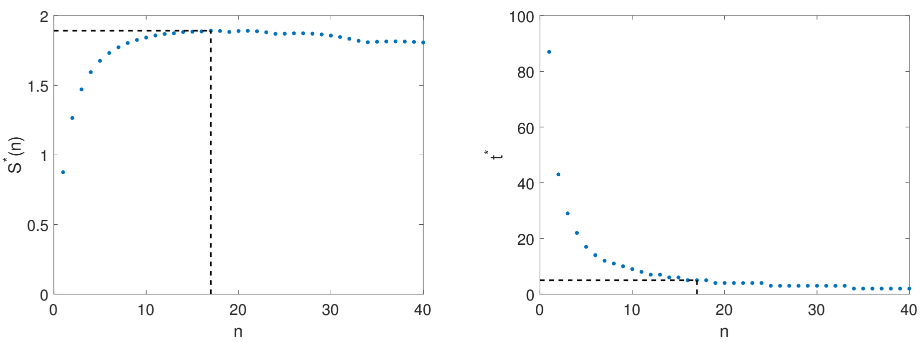

Figure 2.3: Optimal hiring choice of a patient principal ($\delta = 0.95$)

2.4 Free-rider problem

This sections considers an optimal contracting under imperfect information about the actions of the agents. Assume the principal only observes the outcome of the project each period rather than the actions of each agent. Additionally, the principal could not distinguish each agent's contributions to the success. Although the agents cannot observe the others' actions, the progress of the project, whether it succeeds or not, is known to them. Thus, once one of them makes a breakthrough, others could just be a free-rider by claiming the reward. If no incentive-compatible contract is imposed, all the agents will be reluctant to exert effort, which impairs the welfare of the principal. With the number of agents n be exogenously fixed, the principal aims to propose an optimal contract to motivate all the agents to work until the socially efficient stopping time t^*. According to Holmstrom (1982), the principal could prevent the free-rider problem by "breaking" the budget: team members will be provided with a certain revenue share only if the first-best level of total production is achieved. Otherwise, the principal could even deprive the entire output from team workers as a serious punishment. Such a non-balancing budget strategy is theoretically proven to eliminate the free-rider problem and could be enforced by various types of contracts.

First, we apply the bonus contract, $C = (W_0, t^*, b)$, where the principal offers b as a reward to every team member once the project succeeds, t^* is the socially efficient stopping time, W_0 is an upfront transfer between the agent and principal, which ensures the participation constraint is satisfied and binding, such that $W_0 = -\sum_{t=1}^{t^*} \delta^t \{(1 - f(n))^{t-1}\beta_0(f(n)b - c) - (1 - \beta_0)c\}$, which is negative. To join the teamwork, each agent needs to pay the principal W_0 at the beginning, and in the event of success, each agent can get compensated by the bonus b.

Second, we need to find out an optimal b^*, which removes the free-rider problem and provides each agent incentive to work until the social efficient termination time t^*. Each agent will keep on exerting effort so long as the expected payoff from working dominates that from shirking. Suppose for each agent, $\bar{t}$ is the optimal stopping time which solves agent's maximisation problem as a best response to others' working behaviour. Under the bonus contract above, an agent's maximum stopping time $\bar{t}$ is characterised as the following:

$$\bar{t} = max_t \{t : t \in \mathbb{Z} \mid \overline{\beta}_t bf(n) - c \geq \overline{\beta}_t bf(n-1)\}$$
$$where\ \overline{\beta}_t = \frac{\beta_0(1-f(n))^{t-1}}{\beta_0(1-f(n))^{t-1} + (1-\beta_0)},$$

where the left hand side of inequality equation denotes the expected payoff from working given the other agents are working this period, and the right hand side denotes the expected payoff from free-riding at period t holding updated belief from previous $t-1$ time's failure. Solving for the optimal stopping time, we get:

$$\bar{t} = 1 + \lfloor \frac{log(\frac{c}{(f(n)-f(n-1))b-c}\frac{1-\beta_0}{\beta_0})}{log(1-f(n))} \rfloor.$$

We then determine an optimal bonus by solving the inequality equation: $\bar{t} \geq t^*$, where t^* is the socially efficient stopping time. The result is as follows:

Proposition 3 *When $b \geq \gamma h$, each agent will work until the socially efficient stopping time t^*, such that $\bar{t} \geq t^*$ always holds for any $n > 1$, where $\gamma = \frac{f(n)/n}{f(n)-f(n-1)} > 1$. Thus, the minimum reward is $b^* = \gamma h$.*

Proof. See Appendix. ■

The above proposition indicates that the smallest reward needed to achieve the socially efficient stopping time is γh under the imperfect information. In other words, γ is the minimum share of the revenue needed to motivate effort until t^*. As the parameter $\gamma > 1$, it indicates that the average success probability always exceeds the marginal success probability in n-agent teamwork for $n > 1$, which is followed from the concavity of the function f. In this case, the principal needs to pay more than the total revenue to every one as long as the project succeeds.

However, the contract is signed in the sense that there is an upfront transfer from the agents to the principal, which is equivalent to sell the project to each agent. The principal in this case acts as an outsider who distributes a relative high bonus to each agent whenever someone succeeds. As both parties are supposed to be fully committed to this long-term contract, agents are willing to sign the contract and work for the principal as long as the total expected payoff is non-negative.

Corollary 1 *Under single-agent-experimentation, the principal should reward the agent h if the project succeeds when the agent's action is unobservable.*

Proof. When $n = 1$, $\gamma = 1$. For $b \geq h$, following the proof of Proposition 3, we could induce: $\bar{t} \geq t^*$. Moreover the equality holds whenever $b = h$, which is consistent with the result in the single-agent-experimentation model by Halac et al. (2016a). ■

2.4.1 Equilibrium analysis

Given the optimal contract $C = (W_0, t^*, \gamma h)$, this section discusses the equilibrium behaviour of the agents if the others' actions are unobservable. In particular, we want to answer whether for any agent regardless of the actions of the others, working each period dominates shirking. If so, working can be proven as an unique equilibrium.

Proposition 4 *Under the contract $C = (W_0, t^*, \gamma h)$, there exists an unique equilibrium where each agent chooses to work until the first-best stopping time t^* regardless of the actions of the others as long as no success arrives.*

Proof. See Appendix ■

The above proposition shows that the agents will always work under the optimal contract no matter what the actions of the others are. On the one hand, we show that if all the others work, the best response of the agent is to work. Intuitively, within the first-best stopping time t^*, if all the others work, the agent will always find work to be the best response as working increases the total success probability and results in a higher expected payoff. On the other hand, in an extreme case, if all the others shirk, the best response for this agent is to work. Given the bonus contract, the agent is willing to work alone until the stopping time, denoted as $\tilde{t}$, which is longer than the teamwork-stopping-time t^*. Thus, within time t^*, the agent works as long as no success arrives. Moreover, if some of the agents shirk occasionally, the agent still prefers working. Intuitively, if someone shirked in the past, the agent's posterior belief of the good state would increase, which gives rise to a higher expected payoff of working this period. If someone shirks this period, the agent would have more incentive to work since the marginal contribution to the total success probability is greater when fewer agents are involved since f is concave in n. Thus, working is proven to be a dominant strategy of the agents given all those three scenarios.

2.5 Extension

In Section 2.3.3, we endogenise the planner's hiring choice n and let the planner maximises the total expected discounted surplus by optimally choosing the number of workers to hire and the stopping time simultaneously at the beginning. As an extension, this section examines the optimal hiring choices of a myopic planner when n is allowed to vary each period. The myopic planner would make the hiring choice each period to optimise the expected payoff of the current period, instead of taking accounting of the expected payoff in the future. Thus, the planner would not commit to a long-term stopping time. As long as the expected payoff in the current period is non-negative, the planner will continue hiring and making short-term contracts with agents according to the differentiate optimum hiring choices each period.

We write down the optimisation of the programme of the planner as follows:

$$\begin{aligned} \underset{n_t}{\text{maximise}} \quad & s_t(n_t) = \beta_t f(n_t)h - n_t c \\ \text{subject to} \quad & s_t(n_t) \geq 0, \end{aligned}$$

where β_t is the updated posterior belief based on the hiring history and previous outcomes of the project, and the constraint is to guarantee a non-negative expected payoff of the planner. The first order condition indicates that:

$$f'(n_t) = \frac{c}{\beta_t h} \Rightarrow n_t^* = \lfloor \frac{1}{\lambda} \ln \frac{\lambda h \beta_t}{c} \rfloor,$$

where n_t^* monotonically increases in β_t and h but decreases over time. The planner hires more agents if the posterior belief of the good state is greater or the revenue is higher. As long as as no success comes, the posterior belief decreases over time under the "good" news model. Thus, the optimal hiring choice of the myopic planner monotonically decreases over time in accordance with the downward adjusted posterior belief.

2.5.1 Simulation

Due to the decreasing feature of the optimal hiring choice, in the last period we would expect the social planner to hire at least one agent, so assume that $n_T^* = 1$ when T is the last period of experimentation. Under this assumption, we run several simulations to check the above optimal hiring choice, and several features can be explored here.

First, the number of agents hired decreases over time. During the process of innovation, the principal values the chance of breakthrough but also the cost of hiring. The decreasing number of agent employed coincides with the growing pessimistic belief about the good state. More agents would be hired initially if the revenue is higher. There would be a sharp decrease of the optimal n^* especially for higher h after the first failure, which is mainly driven by the intensively downward adjusted belief. As in Figure 2.4, for $h = 80$, n^* dropped by 50 if it fails at the first period. In comparison, for $h = 20$, n^* only dropped by 20 after the first failure. Thus, n^* is flatter when the project revenue is smaller.

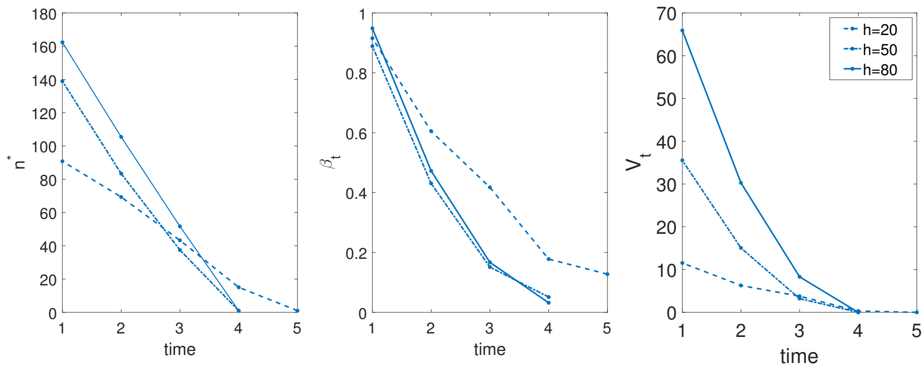

Figure 2.4: Optimal hiring choice for different lump-sum h controlled for other parameters, i.e. $(\lambda, c, \delta, \beta_0) = (0.02, 0.05, 0.5, 0.8)$.

Secondly, the principal's time preference impacts on the optimal hiring choice. With a higher discount factor, the would hire fewer agents and runs the experiment longer, which is similar as the fixed-n case. In the far right of Figure 2.5, when

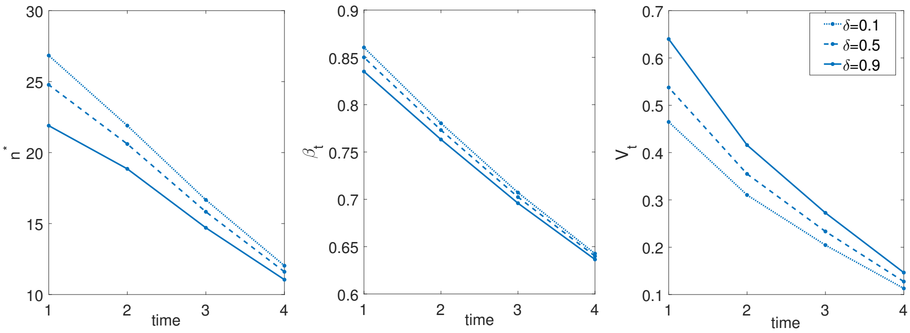

Figure 2.5: Optimal hiring choice for different time preferences δ controlled for other parameters, i.e. $(\lambda, c, h, \beta_0) = (0.02, 0.05, 5, 0.8)$.

$\delta = 0.9$, the principal values future as much as the present, the total discounted present value of the project (V_t) is the highest compared with the smaller discount factors. The monotonic decreasing tendency of the optimal hiring choice still holds since it is driven by the growing pessimistic posterior belief about the good state.

2.6 Conclusion

In conclusion, this paper addresses the contracting issue under the innovative project holder and a team of agents where endogenous learning is required to explore the nature of the state. To motivate the agents, the principal would require an upfront transfer by the agent and give a large bonus in the event of success. As long as the agents find shirking unprofitable, working can be proven as a dominant strategy. The principal would be able to implement the first-best stopping time under this contract. We also show that various decisions should be made contingent on the posterior updated belief about the good state. In general, fewer agents will be employed as long as the posterior belief decreases over time. Moreover, the time preference of the principal plays an important role in making the hiring choice: for a relatively patient principal, less agents will be hired initially and the experiment will last longer. In all, this paper proposes an optimal contract for team experimentation under imperfect information and offers recommendations on the optimal hiring choice in innovation. Team experimentation with heterogeneous agents and the limited liability constraint are left for future research.

Appendix

2.A Proof for Proposition 2

Lemma 2 $\frac{f(x)}{x} = \frac{1-e^{-\lambda x}}{x}$ *is monotonically decreasing in* x, $\forall x \in \mathbb{R}^+$.

Proof. Taking derivative respect to x:

$$\begin{aligned}\frac{d\frac{1-e^{-\lambda x}}{x}}{dx} &= \frac{e^{-\lambda x}\lambda x - (1-e^{-\lambda x})}{x^2} \\ &= -\frac{e^{-\lambda x}(e^{\lambda x} - \lambda x - 1)}{x^2}\end{aligned}$$

According to the definition of the exponential function:

$$\begin{aligned}e^{\lambda x} = \sum_{k=0}^{\infty}\frac{(\lambda x)^k}{k!} &= 1 + \lambda x + \frac{(\lambda x)^2}{2} + \frac{(\lambda x)^3}{6} + \frac{(\lambda x)^4}{24} + \cdots \\ &> 1 + \lambda x\end{aligned}$$

Thus, $\frac{d\frac{1-e^{-\lambda x}}{x}}{dx} < 0$ holds. ■

As $t^* = 1 + \lfloor \frac{log(\frac{nc}{f(n)h-nc}\frac{1-\beta_0}{\beta_0})}{log(1-f(n))} \rfloor$, substituting $f(n) = 1 - e^{-\lambda n}$ into the denominator and making some rearrangement:

$$t^* = 1 + \lfloor \frac{log(\frac{f(n)h}{nc} - 1)\frac{\beta_0}{1-\beta_0}}{\lambda n} \rfloor.$$

For $n \in \mathbb{Z}^+$, according to the lemma above, $\frac{f(n)}{n}$ decreases in n. As the denominator of t^* increases in n, t^* decreases in n can be proven.

2.B Proof for Proposition 3

We solve the inequality relationship: $\bar{t} \geq t^*$ by substituting the expression of $\bar{t}$ and t^* such that:

$$
\begin{aligned}
1 + \lfloor \frac{log(\frac{c}{(f(n)-f(n-1))b-c}\frac{1-\beta_0}{\beta_0})}{log(1-f(n))} \rfloor &\geq 1 + \lfloor \frac{log(\frac{nc}{f(n)h-nc}\frac{1-\beta_0}{\beta_0})}{log(1-f(n))} \rfloor \\
\frac{log(\frac{c}{[f(n)-f(n-1)]b-c}\frac{1-\beta_0}{\beta_0})}{log(1-f(n))} &\geq \frac{log(\frac{nc}{f(n)h-nc}\frac{1-\beta_0}{\beta_0})}{log(1-f(n))} \\
log(\frac{c}{[f(n)-f(n-1)]b-c}\frac{1-\beta_0}{\beta_0}) &\leq log(\frac{nc}{f(n)h-nc}\frac{1-\beta_0}{\beta_0}) \\
\frac{c}{(f(n)-f(n-1))b-c}\frac{1-\beta_0}{\beta_0} &\leq \frac{nc}{f(n)h-nc} \\
[f(n)-f(n-1)]b-c &\geq \frac{f(n)}{n}h-c \\
[f(n)-f(n-1)]b &\geq \frac{f(n)}{n}h \\
b &\geq \frac{f(n)/n}{f(n)-f(n-1)}h.
\end{aligned}
$$

Since $f(n) = 1 - e^{-\lambda n}$, $f'(n) = 1 + \lambda e^{-\lambda n} > 0$ and $f''(n) = -\lambda^2 e^{-\lambda n} < 0$, $\forall n \in \mathbb{Z}^+$. Thus, $f(n)$ is concave in n. Apply the concavity property we get the following:

$$
\frac{f(n)-f(n-1)}{n-(n-1)} < \frac{f(n)-f(0)}{n-0},
$$

where $f(0) = 0$. Thus, for any $n > 1$, $\frac{f(n)}{n} > f(n) - f(n-1)$ holds as $f(n)$ is concave in n. Let $\gamma = \frac{f(n)/n}{f(n)-f(n-1)}$, so $\gamma > 1$ holds.

2.C Proof for Proposition 4

To determine the existence of an equilibrium, we discuss the best response of agent i regardless of the others' actions. At any period $t \leq t^*$, if all the others choose to work, the agent i would compare the expected payoff this period from working and shirking as follows:

$$
\begin{aligned}
u_{i,t}(x_{i,t} = 1 | x_{-i,t} = 1) &= \overline{\beta}_t f(n)\gamma h - c \\
u_{i,t}(x_{i,t} = 0 | x_{-i,t} = 1) &= \overline{\beta}_t f(n-1)\gamma h.
\end{aligned}
$$

According to the optimal stopping rule, the agent i will work until t^* as long as no success arrives, which means $u_{i,t}(x_{i,t} = 1|x_{-i,t} = 1) \geq u_{i,t}(x_{i,t} = 0|x_{-i,t} = 1)$, $\forall t \leq t^*$. Thus, the agent i would work if the others work.

At any period $t \leq t^*$, if the others shirk this period as well in the past, the agent also makes the choice between working and shirking according to the expected payoff such that:

$$u_{i,t}(x_{i,t} = 1|x_{-i,t} = 0) = \tilde{\beta}_t f(1)\gamma h - c,$$
$$u_{i,t}(x_{i,t} = 0|x_{-i,t} = 0) = 0,$$

where $\tilde{\beta}_t$ is the belief updating when he works but no success comes, expressed as $\tilde{\beta}_t = \frac{\beta_0(1-f(1))^{t-1}}{\beta_0(1-f(1))^{t-1}+(1-\beta_0)}$. Under the single-agent-experimentation, the optimal stopping time for agent i can be solved as follows:

$$\tilde{t} = max_t \{t : t \in \mathbb{Z} \mid \tilde{\beta}_t \gamma h f(1) - c \geq 0\}$$
$$where \ \tilde{\beta}_t = \frac{\beta_0(1 - f(1))^{t-1}}{\beta_0(1 - f(1))^{t-1} + (1 - \beta_0)},$$

so $\tilde{t} = 1 + \lfloor \frac{log(\frac{c}{f(1)\gamma h - c}\frac{1-\beta_0}{\beta_0})}{log(1-f(1))} \rfloor$, and the agent i will work until $\tilde{t}$, which satisfies $\tilde{t} > \bar{t} = t^*$. Thus, $u_{i,t}(x_{i,t} = 1|x_{-i,t} = 0) \geq u_{i,t}(x_{i,t} = 0|x_{-i,t} = 0)$ holds as long as $t \leq \tilde{t}$, which proves that working is the best response to the others' shirking behaviour.

If some of agents shirk this period as well as in the past, we want to check whether the equilibrium still holds. Let m denotes the number of agents choose to work at period t such that $m < n$. Due to the concavity of function f, $f(m) - f(m-1) > f(n) - f(n-1)$ holds as $m < n$. This inequality shows that the incentive of working increases if someone shirks this period since the marginal success probability is higher in m. Moreover, if in the past, someone shirked and was speculated by the agent, the agent would increase the posterior belief of the good state this period, which gives rise to a higher expected payoff as well as higher incentive to work. Thus, the agent will always choose work no matter someone has shirked this period or in the past.

Chapter 3

Financing innovation under asymmetric information: signalling through internal financing

3.1 Introduction

Under-investment in innovation has become a widely recognised problem, especially for those small and medium-sized innovative firms. Early theoretical studies have suggested that severe asymmetric information between innovators and investors may result in under-provision of innovation investment according to the survey by Hall and Lerner (2010). A moral hazard problem arises when the results of innovation can easily be copied by the free-riders, which makes the innovators reluctant to undertake the investment and reveal the innovation progress (Nelson, 1959; Arrow, 1962). On the other hand, investors may not observe the quality of the innovation project to the same extent as innovators do, which gives rise to the adverse selection issue (Brealey et al., 1977; Myers and Majluf, 1984). Compared with external financing sources (venture capitalists or business angels), internal funding is a more convenient and effortless way of financing, which particularly works well at the initial learning stage of innovation. In fact, internal financing plays a central role in R&D spending among small high-tech firms in the US and UK, and the development of innovation is primarily based on the availability of internal financing (Spence, 1979; Himmelberg and Petersen, 1994; Bougheas, 2004). Although the importance of internal financing has been corroborated empirically, there is hardly any theoretical study exploring the crucial role of internal financing in innovation, especially under asymmetric information. This chapter aims to provide insights into firms' optimal decision of internal and external financing. It also shows that internal financing could be adopted as a signalling tool in innovation financing that alleviates the adverse selection problem.

The chapter first discusses the existence of the adverse selection problem within the innovation process. In practice, innovation often goes through a trial and error learning phase before major investment occurs, where innovators usually use the internal funding to learn how good the project is. Assume innovators are initially endowed with an identical innovative project with the same payoff in the event of success or failure but have either a high or low prior belief about success[1]. The prior beliefs are unobservable to the investors. However, for the high-type project, it is normally more attractive to investors and get funded easier than the low-type. In this case, an adverse selection problem arises where the low-type innovators may pretend to be a high-type. Without any signalling scheme, it may even result in the "lemons" market problem where no investor is willing to fund the project, and the high-type innovators are driven out of the market (Akerlof, 1970). To address this issue, we claim that not only does the internal funding support the effort cost but also reveals the innovator's prior belief to the investors.

To see this, we first look into the first-best internal financing of the high and low types. The model adopts the "bad news" exponential-bandit game by Keller and Rady (2010) in accordance with the unknown state of the innovation and the exogenously fixed experimentation time. Internal funding is allocated by the innovators in order to explore the unknown state of the world. In the good state, a risky project succeeds and generates a higher revenue once the investment is made; in the bad state, it fails and receives a lower revenue. Experimentation would randomly generate bad signals which indicate the state is bad. Thus, one can be certain that the state is bad as long as a bad signal is observed, otherwise, the state remains unknown before the investment. Assume that the arrival of bad news has an intensity λ, which depends on the amount of internal funding

[1] In this chapter, we sometimes refer to innovators as the high-type or low-type according to their endowed beliefs about the project, but there is no heterogeneity in innovators' learning ability or marginal cost of learning.

provided, such that more internal funding accelerates learning and provides a more accurate prediction of the state. According to Bayes rule, given a fixed time of experimentation, the posterior belief of the good state increases in the amount of internal funding as long as no bad news arrives. Under perfect information, the first-best internal financing is determined by the profit-maximisation problem of the innovators. As a result, for the project with a low prior belief of the good state, a larger amount of internal funding is required. Conversely, given a high prior belief, the innovators use less internal funding. Intuitively, the project with a higher prior of success requires less learning input as it is expected to succeed more easily, so the innovators are prone to save their funding for other projects; however, the innovation with a lower prior requires more experimentation and attention in order to avoid future failure. The above intuition appears to be in line with the general result of a real-option signalling model where innovators hold the option of investment timing, such that innovators with a higher prior belief are apt to provide less learning and invest earlier (Grenadier and Wang, 2005; Bouvard, 2014; Bobtcheff and Levy, 2014). Instead, internal financing is treated as the optimal decision of innovators in this chapter, which conveys private information about the quality of the innovative project. Despite a larger internal funding being allocated by the low-prior project holders, by the time of investment, they still hold a lower expectation on the state being good than those with a higher prior. Indeed, for those cash-constrained innovative companies, it is fairly risky to spend too much internal funding on the project with a relatively low prior of success at the early stage.

The chapter then explores external financial contracting between the informed innovators and the less-informed investors. I assume that the innovative project, irrespective of the prior beliefs, requires the same amount of external investment, which cannot be financed internally. The innovators are expected to share a certain proportion of revenue with the investor if the project succeeds in return for the external investment. Hence, revenue shares conditioning on the outcome of the innovation will be included ex-ante in the financial contract. Specifically, the innovators with a higher prior belief will keep a higher revenue share given a higher expected possibility of success compared with the low-type. Thus, the equilibrium contract is composed of the revenue shares and internal funding, which jointly convey the information about the quality of the project. We can find the relevant contractual terms appear in the real-world financial contracts within venture capital and start-up firms. As documented in Kaplan and Strömberg (2003), cash flow rights contingent on the accomplishment of milestones are widely used in incentivising entrepreneurs under asymmetric information, which function the same as revenue shares in our model. Internal funding is normally implicitly regarded as a criterion by which investors may judge how much confidence innovators have in innovation.

The information asymmetry may generate distortions in the sense that the first-best contracts are no longer attainable for both types. With the first-best contracts, the low-type innovators will deviate by choosing the high-type contract by offering less of the revenue shares and internal financing to the investor. To resolve this problem, the high-type would endeavour to separate from the low ones by signalling. Proposition 5 states that a larger amount of internal financing will be adopted by the high-type in this case as compared with the high-type full-information optimum as well as the low-type's first-best. In fact, increasing internal funding would adjust upwards the project's expected probability of success, if no bad news is observed during experimentation, which in return builds up the investor's expectation of investing in this innovative project. Thus, the expansion of the high-type internal financing compensates the investor's potential loss on the low-type contract, which avoids the market-breakdown. More importantly, internal financing functions as a separation tool for the high-type which deters the low-type from mimicking. Regarding high-type and low-type innovators' preferences over internal funding and revenue shares, a single-crossing property is shown to be held. It indicates that it costs low-type less to exchange internal funding for an additional increase in revenue shares in comparison with a high-type. Details of this property will be shown in the analysis later. Consequently, mimicking the high-type contract by exerting more internal funding is no longer profitable for the low-type as the low-type incentive compatibility constraint is relaxed. Proposition 5 formally describes this separating equilibrium which manages to distinguish the agent's type and make the investor break-even on average. This chapter also covers other pooling equilibria, which involve cross-subsidisation from the high to the low, when either one of the contractual tools is available. The above-mentioned separating equilibrium is proven to be the least-costly separating equilibrium according to the criteria proposed by Maskin and Tirole (1992) as it is interim efficient and weakly Pareto dominates all other equilibria in the sense that both types of agents would perform their type-equivalent contract and the low-type is weakly better-off.

Supportive evidence regarding an innovator's distinctive choices over internal funding could be found in many empirical studies of various industries where R&D plays a pivotal role. According a study of the Italian manufacturing industry by Succurro and Costanzo (2016), internal financing plays an important role in R&D financing such that there is a strong positive correlation between the size of the internal financing and whether the firm will take any R&D and the size of overall spending on R&D. Moreover, the high-tech firms in Italy still rely largely on the internal financing sources rather than external financing due to the information asymmetry problem. Danzon et al. (2005) and Lerner et al. (2003) also confirm the existence of a strong reliance on the internal financing in the large and experienced pharmaceutical companies. Those companies often have strong confidence in the prospect of their research, hence more internal input is allocated at the early stage. According to data, drugs of those big companies get approved more often by the administration. Internal financing in this chapter not only represents the monetary input, but also indicates the intangible effort, managerial skills and marketing strategies of the innovator. Galende and de la Fuente (2003) conduct an empirical study on the factors that determine firms' innovation process based on the data of Spanish innovation firms. They investigate various internal factors, including tangible capital and intangible managerial skills, human capital and strategies, and their relationship with the innovative process. Their regression result indicates that the factors regarding innovator's marketing strategies and intellectual skills can explain 50% of a firm's innovative process.

Financing innovation under adverse selection has been examined in several recent papers that are closely related to ours (Grenadier and Wang, 2005; Bouvard, 2014; Bobtcheff and Levy, 2014). In their models, the duration of experimenting (investment timing), rather than internal financing, is an active contractual variable and served as a signalling tool. ‘Although this chapter shares some similarities with Bouvard (2014) where the high-type tend to provide excess effort in terms of either a larger internal financing or a longer experimentation time than the first-best to get themselves separated under imperfect information, several important distinctions should be pointed out here: firstly, in their paper, the innovators hold the option to delay or hurry the experimentation prior to investment, however, the experimentation time monotonically decreases innovators expected payoff but non-monotonically affects the expected payoffs of the investors, because there is a trade-off between learning and discounting. In contrast, in our model, the internal financing monotonically increases in the expected payoffs of the investor but monotonically decreases in the payoffs of the innovator, thus, there exists no such trade-off. Also, we assume that the internal funding is spent upfront as a lump-sum, which represents not only the cost of learning but also non-financial observed skills or effort. On the other hand, in their paper, the cost grows proportionally as experimentation is extended, hence both types of agents will stop at the same cutoff belief level where marginal benefit reaches zero. Bobtcheff and Levy (2014) model innovators with different learning speed. Because the optimal investment date is non-monotonic in learning speed λ, it results in either under or over investment under imperfect information. Although the papers mentioned above have pointed out that internal funding may alleviate the information distortion due to adverse selection, none of them have regarded internal financing as a signalling tool nor related it with learning. Thus, this chapter fills the gap in the innovation-financing literature by introducing internal financing as a contractible and signalling tool of heterogeneous innovators. Furthermore, it attempts to answer how information asymmetry affects innovators’ internal financing decision.

This article belongs to a broader literature that combines experimentation and agency problems. The seminal papers by Bergemann and Hege (1998, 2005) have examined the agency conflict under imperfect information and its correlation with the speed of investment based on the “good news” bandit. Under the arm’s-length financing when agents’ actions are observable, the investment rate decreases since the posterior belief of the good state decreases as long as no good news arrives over time. However, when actions are unobservable under the relationship financing mode, a moral hazard problem arises. In this case, the investor tends to downgrade his posterior belief of the good state after a deviation is detected, which leads to a shorter financing horizon. In a two-period experimentation model by Drugov and Macchiavello (2014), investment cost is revealed after the first period's experimentation. More recently, the mechanism designed to alleviate moral hazard in financing innovation has also received attention (Manso, 2011; Yu et al., 2012; Sannikov, 2014; Horner and Samuelson, 2013). However, this chapter is more related to another strand of the innovation literature concentrating on the adverse selection conflict between informed innovators and less-informed investors: Gomes et al. (2013) examine screening contracts which help the monopolistic investor to distinguish private information of the innovators. Free-riding and communication between a group of innovators have been studied under an innovation competition set-up (Halac et al., 2016a; Heidhues et al., 2015; Akcigit and Liu, 2015). Compared with the above-mentioned papers, the distinctive feature of this chapter is inverting the roles of innovator and investor. As we assume there are a lot of investors in the market, the innovator, the owner of the innovative project, captures the entire bargaining power and proposes the financing contracts as a take-it-or-leave-it offer. Thus, we draw attention to the equilibrium behaviour of innovators with heterogeneous prior beliefs about the future success probabilities.

As motivation and related literature are illustrated above, this chapter proceeds as follows. Section 3.2 introduces model set-up and the first-best internal financing. Section 3.3 explores the equilibrium contracts under perfect and imperfect information respectively, including the analysis of the least-costly separating equilibrium and pooling equilibria. Section 3.4 concludes.

3.2 Model

3.2.1 Set-up

Consider a two-stage innovation investment project as illustrated in Figure 3.1. At the start-up testing stage, innovators use internal funding to learn the quality of the project. As the funding is allocated prior to learning, it is fixed and irreversible even if a bad signal arrives during the experiment process afterwards, which indicates the quality of the project is bad. Let λ denote the internal

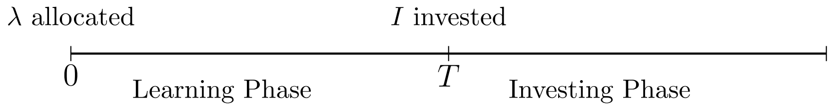

Figure 3.1: Timing of events

financing, which represents the entire input of learning. It involves not only financial cost but also non-financial inputs, such as managerial skill, observable effort, etc. Thus, it is assumed that internal funding cannot cover the investment cost in the future, which amounts to I and $I > \lambda$. At the investing stage, the investor lends the funding I to the innovator according to the pre-negotiated contract. The details of the contract will be discussed below. The result of the investment, either success or failure, is realised immediately upon investing. In the good state, the project succeeds and generates revenue R. However, it fails in the bad state and pays 0. Whenever the expected revenue of the project is greater than the investment cost together with internal funding, it's worthwhile to initiate the experiment. Let T be the exogenously fixed investment timing. Although this assumption is absent in most of the real option theoretical literature, it is quite common in real life. For instance, a machine for innovation may be only available for a certain period of time, which determines timing of the investment.

Innovators are privately assigned with a prior belief p_0^θ of the project being good, which could either be high or low, such that $p_0^h > p_0^l$. Note that the prior belief of the project is the only information privately known to innovators, but separating contracts designed in this chapter aim to help investors distinguish the high-type and low-type. Within the distribution, it is commonly known that there is a proportion q of high-type innovators and $(1 - q)$ of low-type ones. We also impose a condition to make both types of innovators have incentives to experiment: assume the expected revenue of the project based on the prior beliefs is higher than total investment, i.e., $p_0^h R > p_0^l R \geq I$.

Uncertainties exist in the innovation in the sense that whether the project will succeed or not is unknown even at the investment point, but learning will provide additional information about the project via the arrival of bad news. The intensity of the arrival is represented by the internal funding λ. Here we adopt the exponential distribution to model to the arrival of a bad signal during experimentation, following Keller and Rady (2010). Thus, learning would randomly generate bad news with the intensity λ if the state is bad following an exponential distribution, which is public observable. As soon as a bad news arrives, the posterior belief of the good state will jump to 0, and both the innovator and investor will abandon the project immediately. However, if no bad news comes, based on the Bayes rule, one would adjust upward their beliefs on the good state.

We use U and V to describe the expected payoff of the innovator and investor respectively. Both parties are risk-neutral and discount future with the discount rate r.

3.2.2 First-best internal financing

This section considers the optimal allocation of internal funding by a social planner who has perfect information on the project's prior belief. At time 0, he chooses internal funding λ^θ given the realisation of type-θ prior belief in order to maximise the expected surplus of the project, denoted by $\Pi(\lambda^\theta)$:

$$\max_{\lambda^\theta} \Pi(\lambda^\theta) = e^{-rT}[p_0^\theta + (1-p_0^\theta)e^{-\lambda^\theta T}](p_T^\theta R - I) - \lambda^\theta. \tag{3.1}$$

In continuous time setting, e^{-rT} is the time discounting and T is exogenously given investment time. The term in the brackets, $[p_0^\theta + (1-p_0^\theta)e^{-\lambda^\theta T}]$, is the probability of getting no bad news before T: in the good state, no bad news will come with certainty; in the bad state, bad news comes with probability $(1-e^{-\lambda^\theta T})$, hence $(1-p_0^\theta)e^{-\lambda^\theta T}$ is the conditional probability of not receiving any bad news in the bad state. The third term, $(p_T^\theta R - I)$, indicates the expected net return from investing at time T, where p_T^θ is the updated posterior belief on good state at time T: $p_T^\theta = \frac{p_0^\theta}{p_0^\theta + (1-p_0^\theta)e^{-\lambda^\theta T}}$ following Bayes rule if no bad news arrives; otherwise, $p_T^\theta = 0$. The last term λ^θ represents the entire experimentation cost of the θ-type project, which is assumed to be linear in λ with the coefficient 1 without loss of generality.

However, there is a trade-off of choosing internal funding λ^θ. On the one hand, a larger amount of internal funding would accelerate one's learning speed, in the event of no bad news before T, one would be more confident in the good state as a higher posterior belief will be perceived. On the other hand, internal investment is allocated before learning and irreversible. Thus, internal funding allocation varies across different types of projects.

Equation 3.1 could be rearranged by substituting the expression of the posterior belief p_t^θ follows:

$$\max_{\lambda^\theta} \Pi(\lambda^\theta) = e^{-rT}\{p_0^\theta R - [p_0^\theta + (1-p_0^\theta)e^{-\lambda^\theta T}]I\} - \lambda^\theta. \tag{3.2}$$

Then the first-best $\bar{\lambda}^\theta$ for θ-type project could be solved from the FOC:

$$\bar{\lambda}^\theta = -r + \frac{1}{T}\ln\left(T(1-p_0^\theta)I.\right) \tag{3.3}$$

Specifically, for $p_0^h > p_0^l$, $\bar{\lambda}^h < \bar{\lambda}^l$. Intuitively, for a high-type, less internal funding is needed to learn the state of nature given the high prior belief of success; however, for a low-type, larger initial funding is required in order to discover whether the project is worth investing.

3.3 The External Financial Contracting

This section focuses on the financial contracting issue between innovators and investors in the perfect competitive capital market. Assume innovators have limited cash at hand, which could only cover the learning expenditure instead of making the final investment. Thus, innovators have to seek external funding to finance the investment. This chapter adopts the ex-ante contracting setting where informed innovators contract with less-informed investors before exerting their internal funding. It is analogous to the informed-principal problem by Maskin and Tirole (1992). Note that the informed-principal setting inverts the timing of typical signalling game, for example, Spence (1973). In Spence's signalling model, workers acquire educational certificates so as to send signals to uninformed employers before signing labour contracts. One can show that this educational signalling game involves multiple perfect Bayesian equilibria involving the pooling and separating ones, some of which are inefficient and fail to fulfil the Cho-Kreps intuitive criterion (Cho and Kreps, 1987, Laffont and Martimort, 2009). This chapter adopts an ex-ante contracting setup which helps to eliminate those inefficient equilibria and selects the least-costly separating equilibrium.

Define α^θ as the revenue share of the investor in the event of success, $\theta = \{h, l\}$. In the event of failure, both parties get zero payoff. Let the contract proposed by

θ-type agent include internal funding and revenue share, such that $C^\theta = \{\lambda^\theta, \alpha^\theta\}$. As there are only two types of innovators, we could concentrate on two possible incentive compatible contracts, i.e., C^h and C^l, corresponding to the revelation principle.

In the environment of imperfect information, the innovator, regardless of his type, is assumed to propose a menu of two contracts to the investor. After updating beliefs, the investor decides whether to accept the offer or not. If accepted, the innovator chooses one of the contracts to execute; otherwise, they propose to another investor. According to Tirole (2006), delivering a menu of two contracts to the uninformed party is to "get rid of their bad expectations". Intuitively, the low-type gains by mimicking the high-type contract, and the investor makes a loss under imperfect information. Thus, the investor is reluctant to take a high-type contract solely as it might be from a low-type agent under information asymmetry. However, such a bad expectation from the investor would be eliminated if the high-type agent includes a low-type incentive-compatible contract into the proposal, which makes the investor break-even on average. After contracting, innovators choose one of the contracts to perform and initiate the experimentation.

3.3.1 Perfect information equilibrium contracts

As a benchmark, we first consider the case where investors have complete information about the prior belief of the project being good when contracting with the innovator. Learning is necessary to explore whether the project indeed is good or bad. From the previous section, we know that $\bar{\lambda}^h$ and $\bar{\lambda}^l$ are the first-best internal financing levels, which are supposed to be achievable under full information contracting. In this setup, investors would sign the contract if the expected net return is no less than the outside option of 0. Thus, α^θ is proposed to make the investor at least break-even under θ-type financial contract. Denote the expected payoff of the investor and innovator at period 0 as $V^\theta(C^\theta)$ and $U^\theta(C^\theta)$ respectively where:

$$\begin{aligned} V^\theta(C^\theta) &= e^{-rT}[p_0^\theta + (1-p_0^\theta)e^{-\lambda^\theta T}](\alpha^\theta p_T^\theta R - I) \\ &= e^{-rT}\{p_0^\theta \alpha^\theta R - [p_0^\theta + (1-p_0^\theta)e^{-\lambda^\theta T}]I\}, \\ U^\theta(C^\theta) &= e^{-rT}(1-\alpha^\theta)p_0^\theta R - \lambda^\theta. \end{aligned}$$

Note that U^θ is monotonically decreasing in λ. As internal funding increases, it builds up one's posterior belief in the project being good in the event of no bad news. However, at period 0, the innovator's expected payoff depends on the ex ante success probability rather than the updated posterior belief. Therefore, the increase of internal funding solely adds extra cost to innovator's total expected payoff.

Here we write down the optimisation programme of θ-type innovators under full information:

$$\max_{C^\theta} U^\theta(C^\theta), \quad \text{subject to } V^\theta(C^\theta) \geq 0.$$

The investor breaks even, so $V^\theta(C^\theta) = 0$. Substituting this condition back into the maximisation problem gives the following:

$$\bar{\lambda}^\theta = -r + \frac{1}{T} \ln T(1 - p_0^\theta) I; \tag{3.4}$$

$$\bar{\alpha}^\theta = \frac{I}{p_T^\theta R} = \frac{I}{\frac{p_0^\theta}{p_0^\theta + \frac{1}{e^{-rT}TI}} R} = \frac{I}{R} + \frac{1}{p_0^\theta R e^{-rT} T}, \tag{3.5}$$

where $\bar{\lambda}^\theta$ is the same as the first-best internal financing in Equation 3.3. For $p_0^h > p_0^l$, $\bar{\lambda}^h < \bar{\lambda}^l$ and $\bar{\alpha}^h < \bar{\alpha}^l$. Thus, the high-type innovators allocate less internal funding and keep a larger fraction of revenue themselves compared to low-type innovators. Moreover, given the optimal internal funding, the posterior beliefs about the good state at time T are $p_T^h = \frac{p_0^h}{p_0^h + \frac{1}{e^{-rT}TI}} > p_T^l = \frac{p_0^l}{p_0^l + \frac{1}{e^{-rT}TI}}$. Thus, at time of investing, the high-type projects still have a higher posterior belief about succeeding than the low-type projects, and they cost less internal funding, resulting in a higher expected returns.

To summarise, the full-information equilibrium contracts for high and low types of innovators are $\bar{C}^h = \{\bar{\lambda}^h, \bar{\alpha}^h\}$ and $\bar{C}^l = \{\bar{\lambda}^l, \bar{\alpha}^l\}$ with first-best internal financing and under which the innovators capture the entire surplus and the investors break even in expectation. As shown in Figure 3.2, the optimal allocations of $\bar{\lambda}$ and $\bar{\alpha}$ are the tangency points of the innovators' indifference curves (U^h and U^l) and investors' zero-expected-payoff curves under the high and low types of equilibrium contracts (V^h and V^l).

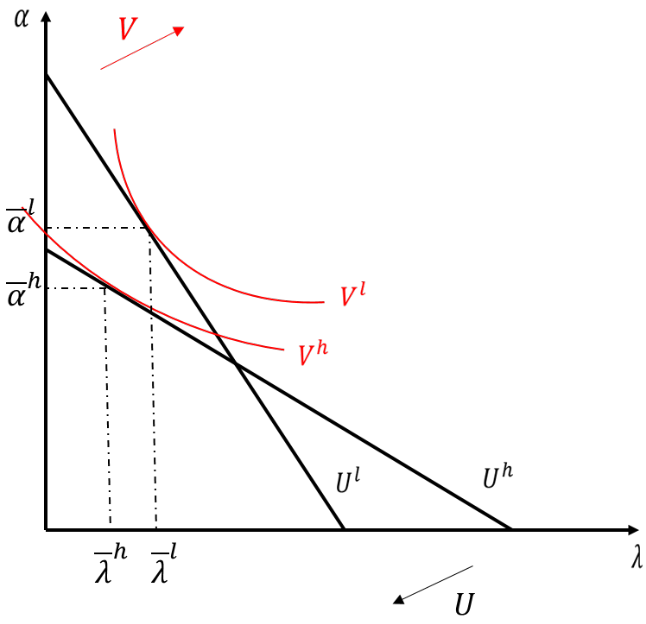

Figure 3.2: Perfect information contracting

3.3.2 Equilibrium contracts under imperfect information

When investors have imperfect information about the type of the project, the low-type innovators have strong incentives to undertake the high-type optimal contract, shown as an inward shifting of the low-type indifference curve in Figure 3.3, which makes full-information equilibrium contracts no longer stable. This sections attempts to develop perfect Bayesian Equilibria under imperfect information and pays particular attention to the Pareto-dominant one, or the least-cost separating equilibrium.

Tirole (2006) formally concludes the procedure of determining a Pareto dominant equilibrium in signalling game in the book, *The Theory of Corporate Finance* (Tirole, 2006, pp. 267–269), and we will follow in this paper. This section starts by characterising a low-information intensity equilibrium, also known as the Rothschild-Stiglitz-Wilson (RSW) allocations for the high-type and low-type (Rothschild and Stiglitz, 1976; Wilson, 1977), that should be by stable and incentive compatible for both types of innovators regardless of the prior belief of the investor by definition (Maskin and Tirole, 1992). The RSW allocations guarantee the payoff that each type should at least receive in any other equilibrium. Later on, we check if the RSW allocations are interim efficient with respect to the

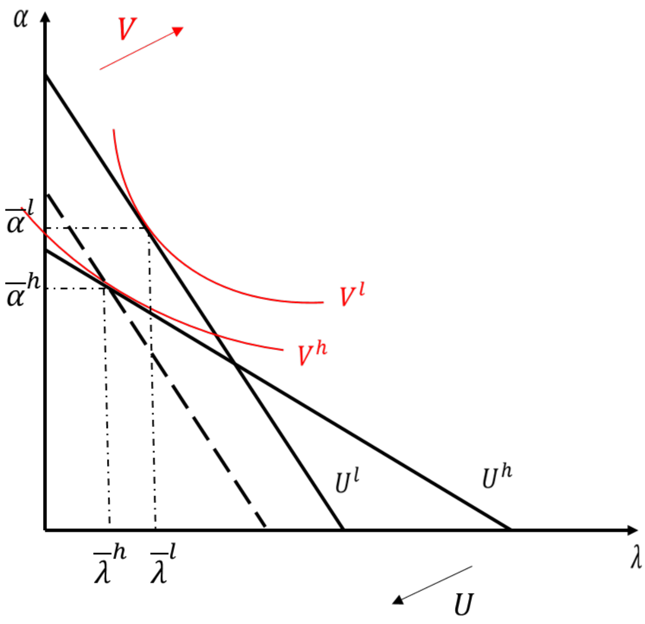

Figure 3.3: Deviation of the low-type

investor's prior belief; if so, a Pareto-dominant equilibrium could be determined.

Definition 3 $(\hat{\lambda}^h, \hat{\alpha}^h)$ *is an RSW-h allocation of the high-type agent if it maximises the total expected payoff of the high-type subject to a set of incentive compatible allocations for both types of innovators and non-negativity expected profit constraints of the investor regardless of his beliefs.*

The RSW-h allocation can be found by solving the following maxmisation problem:

Programme A (high-type)

$$\begin{aligned} \max_{C^h, C^l} \quad & U^h(C^h) \\ \text{subject to} \quad & U^h(C^h) \geq U^h(C^l), && (IC^h) \\ & U^l(C^l) \geq U^l(C^h), && (IC^l) \\ & V^h(C^h) \geq 0, && (IR^h) \\ & V^l(C^l) \geq 0, && (IR^l) \end{aligned}$$

where IC^h and IC^l are incentive compatibility constraints which make sure innovators in the equilibrium choose their type-equivalent contracts, and IR^h and

IR^l are to guarantee the investor earns non-negative profits on each type. In intuition, the high-type cannot achieve the first-best in this equilibrium and has to sacrifice part of the revenue in order to make the high-type contract unattractive to the low-type. There is an analogous maximization programme for the low-type agent. Let $(\hat{\lambda}^l, \hat{\alpha}^l)$ be the RSW-l allocation of the low-type innovators. As an RSW equilibrium could be achieved regardless of the investor's prior beliefs in the agent's type under imperfect information, the agent could guarantee himself an RSW payoff by proposing $(\hat{\lambda}^\theta, \hat{\alpha}^\theta)$, and the corresponding RSW payoff can be regarded as one's reservation utility.

Lemma 4 *The low-type agent receives no more than his full-information equilibrium payoff in the RSW-l allocation, that is,* $U^l(\hat{\lambda}^l, \hat{\alpha}^l) \leq U^l(\bar{\lambda}^l, \bar{\alpha}^l)$.

Proof. For $(\bar{\alpha}^l, \bar{\lambda}^l)$ solves the first-best maximisation problem of low-type with the individual rationality constraint of the investor being binding, thus, $V^l(\bar{\lambda}^l, \bar{\alpha}^l) = 0$. In the Programme A of the low-type, $V^l(\hat{\lambda}^l, \hat{\alpha}^l) \geq 0$, which means $V^l(\hat{\lambda}^l, \hat{\alpha}^l) \geq V^l(\bar{\lambda}^l, \bar{\alpha}^l)$. Thus, $U^l(\hat{\lambda}^l, \hat{\alpha}^l) \leq U^l(\bar{\lambda}^l, \bar{\alpha}^l)$ holds as the agent claims the remaining surplus of the project. ■

Intuitively, the low-type innovators maximise profit given that the individual rationality constraint of the investor binds under full information. That is, for $\bar{C}^l$ to be the first-best contract, the investor makes zero profit: $V^l(\bar{C}^l) = 0$. Under Programme A, the investor makes a non-negative profit on both types of innovators irrespective of his expectation. Thus, the low-type agent can offer a menu with just one low-type contract $\{\bar{C}^l, \bar{C}^l\}$, and the investor will make either zero payoff if the agent is a low-type or positive profits if the agent is a high-type, hence the investor will accept the offer. This forms an RSW equilibrium and guarantees the first-best payoff for the low-type.

Figure 3.4 shows the possible RSW equilibrium where none of the agents has incentives to deviate. Firstly, following from the lemma, the low-type could achieve the first-best under RSW equilibrium, showing as the point $(\bar{\lambda}^l, \bar{\alpha}^l)$ in the graph. Secondly, to satisfy incentive compatibility, the high-type must sacrifice part of the welfare in order to prevent the low-type from mimicking, and the best they could do is to make the low-type indifferent between choosing either of the contracts. In other words, IC^l must be binding in this case. As shown in the figure,

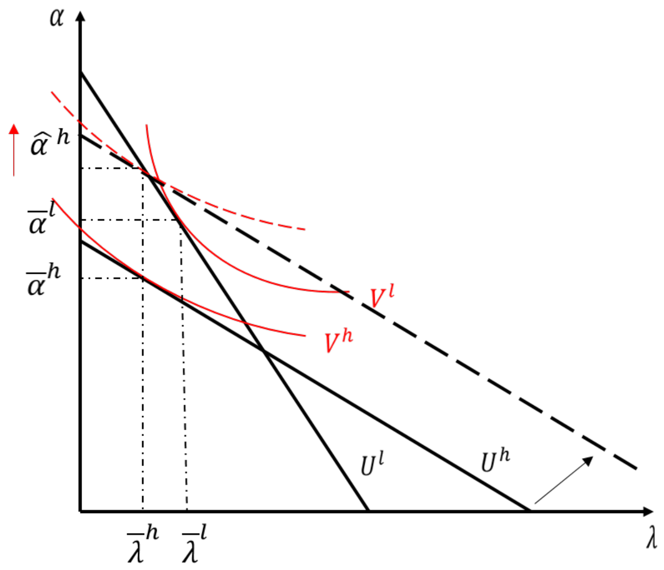

Figure 3.4: RSW equilibrium

$(\bar{\lambda}^h, \hat{\alpha}^h)$ is the RSW^h allocation as it is the utility maximisation of the high-type with the non-negativity profit constraint of the investor being satisfied and the IC^l constraint binding. As the indifference curves of V^h are vertical translations of each other, $\bar{\lambda}^h$ stays at the first-best, however, with revenue shares increasing from $\bar{\alpha}^h$ to $\hat{\alpha}^h$, so the high-type suffers a welfare loss under this equilibrium relative to the first-best. Then we need to check if there exists any equilibrium that Pareto dominates this RSW^h allocation, makes the high-type better off and achieves a full separation between the high-type and low-type. Then we combine the investor's prior belief of the agent's type and look for an equilibrium makes the investor break even on average.

Definition 5 *An incentive compatible set of contracts $\{\tilde{C}^h, \tilde{C}^l\}$ is interim efficient if it make the investor earn non-negative profit in expectation relative to the prior belief about the agent's type, $\{q, 1-q\}$.*

As the investor has a prior belief about the agent's type, the low-information intensity equilibrium cannot satisfy the interim efficiency according to the definition above: RSW^h leaves the investor positive profit, and RSW^l allocates zero profit to the investor. Whereas, the high-type innovator can improve his profit by making the investor break-even in expectation. We then relax the type-by-type non-negativity constraints of the investor to check if there exists another equilibrium that is Pareto dominating the low-information intensity equilibrium as well as interim efficient. With probability q, the investor believes the agent is a high-type, and conversely, with probability $(1-q)$, the agent is regarded as a low-type. Given this belief, the investor would accept the contract if he breaks-even on average. In this case, we can find an efficient equilibrium that is interim efficient relative to the prior belief $\{q, 1-q\}$ and incentive compatible for both types of agents. The term "interim" refers to the stage where innovators propose the financial contract to the investor with private information on their types. Moreover, an interim efficient equilibrium weakly Pareto dominates the RSW equilibrium, which is the key results of Maskin and Tirole (1992). Intuitively, as the individual rationality constraints of the investor are relaxed, the interim efficient equilibrium allows the investor to make a loss on one of the types of the innovators, which leads to a Pareto improvement without hurting the other type's benefit.

Following Lemma 4, the low-type could at most receive their first-best payoff in the RSW equilibrium. Thus, the first-best payoff should be guaranteed in the Pareto dominating equilibrium. The following the maximisation programme leads to the interim efficient and Pareto dominating outcomes:

Programme B (high-type)

$$\begin{aligned}
\max_{C^h, C^l} \quad & U^h(C^h) & \\
\text{subject to} \quad & U^h(C^h) \geq U^h(C^l), & (IC^h) \\
& U^l(C^l) \geq U^l(C^h) & (IC^l) \\
& qV^h(C^h) + (1-q)V^l(C^l) \geq 0 & (IR) \\
& U^l(C^l) \geq U^l(\bar{C}^l) & (IE^l)
\end{aligned}$$

where IC^h and IC^l are the incentive compatibility constraints of high and low types of innovators respectively, IR is the individual rationality constraint of the investor, which ensures he gets non-negative profit in expectation and guarantees the interim efficiency, and IE^l constraint indicates that the low-type agent should receive at least their first-best expected profit in this equilibrium.

3.3.2.1 The least costly separating equilibrium

Solve the Programme B fully and get the separating equilibrium below.

Proposition 5 *There exists a least-costly separating equilibrium:* $\{\lambda^{l^*}, \alpha^{l^*}\}$, $\{\lambda^{h^*}, \alpha^{h^*}\}$ *such that:*

- *Regarding internal funding, the low-type innovators remain at their first-best optimum, whereas, the high-type innovators invest even more than the low-type first-best internal funding such that* $\lambda^{l^*} = \bar{\lambda}^l$*, and* $\lambda^{h^*} > \bar{\lambda}^l > \bar{\lambda}^h$*;*

- *For the revenue share, the low-type innovators allocate no more than the full-information revenue shares, but more than the high type revenue shares to the investor; thus the following holds:* $\alpha^{h^*} < \alpha^{l^*} \leq \bar{\alpha}^l$*.*

Proof. See appendix. ■

In practise, the separating equilibrium works as follows: the innovator, either high or low type, proposes a menu of $\{\lambda^{l^*}, \alpha^{l^*}\}$ and $\{\lambda^{h^*}, \alpha^{h^*}\}$ contracts to the investor. Upon observing these two contracts, the investor still stays uninformed of the type of the innovator but will accept the offer as it makes him break-even in expectation. Then the innovator will execute one of the contracts according to her type as there is no profitable deviation for either party. However, in the case when the informed innovator proposes only one contract of her type, the investor will update his belief in agent's type accordingly. The low-type would have incentive to deviate as the high-type contract is most likely to be accepted by the investor. Thus, to avoid further distortion, a set of two contracts should be proposed ex ante.

Proposition 5 states that the low-type innovators remain at their full-information equilibrium internal financing, whereas, the high-types invest more than their first-best internal financing. Firstly, in terms of the internal financing allocation, $\{\lambda^{h^*}, \bar{\lambda}^l\}$ allows the high-type being fully separated from the low-type. By over-investing on learning, the high-type refrains the low-type from mimicking under imperfect information. As the internal financing of the high-type increases from $\bar{\lambda}^h$ to λ^{h^*}, the incentive compatibly constraint of the low-type agents has been relaxed. Thus, it makes the low-type agents prefer choosing their full-information contracts to the high-type contract, which ensures as much as an RSW optimum payoff.

Moreover, as $\lambda^{h^*} > \bar{\lambda}^l$, the internal funding of the high-type increases even further and turns to be greater than low-type's first-best. By definition of separating equilibrium, all innovators would prefer their type-equivalent contracts. The separation would not be possible if the high-type chose less than the low-type's first-best internal funding, which gives rise to a possible deviation by the high-type. This can be verified intuitively as the single-crossing property holds in this case. Let MU_α^θ and MU_λ^θ denote the marginal utility of revenue share and internal funding given a certain utility level for the θ-type agent. We can calculate the marginal rate of substitution of revenue for internal funding as a ratio of MU_α^θ over MU_λ^θ respectively for high-type and low-type. Taking the difference between the ratio of high-type and low-type , we find the following relationship:

$$\frac{MU_\alpha^h}{MU_\lambda^h} - \frac{MU_\alpha^l}{MU_\lambda^l} > 0, \tag{3.6}$$

which indicates that the high-type are willing to sacrifice more internal funding in exchange for a further decrease of revenue shares compared with the low-type. As the high-type expects success with a higher probability, so the revenue share matters more to the high-type than to the low-type. As a result, the high-type would allocate more internal funding than the low-type first-best in order to preserving a larger proportion of revenue shares in the equilibrium.

Secondly, we analyse the welfare under this separating equilibrium. Under the financial contract $\{\bar{\lambda}^l, \alpha^{l^*}\}$, as $\alpha^{l^*} \leq \bar{\alpha}^l$ the low-type would achieve at least their full-information payoff and have no incentive to deviate. $\bar{\lambda}^l$ amount of internal funding guarantees the low-type a sufficient learning and generates the optimum expected payoff, which means they cannot get any better by adjusting the spending. Whereas, reducing the revenue share from $\bar{\alpha}^l$ in the RSW equilibrium to α^{l^*} makes the low-type weakly better off. Moreover, although more internal funding than the first-best is provided by the high-type, the revenue shares they need to sacrifice to make the separation is less than the RSW equilibrium, which also results in a welfare improvement. Thus, this interim efficient equilibrium weakly Pareto dominates the above-mentioned RSW equilibrium in terms of the welfare of both agents. In addition, it manages to make the separation between two types of the agents. Although the investor is not perfectly informed ex ante, proposing this set of two contracts ensures his participation, and he would also believe that each type of innovators will act according to their types after contracting.

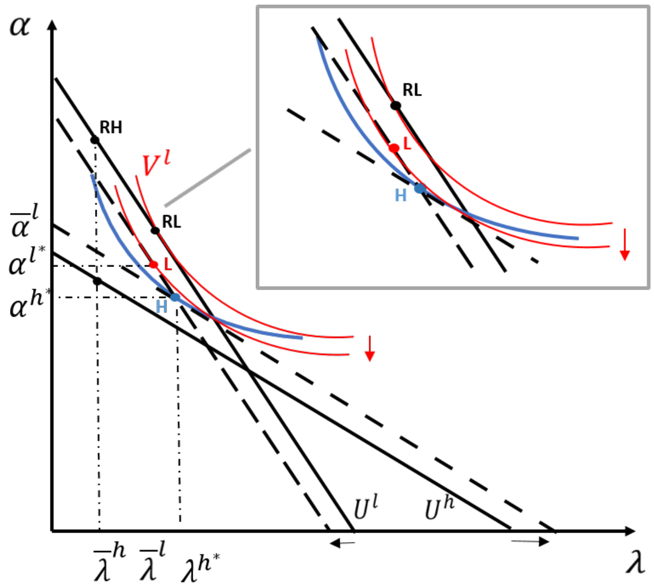

Figure 3.5: Separating equilibrium

In Figure 3.5, the point RH and RL shows the RSW equilibrium allocations for the high-type and low-types respectively. An outward shift of the high-type's indifference curve under full information U^h (the solid black line) generates a tangent at the point H to investor's zero-average-profit curve (the blue line), which represents a cross-subsidisation from the high to the low and a fulfilling of the interim efficiency of the IR constraint in Programme B. As a response, the low-type could make themselves better off by inward shifting the indifference curve U^l until it crosses high-type's new allocation. Any point on that dotted new indifference curve makes the low-type receive as much payoff as if taking the high-type contract. The tangency point L with the red indifference curve of the investor V^l maximises the investor's profit subject to the incentive constraint holding given that she is facing the low-type. As a set of red indifference curves of the investor V^l are vertical translations of each other, $\bar{\lambda}^l$ remains to be the optimum internal funding. Whereas, the inward shifting results in a revenue share decrease from $\bar{\alpha}^l$ to α^{l*}, which, in all, illustrates that point L, $(\bar{\lambda}^l, \alpha^{l*})$, is low-type's best attainable payoff under separation. Compared with the RSW equilibrium allocations (pont RH and RL), both types of agents are better off at this new separating equilibrium (point H and L). However, under this separating equilibrium, the investor makes strictly positive profit since at point she makes more profit than she would have against the low-type. Thus, the high-type can do better by reducing the revenue share α^h, hence leave the investor zero profit in expectation, which is not included in the graph.

In the following, we deliver some implications and thoughts from Proposition 5. To begin with, the above separating equilibrium can be related with Spence's (1973) labour market signalling model where heterogeneous employees decide how much costly education to acquire under adverse selection. The paper also proposes a separating equilibrium under which the high-type agents are willing to get an education just for signalling purpose and the low-type agents get no education. In this case, education severs as a signal to less-informed employers as it's more costly for the low-type to acquire the same education compared with the high-type. In our model, the high-type innovators would otherwise undertake less internal funding than the low-type given a higher expected probability of success. Directly driven by the purpose of separating, the high-type innovators are willing to spend even more internal funding than the low-type, which leads to a Pareto dominating separating equilibrium.

Proposition 5 also echoes with the well-known pecking order theory according to which firms prefer internal funding over external resources and debt over equity at the external financing stage in the presence of information frictions (Myers and Majluf, 1984; Myers, 1984). This theory has also be confirmed by some empirical studies as surveyed in Frank and Goyal (2011). Our model predicts that innovators regardless of their types would allocate at least the first-best level of internal financing to support learning under imperfect information. Although external financing happens endogenously between investor and innovator without any cost, if borrowing is costly we expect that the high-type innovators would adopt more of the internal funding and borrow a smaller proportion of the external funding in comparison with the low-type ones. As the innovator's investing decision is contingent on the expected probability of success, it follows that innovators with high-type project tend to provide more internal funding based on their higher probability of success. Additionally, by doing so, the high-type could save extra borrowing cost, and less compensation is needed for the less-informed outside

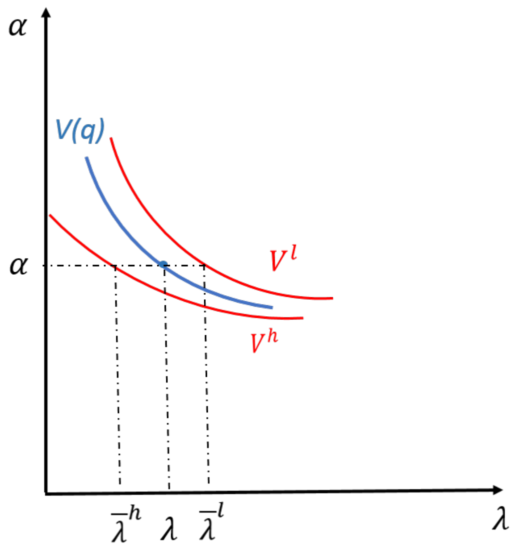

Figure 3.6: Pooling equilibrium with α fixed

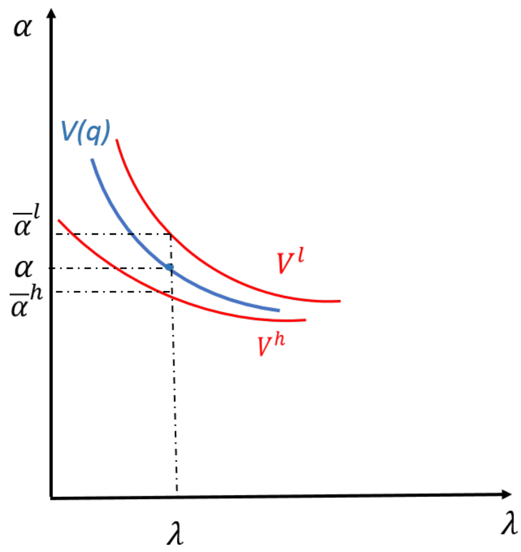

Figure 3.7: Pooling equilibrium with λ fixed

investor. Thus, implications could be found in our model as well such that heterogeneous agents reveal a slightly different preference over internal and external financing.

3.3.2.2 Pooling equilibria

This section we discuss if there exists any pooling equilibrium when only one of the contractual terms is available, either revenue share or internal funding. The key criterion of a pooling equilibrium is to make sure there is no profitable deviation for either type of innovators and guarantee the investor on average to break-even. To achieve this, the high-type may need to compromise on their revenue in order to subsidise the investor's welfare losses on the low-type.

Proposition 6 *Assume the revenue share to be exogenously fixed, the high and low types of innovators choose the same amount of internal funding as a pooling equilibrium, such that* $\lambda^h = \lambda^l = \lambda$.

Proof. See appendix. ■

Proposition 7 *Assume the internal funding to be exogenously fixed, the high and low types of innovators choose the same revenue share as a pooling equilibrium, such that* $\alpha^h = \alpha^l = \alpha$.

Proof. See appendix. ■

As shown in Figure 3.6 and 3.7, $V(q)$ is the investor's zero-profit indifference curve as a linear combination of his indifference curves when contracting with the high and low respectively, denoted as $V(q) = qV^h(C) + (1-q)V^l(C) = 0$. Notice that both pooling equilibria involve a cross-subsidy from the high to the low. In the first case, with exogenously fixed α, high and low types pool at same internal funding level, λ. It turns out that the high-types tend to undertake more internal financing than needed under full-information, i.e., $\lambda > \bar{\lambda}^h$, as shown in Figure 3.6, which results in a positive gain of the investor. However, the low-type would choose a lower amount of internal financing than their full-information equilibrium, i.e., $\lambda < \bar{\lambda}^l$, which induces a loss of the investor. In this case, as there is no profitable deviation for both types, it can be regarded as a stable equilibrium outcome. Moreover, as the allocation $\{\alpha, \lambda\}$ is on the $V(q) = 0$ indifference curve, the investor is willing to accept the financial contract and exert external investment.

Similarly, as stated in Proposition 7, with internal funding exogenous, the only signalling tool available is the revenue share. The high-types are willing to give a higher revenue share to the investor in order to compensate her payoff-loss from the low-type. On the other hand, the best reaction of the low-types is to adopt this α-revenue share, which generates a positive gain without being distinguished from the high-type.

As the above two pooling equilibria are achieved with either of the signalling tool available, it's worth exploring the existence of any other pooling equilibrium under two signalling dimensions.

Proposition 8 *There exists no other Pareto dominating pooling equilibrium when both revenue shares and internal funding are available under imperfect information, which proves the uniqueness of the least costly separating equilibrium.*

Proof. See appendix. ■

3.4 Conclusion

This chapter presents an exponential innovation model to show heterogeneous innovation firms' equilibrium options of internal and external financing. Given that prior belief of the project being good is private information, a separating equilibrium contract indicates that the high-type could be distinguished from the low-type by committing to a higher level of internal financing. Otherwise, heterogeneous innovators will be pooled at same revenue share and internal financing level where the high-type is strictly worse-off due to cross-subsidisation of the low-type. The model manages to capture several features of real-world innovation: costly learning is required before investment; the outcome of an innovation project is uncertain; the more in depth the innovator learns, the higher the posterior belief of success conditional on no bad news gets. Several important predictions could be made from our model: innovators with the high-type project usually put more effort in the initial learning stage, hence are more likely to succeed; investors could judge a project's success likelihood by looking into the innovator's internal financing on experimentation; internal financing is more preferred to the high-type especially when external borrowing is more costly, or information asymmetry problem is severer compared with the low-type.

This article can be extended in the following ways: first, let the innovation firms become less cash-constrained where internal funding can cover part of the investment. Attention could be focused on heterogeneous firms' optimal internal and external financing allocation under information frictions. Second, initial learning stage of innovation could be bought under a contest environment where only the most promising project could be funded. It is similar to Halac et al. (2016b) where the principal plays a role of attributing prizes and providing incentives to contestants via partly disclosing information. However, innovators will hold private information on their prior beliefs in our contest setting.

Appendix

3.A Proof of Proposition 5

Here we illustrate the way of solving the Programme B :

$$\begin{aligned}
\max_{\lambda^h,\alpha^h,C^l} \quad & U^h(\alpha^h,\lambda^h) = e^{-rT}(1-\alpha^h)p_0^h R - \lambda^h \\
\text{subject to} \quad & U^h(C^h) \geq U^h(C^l), && (IC^h) \\
& U^l(C^l) \geq U^l(C^h) && (IC^l) \\
& qV^h(C^h) + (1-q)V^l(C^l) \geq 0 && (IR) \\
& U^l(C^l) \geq U^l(\bar{C}^l) && (IE^l)
\end{aligned}$$

Firstly, as the IR constraint binds at the optimum, we obtain the following condition:

$$qV^h(\alpha^h,\lambda^h) + (1-q)V^l(\lambda^l,\alpha^l) = 0. \tag{3.7}$$

As the expected payoffs of high-type innovators could be expressed as the difference between expected surplus of the high-type project and expected payoff of the investor, $U^h(\alpha^h,\lambda^h)$ can be rewritten as the following:

$$\begin{aligned}
U^h(\alpha^h,\lambda^h) &:= e^{-rT}(1-\alpha^h)p_0^h R - \lambda^h \\
&= e^{-rT}\{p_0^h R - [p_0^h + (1-p_0^h)e^{-\lambda^h T}]I\} - \lambda^h - V^h(\alpha^h,\lambda^h)
\end{aligned}$$

Substitute (7) into $U^h(\alpha^h,\lambda^h)$ and cancel variable α^h:

$$U^h(\lambda^h) = e^{-rT}\{p_0^h R - [p_0^h + (1-p_0^h)e^{-\lambda^h T}]I\} - \lambda^h + \frac{1-q}{q}V^l(\lambda^l,\alpha^l)$$

We ignore IC^h at this stage as it will be shown to be satisfied ex post.

Lemma 6 *In Programme B, if the IE^l constraint is binding, the IC^l constraint must be binding.*

Proof. Under unconstrained maximisation problem, $\bar{\lambda}^h$ and $\bar{\alpha}^h$ are the first-best solution of high-type agents. If the IE^l constraint is binding under this constrained optimisation problem of the high-type, thus for $\{\lambda^l, \alpha^l\}$:

$$U^l(\lambda^l, \alpha^l) = U^l(\bar{\lambda}^l, \bar{\alpha}^l) < U^l(\bar{\lambda}^h, \bar{\alpha^h}) \quad holds,$$

which means IC^l would be violated at the first-best contracts. This implies that IC^l must be binding under constrained optimisation problem of the high-type. In other words, $\{\bar{\lambda}^h, \bar{\alpha^h}\}$ are no longer achievable as a second-best solution for the high-type agents. ■

As IE^l constraint binding is the sufficient condition that leads IC^l to be binding according to Lemma 6, so Programme B becomes:

$$\begin{aligned} \max_{\lambda^h} \quad & U^h(\lambda^h) = e^{-rT}\{p_0^h R - [p_0^h + (1-p_0^h)e^{-\lambda^h T}]I\} - \lambda^h + \frac{1-q}{q}V^l(\lambda^l, \alpha^l) \\ \text{subject to} \quad & U^l(\lambda^h) \geq U^l(\bar{\lambda}^l, \bar{\alpha}^l) \qquad (IE^{l\prime}) \end{aligned}$$

The key to this high-type maximisation problem is defining a λ^h that makes $IE^{l\prime}$ binding. So we rewrite the left-hand-side of $IE^{l\prime}$ as the following:

$$\begin{aligned} U^l(\lambda^h) &= \frac{p_0^l}{p_0^h}[e^{-rT}p_0^h(1-\alpha^h)R] - \lambda^h \\ &= \frac{p_0^l}{p_0^h}(U^h(\lambda^h) + \lambda^h) - \lambda^h \\ &= \frac{p_0^l}{p_0^h}U^h(\lambda^h) + (\frac{p_0^l}{p_0^h} - 1)\lambda^h \end{aligned} \tag{3.8}$$

The objective function $U^h(\lambda^h)$ is concave in λ since $U^{h\prime\prime}(\lambda^h) < 0$, which guarantees a unique $\tilde{\lambda}^h$ that maximises U^h. Moreover, for any $U^{h\prime}(\lambda^h) < 0$, we could deduce that $\lambda^h > \tilde{\lambda}^h$.

Then we study the properties of $U^l(\lambda^h)$ in order to identify the λ^h which makes $IE^{l\prime}$ binding. $U^l(\lambda^h)$ shows the same concavity property in λ^h. As $\frac{dU^l}{d\lambda^h} = \frac{p_0^l}{p_0^h}U^{h\prime}(\lambda^h) + (\frac{p_0^l}{p_0^h} - 1)$, and for $\bar{\lambda}^h$ to be the first-best, $U^{h\prime}(\bar{\lambda}^h) = 0$ holds. Thus, we find out that: $U^l(\lambda^h)$ decreases at the $\bar{\lambda}^h$-neighbourhood as $\frac{dU^l}{d\lambda^h} < 0$. Moreover, since $U^l(\bar{\lambda}^h) > U^l(\bar{\lambda}^l)$ and $U^{l\prime}(\bar{\lambda}^h) < 0$, there exists a $\lambda^{h*} > \bar{\lambda}^h$ that satisfies $U^l(\lambda^{h*}) = U^l(\bar{\lambda}^l)$, and IE^l binds.

We then turn our attention to the IC^h constraint: $U^h(C^h) \geq U^h(C^l)$. Rearrange the right-hand-side $U^h(C^l)$ as following:

$$U^h(C^l) = \frac{p_0^h}{p_0^l} U^l(C^l) + (\frac{p_0^h}{p_0^l} - 1)\lambda^l \tag{3.9}$$

Substituting the condition that IC^l binds in the equilibrium, i.e., $U^l(C^l) = U^l(C^h)$, into the above $U^h(C^l)$:

$$\begin{aligned} U^h(C^l) &= \frac{p_0^h}{p_0^l} U^l(C^h) + (\frac{p_0^h}{p_0^l} - 1)\lambda^l \\ &= e^{-rT} p_0^h (1-\alpha^h) R - \lambda^h + (\lambda^h - \lambda^L)(1 - \frac{p_0^h}{p_0^l}) \\ &= U^h(C^h) + (\lambda^{h^*} - \lambda^L)(1 - \frac{p_0^h}{p_0^l}) \end{aligned}$$

For IC^h to hold: $U^h(C^h) \geq U^h(C^l) = U^h(C^h) + (\lambda^{h^*} - \lambda^L)(1 - \frac{p_0^h}{p_0^l})$, as $(1 - \frac{p_0^h}{p_0^l}) < 0$, thus we deduce that $\lambda^{h^*} > \lambda^L$ should be satisfied.

3.B Proof of Proposition 6 and 7

According to the assumption, $\alpha^h = \alpha^l = \alpha$. As $U^\theta(\lambda)$ is monotonically decreasing in λ, therefore, both incentive compatibility constraints indicate that $\lambda^h \leq \lambda^l$ and $\lambda^l \leq \lambda^h$. Thus, $\lambda^h = \lambda^l = \lambda$ in this case.

As the individual rationality constraint of the investor binds when innovators receive the optimal payoff, we choose the smallest λ that makes IR binds as following:

$$qe^{-rT}\{p_0^h \alpha R - [p_0^h + (1-p_0^h)e^{-\lambda T}]I\} + (1-q)e^{-rT}\{p_0^l \alpha R - [p_0^l + (1-p_0^l)e^{-\lambda T}]I\}$$
$$= 0$$
$$\implies \tilde{\lambda} = \frac{1}{T} ln \frac{I[q(1-p_0^h) + (1-q)(1-p_0^l)]}{(\alpha R - I)[qp_0^h + (1-q)p_0^l]}$$

Since the full-information internal financing of the low-type under fixed revenue share is

$$\bar{\lambda}^l = \frac{1}{T} ln \frac{I(1-p_0^l)}{(\alpha R - I)p_0^l}$$

Thus, $\tilde{\lambda} \leq \bar{\lambda}^l$ for $q \in [0,1]$ and IE^l holds. Thus Proposition 6 is proved.

Following the same reasoning. For $\lambda^h = \lambda^l = \lambda$, two IC constraints imply that $\alpha^h = \alpha^l = \alpha$. Then according to the IR constraint's binding condition, one could show that $\tilde{\alpha} \leq \bar{\alpha}^l$, then the IE^l constraint holds as well.

3.C Proof of Proposition 8

Prove by contradiction. Suppose that there exists a pooling equilibrium where high and low types of innovators propose the same financial contract, i.e $C^h = C^l$. In the following we first solve the Programme B given $C^h = C^l$ to see the most efficient outcome in the high-type's best interest. Then we attempt to analysis the symmetric Programme B of the low-type when $C^h = C^l$. If the outcomes of these two programmes coincide, there is a pooling equilibrium as a solution of Programme B. If not, there is always profitable deviation for either type, which results in no pooling result under imperfect information and confirms the uniqueness of the separating equilibrium.

In Programme B, when $C^h = C^l$, IC^h and IC^l hold with equality and could be ignored. We first ignore IE^l and check whether it will be satisfied later. Write down the reduced version of Programme B as the following:

$$\begin{aligned} \max_{\lambda,\alpha^h} \quad & U^h(\alpha^h,\lambda^h) = e^{-rT}(1-\alpha^h)p_0^h R - \lambda^h \\ \text{subject to} \quad & V(q) = qV^h(\lambda^h,\alpha^h) + (1-q)V^l(\lambda^h,\alpha^h) \geq 0 \end{aligned} \qquad (IR)$$

As the optimum allocation is where indifference curve of the high-type tangent with investor's zero-profit curve where $MRS^h_{\alpha^h,\lambda^h}(\bar{U}) = MRS_{\alpha^h,\lambda^h}(V(q))$ for given utility level $\bar{U}$ and 0, we calculate in the following way:

$$\begin{aligned} MRS^h_{\alpha^h,\lambda^h}(\bar{U}) &= \frac{\partial U/\partial \alpha^h}{\partial U/\partial \lambda^h} = e^{-rT}p_0^h R \\ MRS_{\alpha^h,\lambda^h}(V(q)) &= \frac{\partial V/\partial \alpha^h}{\partial V/\partial \lambda^h} = \frac{qp_0^h R + (1-q)p_0^l R}{[q(1-p_0^h) + (1-q)(1-p_0^l)]TIe^{-\lambda^h T}} \end{aligned}$$

then equalise the two MRS finding the tangent point:

$$\begin{aligned} \lambda^h &= -r + \frac{1}{T}ln\frac{TIp_0^h[q(1-p_0^h) + (1-q)(1-p_0^l)]}{qp_0^h + (1-q)p_0^l}, \\ \alpha^h &= \frac{I}{R} + \frac{1}{p_0^h R e^{-rT} T} = \bar{\alpha}^h. \end{aligned}$$

According to the same logic, we solve the low-type optimisation programme:

$$\begin{aligned} \max_{\lambda,\alpha^l} \quad & U^l(\alpha^l,\lambda^l) = e^{-rT}(1-\alpha^l)p_0^l R - \lambda^l \\ \text{subject to} \quad & V(q) = qV^h(\lambda^l,\alpha^l) + (1-q)V^l(\lambda^l,\alpha^l) \geq 0, \end{aligned} \qquad (IR)$$

get the following results:

$$\lambda^l = -r + \frac{1}{T} ln \frac{TIp_0^l[q(1-p_0^h) + (1-q)(1-p_0^l)]}{qp_0^h + (1-q)p_0^l},$$
$$\alpha^l = \frac{I}{R} + \frac{1}{p_0^l Re^{-rT}T} = \bar{\alpha}^l.$$

We can see the contradictions here where $\lambda^h \neq \lambda^l$ and $\alpha^h \neq \alpha^l$. Moreover, whether IE^l will hold or not depends on the q parameter value, so it's not satisfied under any circumstances.

Chapter 4

Multi-stage innovation financing under imperfect information

4.1 Introduction

4.1.1 Motivation and analysis

Successful innovation often requires completion of multiple phases. Launching a new medicine in the UK requires passing three phases of clinical trials, which usually takes around 10-12 years and costs 1 billion pounds per new medicine. Due to safety and risk concerns, only 1 in 10,000 candidate drugs will eventually succeed (The Faculty Of Pharmaceutical Medicine, 2017). An increasing number of pharmaceutical companies are seeking outside financing sources due to the high development cost. However, it is getting increasingly difficult for those small companies to raise funding than the large ones, which is mainly due to the imperfect information between innovators and investors in the market (Hall and Lerner, 2010). According to a study by McKinsey & Company, over half of the late-stage research is funded externally, and many of them are pursuing novel financing and collaboration modes, such as sharing control right and selling future options (David et al., 2010). With outside funding of multistage innovation projects and asymmetric information about stage completion, it is important that innovators are incentivised to reveal progress to investors in a trustworthy way.

To address the question of how innovators can be incentivized to reveal progress in multistage innovation project, this chapter analyses a simple two-stage innovation project funded by an outside investor when the progress of the project is the private information of the innovator. This is a multistage principal-agent problem with innovation where the innovator is the agent and the investor is the principal. To successfully conclude, the project requires completion of two stages. The first stage could be considered as a development phase that involves initial investigation, planning, hiring of researchers, setting up of processes and equipment and so on. The time of completion of this first phase is unobservable to the investor. For simplicity, we assume that the time of completion is binary and either early or late. The second stage requires investigation or experimentation that may or may not be successful. The probability of the success in the experimentation stage depends on an unknown state of the world. If the state is good, the project will generate a relative high revenue compared to the investment cost; if the state is bad, it realises nothing. Thus, in the second stage, the agent needs to explore the unknown state of the world and may potentially abandon the project if no good news comes out as time passes[1]. The information imperfection on the timing of completion of the development stage and uncertainty on the result may deter the investor from committing to a long-term contract and financing the innovation. Therefore, the question of interest is how to design an optimal financial contract which is incentive compatible for the agent to truthfully reveal the progress of the innovation project with respects to the participation constraint of the investor.

The project progress plays a crucial role in the long-term financial contracting as it determines the funding deadline and relevant amount of compensation. As the agent is short of funding and continuously requires investment funding from the investor, the contract should include decent compensation to make the investor willing to sign the contract. The reward should be paid as soon as the second stage is completed for the following two reasons: first, it is difficult for the two parties to contract on the accomplishment of the first stage as it is the private information of the agent. Second, the agent is short of funding to repay the investor and relies on the breakthrough of the innovation, which will generate a large amount of revenue. Thus, the financial contract under perfect information needs to specify two funding periods with different reward in the event of early-completion and late-completion respectively. We assume that, for simplicity, renegotiation and short-term contract are not allowed in our model. The contract should take all

[1]The state of the world does not affect the accomplishment of the development stage, but the project can only be completed in the second experimentation stage if the state is good.

possible information imperfection scenario into consideration and especially prevent potential deviations of the agent. We show that with some moderation on the bonus contract mentioned above, a truthful revealing of the project progress and the first-best experimentation can be achieved under imperfect information.

This chapter contributes to the theoretical literature on the financing of innovation and addresses the imperfect information of project progress. There are several common features of the models in this literature strand: firstly, the agents are normally short of funding and need financial support from investors. Secondly, the outcomes of the innovation are uncertain, which requires agents to conduct experiment and learn the quality of the innovation. Thirdly, investors are generally less informed than agents, which gives rise to the problem of information asymmetry. However, most of the papers in this literature strand focus on the moral hazard problem where the agent may potentially shirk and delay exert effort under imperfect monitoring by the investor (Bergemann and Hege, 1998, 2005; Horner and Samuelson, 2013). This chapter enriches the discussion of the financing innovation literature by introducing the imperfect information problem of project progress.

This chapter uses the "good news" exponential bandit model to describe the random arrival of a success in the experimentation stage, which is widely used in the theoretic literature of strategic experimentation (Keller et al., 2005; Horner and Samuelson, 2013; Moroni, 2016). This chapter is novel in this literature strand that demonstrates the information spillover effect between the development stage and experimentation stage during innovation. Experimentation is adopted by the agent to explore the state of the world. As the state of the world can either be good or bad, the agent exerts effort, which is equivalent of pulling the risky arm. Whereas, taking the safe arm, shirking or delaying effort, generates zero profit and is not preferable to the agent as actions are fully observable to the investor, and as the project owner, the fundamental interest of the agent is to generate the final breakthrough and get the revenue. The innovation project is assigned with a prior belief of the good state, denoted as p_0, which is public information. The agent will initiate innovation as long as the expected revenue exceeds the investment cost. The agent experiments and waits for the arrival of a good news indicating an innovation breakthrough, but the belief of the good state decreases as long as no good news comes. Due to the downward-adjust posterior belief, the agent would abandon the project when the expected revenue can not cover the cost, which forms a stopping time of the innovation.

The chapter first shows that the first-best financial contract involves two stopping times and a bonus paid from the agent to the investor at the completion of the experimentation stage, and the amount of the bonus depends on whether the development phases finishes early or late. Assume that in the baseline model, the completion of the first stage conveys no useful information on the state of the world. Thus, the experimentation in the second stage lasts the same length of time in both cases, and at the early-completion, the project would optimally be terminated at an earlier time than the late-completion. The optimal stopping time is determined when the marginal payoff of the project becomes zero, and the posterior belief of the good state reaches to a minimum cutoff. The bonus is paid to compensate the total investment cost from the investor given the completion of the first stage. Since the investor pays more investment costs when the first stages is completed late, hence a larger bonus should be paid compared with the early case. However, deviation from this first-best contract is possible when the completion of the first stage is not observable: the agent with the early-completion would want to hide and claim a late-completion. By doing so, he could obtain funding until the late-contract stopping time. Moreover, deviation of the agent will also result in an asymmetric belief adjusting such that the investor is more optimistic about the good state than the agent. Proposition 9 demonstrates an optimal separating contract which provides incentives for the agent to reveal the progress of the project truthfully and chooses either the early-completion or late-completion relevant contractual terms accordingly. The optimal contract indicates that the bonus of the late-completion should be larger than it of the early-completion and specifies a minimum distance between the two bonuses so as to prevent the agent from hiding the early-completion efficiently. Moreover, the first-best experimentation time can be achieved under this optimal contract, which means that the agent can still claim the entire surplus of the project and the investor break-even on average. We also consider an optimal pooling contract which is restricted with a single stopping time and a bonus exogenously. For example, the investor may prefer a more convenient and straightforward pooling contract as it may take longer to negotiate than a rather complicated separating contract. Proposition 10 proves that the first-best experimentation is not attainable under the optimal pooling contract: the project will be terminated earlier than the stopping time of the late-completion case, and the shortened experimentation time in the late-completion is to subsidise the potential welfare loss of the investor as a result of information imperfection. Thus, the agent without knowing the first stage will finish early or late ex ante would prefer the separating contract to the pooling contract.

The chapter then proceeds to the case when the completion time of the first stage is neither observable nor verifiable by the investor. Since it is very costly and difficult to monitor the innovation progress, the investor relies on the announcement or report from the agent to be informed of the progress of the innovation. However, we assume such reports cannot be verified by the investor. In this case, the agent may deviate from the optimal separating contract: when the first stage finishes late, he may lie and pretend to have an early-completion mainly due to the smaller bonus transfer specified in the early-contract. To solve this deviation problem, we impose another incentive compatibility constraint in the optimal separating contract and find out an upper bound on the difference between the early-bonus and late-bonus. As long as the difference between the two possible bonuses is below this upper bound, the agent will not deviate from this optimal separating equilibrium, and the first-best is still achievable; otherwise, if the bonus difference is not within the range, the agent will deviate.

As an extension, we discuss an information spillover effect across two stages. Suppose that the completion of the first stage is informative such that the early-completion indicates a higher prior belief of the good state than the late-completion. In this case, we want to explore whether this additional information embedded in the first stage would influence the way of long-term contracting and the welfare of both parties. The result in Proposition 12 demonstrates an optimal incentive-compatible contract under imperfect information which specifies two cases regarding the undetermined relationship between the first-best stopping times under the early-completion and late-completion. Due to the heterogeneity of the prior beliefs under the early-case and late-case, the experimentation required in the second stage differs: under the early-completion, the project with a higher prior deserves a longer experimentation than the late-completion. In this case, the relationship between the optimal stopping times depends on the values of exogenous prior beliefs and completion times of the first stage, hence cannot be determined. When the stopping time under the early-completion exceeds the late-completion, a larger bonus should be imposed in the early contract, which is opposite against the previous optimum. If the bonuses is confined within a certain range according to Proposition 12, the first-best is feasible. When the late-completion stopping time exceed the early, the optimal contract remains the same as in Proposition 11 but the distance between the two bonuses becomes smaller. As the early contract contains positive information of the good state, hence has a longer experimentation time than the late, the incentive of mimicking the early contract increases, which gives rise to a pressure of increasing the early-completion bonus. Conversely, there is a pressure of downward adjusting the late-completion bonus. Thus, the difference between the bonuses gets smaller.

4.1.2 Related literature

The chapter first contributes to the literature on the contracting of a multistage project. Within this strand of literature, Green and Taylor (2016) is one of the closest to ours where they build a two-stage-breakthrough model with unknown actions of the agent and imperfect information on the progress of the project. To resolve the imperfect information problem, the optimal contract proposed by the principal consists of a "soft" and "hard" deadline. The principal would expect a report of the innovation progress from the agent before this soft deadline, otherwise, the agents will be punished to undertake a probation, which involves a random probability of being fired. In this case, the agent is incentivized to work and truthfully report the progress of the project. While the hard deadline is to guarantee the completion of the innovation, which is comprised of two successes. As another similarity, our paper also focuses on the discussion of the imperfectly observed progress of the project. However, we consider the agent-optimal contact where the investor holds no bargaining power, it turns out that the optimal incentive compatible contracts for the agent to reveal the progress involve no distortion on the experimentation time. As a difference, moral hazard problem is not a major concern in this chapter. Moreover, according to our model, the arrival of the experimentation success is random and uncertain, which requires learning of agent. Whereas, the project can succeed with certainty such that there is no learning and belief updating in Green and Taylor (2016).

Information imperfection issue has been wildly discussed in the literature of innovation financing. The seminal papers by Bergemann and Hege (1998, 2005) study the dynamic moral hazard in the financing of innovation. In the presence of imperfect information on the actions of the agent, the agent may find it profitable to divert funding to private ends, which may cause an asymmetric belief updating between the investor and agent. Their papers show that the investor can shorten the funding period to prevent such deviation if renegotiation and short-term contract are allowed. Horner and Samuelson (2013), Demarzo and Sannikov (2016), He et al. (2017) recently approach to the incentivising problem in context of long-term financial contracting. They show that the equilibrium effort of the agent is induced to be front-loaded so as to achieve the maximum experimentation in the early stage. Moreover, the termination rule is set by the investor so as to improve the incentives of the agent. In our model, a similar asymmetric information problem arises where the investor can not perfectly observe the progress of the project and his belief might be distorted if the agent hides or lies about the innovation progress. However, our model differs from theirs in several aspects: first, our paper focuses on the imperfect information of the innovation progress and the agent's potential deviation of hiding progress from the investor, rather than shirking and delaying exert effort. Second, multistage innovation is investigated in this chapter where we concentrate on the optimal long-term contract that eliminates the information asymmetry issue.

According to another strand of the literature on the stage financing (Neher, 1999; Bergemann et al., 2010; Dahiya and Ray, 2012), agency problem resulted from imperfect monitoring of the investor can be resolved by providing investment funding sequentially and splitting the financing horizon into different stages. The releasing of funding controlled by the investor varies over time, which can be used to prevent the agent from procrastinate. Whereas, in my model, the investor has no power of the funding releasing speed or timing, and the innovation is restricted exogenously to be completed within two stages .

Under the similar multistage-experimentation-paradigm, Moroni (2016) and Wolf (2017) approach the moral hazard agency issue, so their focus is to design a contract to monitor and incentivise the agent to work, which is different from the imperfect information problem of the project progress addressed in this chapter. As the actions of the agents can not be observed by the principal, the agents may shirk and free-ride on the findings of the others. The optimal contract designed in Moroni (2016) shows that information rent should be paid to the agent who succeeds at early stage, hence the first-best experimentation can not be obtained. Rewards in the contract include the monetary payoff, a extended deadline in the later stage and the lessening of competition as some agents may be excluded by the principal. Wolf (2017) introduces the information spillover effect across two innovation stages, so the agents have additional information rent as the early stage failure may be resulted from the bad state of the world rather than shirking.

The remainder of the chapter is organised as follows: Section 4.2 introduces the model set-up and defines the first-best experimentation. Section 4.3 is the main body of this article, which begins with the optimal verifiable contract under imperfect information, including the discussions of separating and pooling contracts, follows with the optimal contract with non-verifiability. Section 4.4 extends the baseline model to the information spillover effect across stages. Section 4.5 concludes.

4.2 Model

4.2.1 Setup

There is an agent and an investor involved in a two-stage innovation financing game. The agent owns the innovation project. The outcome of the project depends on the unknown state of the world and is either good or bad. In the good state, the project will succeed and generate R amount of revenue with a random arrival rate; in the bad state, it will fail and produce nothing. The agent has a prior belief p_0 that the state is good, and needs to strategically conduct the experiment and discover the state of the world. The agent has no funding to run the project and requires a continuous amount i per unit of time of investment funding from the outside investor. For simplicity, both the agent and the investor have a zero discount rate[2].

The completion of the innovation project goes through two stages, a development stage and an experimentation stage. The first stage is relatively easy, and one can succeed with certainty, which is independent of the state of the world. In other words, the completion of the first stage does not provide any information about the state of the world. The completion of the first stage can only be observed by the agent. For simplicity, we assume that the first stage can be completed either early or late and follows a binary distribution: with probability α, the first stage finishes at time T_1^e and with probability $(1-\alpha)$, it is completed late at T_1^l, where $T_1^e < T_1^l$[3]. Thus, the agent knows whether the first stage finishes early or late at T_1^e. T_1^e and T_1^l are exogenously fixed parameters, and and whether the development stage is completed early or late is only observed by the agent and not the investor. Once the development stage is completed, the project proceeds to the second experimentation stage. In the second stage, a successful outcome of the project is a random variable that depends on whether the state of the world is good or bad. Assume that the arrival of a success in the second stage follows an exponential distribution with the intensity rate λ in the good state. Given the prior belief of the good state p_0, during the second experimentation stage, the belief that the state is good decreases as long as no success comes. With this decreasing posterior belief, the agent will optimally terminate the project when the expected marginal payoff becomes non-positive. Denote T_2^e and T_2^l as two termination times conditional on the early or late completion of the development stage respectively such that $T_2^e \leq T_2^l$.

Under the benchmark model, the prior belief of the state being good would not vary with the early or late completion of the first stage. In the extension part, we allow the completion of first stage to be informative such that an early completion indicates a higher prior of the good state. In this case, the agent would extend

[2]It can be generalised to the case with discounting but the optimal bonus contract of our paper still apply.

[3]The binary distribution of first-stage completion time could be generalised to more complicated distributions, such as exponential distribution, Brownian motion. However, using the binary distribution enables us to relate the incentive problem to just two stopping times, which avoids extra complications.

the deadline of the second stage under an early-completion. We will discuss the spillover effect and incentive problems when $T_2^e \leq T_2^l$ and $T_2^e > T_2^l$ respectively.

Since only the agent observes when the development stage is completed, the agent may have an incentive to delay announcing or lie about the completion of the first stage. The financial contract negotiated by the agent and the investor should take this incentive problem into account.

4.2.2 First-best stopping time

This section considers the optimal choice of a fully-informed social planner. The social planner bears the entire investment cost and decides the optimal termination times T_2^e and T_2^l conditioning on the end of the development phase that maximise the total expected surplus of the innovation project. Denote Π as the total expected payoff of the social planner:

$$\begin{aligned}\Pi = & -\alpha T_1^e i + \alpha \int_{T_1^e}^{T_2^e} (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0)(p_t \lambda R - i)\mathrm{d}t \\ & -(1-\alpha)T_1^l i + (1-\alpha)\int_{T_1^l}^{T_2^l} (p_0 e^{-(t-T_1^l)\lambda} + 1 - p_0)(p_t \lambda R - i)\mathrm{d}t,\end{aligned}$$

where the first two terms are the expected surplus under the early-completion case with the probability of α according to the binary distribution, and the last two terms are the expected surplus under the late completion. Note that because there is no direct benefit generated from the development stage, the social planner as the project owner will immediately commence the experimentation stage as soon as the development phase is completed. Using the properties of the exponential distribution, with arrival rate λ, the cumulative probability of success from T_1^e to time t is $(1 - e^{-\lambda(t-T_1^e)})$. $p_0 e^{-\lambda(t-T_1^e)}$ accounts for the probability of no arrival of the success up to t if the state is good and $1-p_0$ is the probability conditioning on the bad state, so the term $(p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0)$ is the unconditional probability of no success arrives before t. p_t is updated posterior belief of the state being good according to Bayesian updating rule, i.e. $p_t = \frac{p_0 e^{-\lambda(t-T_1^e)}}{p_0 e^{-\lambda(t-T_1^e)} + 1 - p_0}$. The analogous belief updating applies to the case when the first stage finishes late at T_1^l.

Π can be rearranged when substituting the posterior belief p_t as follows:

$$\begin{aligned}\Pi = & -\alpha T_1^e i \mathrm{d}t + \alpha \int_{T_1^e}^{T_2^e} p_0 e^{-(t-T_1^e)\lambda} \lambda R - (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0) i \mathrm{d}t \\ & - (1-\alpha) T_1^l i + (1-\alpha) \int_{T_1^l}^{T_2^l} p_0 e^{-(t-T_1^l)\lambda} \lambda R - (p_0 e^{-(t-T_1^l)\lambda} + 1 - p_0) i \mathrm{d}t.\end{aligned}$$

The optimal T_2^e and T_2^l are chosen such that Π is maximised according to the first-order condition as follows:

$$\begin{aligned}\frac{d\Pi^e}{dT_2^e} = & p_0 e^{-(T_2-T_1^e)\lambda} \lambda R - (p_0 e^{-(T_2-T_1^e)\lambda} + 1 - p_0) i = 0 \\ \Rightarrow \bar{T_2^e} = & T_1^e + \frac{1}{\lambda} \ln \left(\frac{p_0}{1-p_0} \frac{\lambda R - i}{i} \right); \\ \frac{d\Pi^l}{dT_2^l} = & p_0 e^{-(T_2-T_1^l)\lambda} \lambda R - (p_0 e^{-(T_2-T_1^l)\lambda} + 1 - p_0) i \\ \Rightarrow \bar{T_2^l} = & T_1^l + \frac{1}{\lambda} \ln \left(\frac{p_0}{1-p_0} \frac{\lambda R - i}{i} \right),\end{aligned}$$

where $\bar{T_2^e}$ and $\bar{T_2^l}$ denote the first-best experimentation times under early and late completion of the first stage respectively.

The first-best total experimentation time at the second stage is independent of the completion of the first stage as indicated by $\bar{T_2^l} - T_1^l = \bar{T_2^e} - T_1^e$. In other words, the result of the first stage does not convey any information on the state of the world, so the prior belief remains unchanged regardless of the completion time of the first stage. Moreover, the project is terminated when the posterior belief of the good state reaches the cutoff if no success comes such that $p_{\bar{T_2}^e} = p_{\bar{T_2}^l} = \frac{i}{\lambda R}$.

4.3 Financial contracting under imperfect information

4.3.1 Optimal separating contract with verifiability

As the project owner, the agent has no funding and seeks outside financing from the investor. This section considers the financial contracting between the investor and the agent when the completion of the first stage is privately known by the agent but the progress announcement can be verified by the investor. The aims of the financial contract are to optimise the expected payoff of the agent and make sure the investor participates and gains non-negative expected payoffs.

Moreover, as the project progress is unobervable but verifiable to the investor, the agent may have incentives to delay reporting the early completion, so the optimal contract is designed also to make the agent voluntarily announce the progress.

We first consider a separating contract which creates incentives for the agent to announce the completion of the first stage by choosing either early or late terms in the contract accordingly. The financial contract proposed specifies maximum funding periods from period 0 to T_2^e or T_2^l in the event of early or late completion of the first stage respectively and non-time varying lump-sum bonuses from the agent to investor in the experimentation stage conditional on the announcement of early or late completion of the development stage, denoted as B^e and B^l. That is, the contract is $C = \{T_2^e, B^e; T_2^l, B^l\}$. The bonuses B^e and B^l are used to compensate the funding provided by the investor and create creditable incentives for the agent to reveal the completion of the development stage. Specifically, if the first stage finishes early, the agent will inform the investor in order to get funded until T_2^e, and the investor receives B^e as soon as success in the experimentation stage arrives; otherwise, the contractual terms concerning the late completion will be triggered. Since the agent cannot self-fund, the bonuses are made to the investor only when the revenue R is realised at the experimentation success, whereas, in the event of no success in the bad state, zero bonus will be paid. We initially assume that these bonuses do not depend on the time of success at the experimentation stage and check whether such bonuses can be used to achieve the first-best outcome.

As a first step, consider the case of full information where the investor knows the time the development stage is completed. Assume the agent maximises his total expected payoff and subjects to offering the investor a break-even utility in both the early and late cases. Under $C = \{T_2^e, B^e; T_2^l, B^l\}$, let V^e and V^l be the total expected payoff of the investor under the early and late completion of the first stage, which are equal to zero at the agent's optimal contract. Thus, the first-best bonuses, also known as full-information bonuses, are denoted as $\bar{B}^e$ and

$\bar{B}^l$, which can be determined by setting V^e and V^l equal to zero such that:

$$
\begin{aligned}
V^e =& - \int_0^{T_1^e} i\mathrm{d}t + \int_{T_1^e}^{T_2^e} p_0 e^{-(t-T_1^e)\lambda} \lambda B^e - (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0) i\mathrm{d}t = 0 \\
\Rightarrow & \bar{B}^e \int_{T_1^e}^{T_2^e} p_0 e^{-(t-T_1^e)\lambda} \lambda = \int_0^{T_1^e} i\mathrm{d}t + \int_{T_1^e}^{T_2^e} (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0) i\mathrm{d}t; \\
V^l =& - \int_0^{T_1^l} i\mathrm{d}t + \int_{T_1^l}^{T_2^l} p_0 e^{-(t-T_1^l)\lambda} \lambda B^l - (p_0 e^{-(t-T_1^l)\lambda} + 1 - p_0) i\mathrm{d}t = 0 \\
\Rightarrow & \bar{B}^l \int_{T_1^l}^{T_2^l} p_0 e^{-(t-T_1^l)\lambda} \lambda = \int_0^{T_1^l} i\mathrm{d}t + \int_{T_1^l}^{T_2^l} (p_0 e^{-(t-T_1^l)\lambda} + 1 - p_0) i\mathrm{d}t.
\end{aligned}
$$

Thus, the bonuses compensate the total investment cost of the investor under full-information, and $\bar{B}^l$ exceeds $\bar{B}^e$ by the first stage extra investment time from T_1^e to T_1^l.

Now, we consider the case where the completion of the development stage is private information of the agent, but where completion can be costlessly verified by the investor. Thus, the agent cannot claim to have completed early if she has not, but the agent may be able to claim to have completed late even if in fact she completed earlier. The advantage of claiming a late completion of the development stage is that the agent might receive extended funding for experimentation until T_2^l. Such a behaviour would distort upward the posterior belief of the investor, where the investor always holds higher posterior than the agent during T_1^l to T_2^l, which would altogether result in a welfare loss to the investor. Thus, we want to find an incentive compatible contract that gives the agent the incentive to report truthfully.

Let U_0 and V_0 denote the total expected payoff of the agent and the investor at period 0 respectively as follows:

$$
U_0 = - \alpha \int_{T_1^e}^{T_2^e} p_0 e^{-(t-T_1^e)\lambda} \lambda (R - B^e) \mathrm{d}t + (1-\alpha) \int_{T_1^l}^{T_2^l} p_0 e^{-(t-T_1^l)\lambda} \lambda (R - B^l) \mathrm{d}t, \tag{4.1}
$$

$$
\begin{aligned}
V_0 =& - \alpha \int_0^{T_1^e} i\mathrm{d}t + \alpha \int_{T_1^e}^{T_2^e} \{p_0 e^{-(t-T_1^e)\lambda} \lambda B^e - (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0) i\} \mathrm{d}t \\
& - (1-\alpha) \int_0^{T_1^l} i\mathrm{d}t + (1-\alpha) \int_{T_1^l}^{T_2^l} p_0 e^{-(t-T_1^l)\lambda} \lambda B^l - (p_0 e^{-(t-T_1^l)\lambda} + 1 - p_0) i\mathrm{d}t,
\end{aligned} \tag{4.2}
$$

where the first and third terms of V_0 are the total investment funding up to the early or late completion of the first stage. The agent optimally chooses stopping times and compensation subject to the break-even (or individual rationality, or participation constraint) of the investor and incentive compatibility constraints as follows:

$$\begin{aligned} \max_{T_2^e, B^e, T_2^l, B^l} \quad & U_0(C) \\ \text{subject to} \quad & U^e(T_2^e, B^e) \geq U^e(T_2^l, B^l) && (IC^e) \\ & V_0(C) \geq 0, && (IR) \end{aligned}$$

where $U^e(T_2^e, B^e)$ is the expected payoff of taking the early contract conditioning on the early completion, $U^e(T_2^l, B^l)$ is the expected payoff of hiding the early completion and choosing the late contract[4]. IC^e ensures the agent truthfully reveal the early completion of the first stage. IR is the individual rationality constraint of the investor. However, there is no analogous IC^l which ensures the truthful reveal of the late completion as the progress of the first stage is verifiable and the agent can not pretend to complete early in the late case. Let the IR constraint bind at the optimum, so set $V_0 = 0$, and the programme becomes:

$$\begin{aligned} \max_{T_2^e, T_2^l} \quad U_0 =& \alpha \int_{T_1^e}^{T_2^e} (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0)(p_t \lambda R - i) \mathrm{d}t \\ & + (1-\alpha) \int_{T_1^l}^{T_2^l} (p_0 e^{-(t-T_1^l)\lambda} + 1 - p_0)(p_t \lambda R - i) \mathrm{d}t. \end{aligned}$$

Note that the IC^e constraint is ignored at the moment, and we would then check at the optimal experimentation time, whether the full information bonuses could be satisfied. Taking the first-order condition of U_0 with respect to T_2^e and T_2^l:

$$\begin{aligned} \frac{dU_0}{dT_2^e} =& p_0 e^{-(T_2^e - T_1^e)\lambda} \lambda R - (p_0 e^{-(T_2^e - T_1^e)\lambda} + 1 - p_0) i = 0, \\ \frac{dU_0}{dT_2^l} =& p_0 e^{-(T_2^l - T_1^l)\lambda} \lambda R - (p_0 e^{-(T_2^l - T_1^l)\lambda} + 1 - p_0) i = 0. \\ \Rightarrow \bar{T_2}^e =& T_1^e + \frac{1}{\lambda} \ln\left(\frac{p_0}{1-p_0} \frac{\lambda R - i}{i} \right); && (4.3) \\ \bar{T_2}^l =& T_1^l + \frac{1}{\lambda} \ln\left(\frac{p_0}{1-p_0} \frac{\lambda R - i}{i} \right), && (4.4) \end{aligned}$$

[4]Here we assume that the agent can conceal the experimentation success when it arrives between T_1^e and T_1^l. In this case, the incentive problem still exists in the sense that the agent may claim a late completion even when the experimentation success comes before which he chooses the late contract at T_1^l.

which are the first-best experimentation times.

Then we consider the choices of B^e and B^l that satisfying the IC^e constraint under the first-best experimentation times, which incentivise the agent reveal the early completion.

Proposition 9 *The optimal separating equilibrium financial contract under imperfect information on the first stage is* $C = \{\bar{T_2}^e, B^e; \bar{T}_2^l, B^l\}$*, where* $\bar{T_2}^e$ *and* $\bar{T}_2^l$ *are the first-best maximum funding time,* B^e *and* B^l *are the bonuses distributed by the agent when success comes at the second stage based on the early or late completion of the first stage respectively. Specifically, the expectation of* B^e *and* B^l *should satisfy the following relationship:*

$$\int_{T_1^l}^{\bar{T_2}^l} p_{T_1^l} e^{-(t-T_1^l)\lambda} \lambda B^l \mathrm{d}t - \int_{T_1^e}^{\bar{T_2}^e} p_0 e^{-(t-T_1^e)\lambda} \lambda B^e \mathrm{d}t \geq R\gamma^e$$

where $\gamma^e = \int_{\bar{T_2}^e}^{\bar{T_2}^l} p_{\bar{T_2}^e} e^{-(t-\bar{T_2}^e)\lambda} \lambda \mathrm{d}t$, $p^e_{\bar{T_2}^e} = \frac{p_0 e^{-\lambda(\bar{T_2}^e - T_1^e)}}{p_0 e^{-\lambda(\bar{T_2}^e - T_1^e)} + 1 - p_0}$*. Additionally, the full information bonuses satisfy the above inequality, which means the full-information equilibrium* $\{\bar{T_2}^e, \bar{B}^e; \bar{T}_2^l, \bar{B}^l\}$ *is achievable.*

Proof. See appendix. ■

The above optimal separating contract provides sufficient incentives for the agent to truthfully reveal the completion of the first stage and take the early or late terms accordingly, and the investor in expectation breaks even. The proposition indicates that the gap between the bonuses in the early and late terms should be at least large enough to cover the expected net gain of the agent if the late contract is taken at the early completion, which is exactly the expected revenue in the extended experimentation time from $\bar{T_2}^e$ to $\bar{T_2}^l$. Intuitively, by hiding the early completion the agent can distort the investor's posterior belief of the good state in which case the investor holds higher posterior belief over time than the agent such that $\forall t \in [T_1^l, \bar{T_2}^l]$, $p_t(investor|agent\ hides) = \frac{p_0 e^{-\lambda(t-T_1^l)}}{p_0 e^{-\lambda(t-T_1^l)} + 1 - p_0} > p_t(agent) = \frac{p_0 e^{-\lambda(t-T_1^e)}}{p_0 e^{-\lambda(t-T_1^e)} + 1 - p_0}$. Thus, the agent continuously gains positive profit until $\bar{T_2}^l$, which makes a loss to the investor. To deter such a deviation, we impose the incentive compatibility constraint and specify different bonuses be distributed in the event of early or late completion. Specifically, a smaller bonus in case of early completion should be claimed by the investor in order to eliminate agent's hiding incentive, and the difference of the expected bonuses on the late and early cases should be at least large enough to cover the net gain of the agent if the early completion is hidden. Moreover, the separating equilibrium contract can fully induce the first-best outcome regardless of the information imperfection in the sense that the first-best experimentation times and bonuses can be obtained in both early and late cases and the agent claims the entire surplus. In other words, the first-best contract forms a separation equilibrium in both early and late cases where the agent is incentivised to truthfully report the progress of the first stage.

4.3.2 Optimal pooling contract with verifiability

In practice, it can be case when the contract is exogenously restricted to a single experimentation deadline and a threshold bonus. Under such restriction, we want to discuss the existence of an optimal pooling contract when the first stage process is imperfect information and compare it with the optimal separating contract in terms of the welfare of both the agent and investor.

Proposition 10 *When the first stage completion is private information of the agent, and the contract is restricted to a single-time and threshold-bonus form, the optimal pooling contract is $C = \{B, \tilde{T}_2\}$, where the investor starts financing until $\tilde{T}_2$ regardless the completion time of the first stage such that:*

$$\tilde{T}_2 = \frac{1}{\lambda}\ln\left(\frac{p_0}{\alpha(1-p_0)+(1-\alpha)(\frac{p_0}{\alpha p_{T_1^l}+(1-\alpha)p_0}-p_0)}\frac{\lambda R - i}{i}[\alpha e^{T_1^e\lambda}+(1-\alpha)e^{T_1^l\lambda}]\right),$$

which is smaller than $\bar{T}_2^{\,l}$.

Proof. See appendix. ■

The above optimal pooling contract states that under either early or late completion, the funding will be provided until $\tilde{T}_2$, which is smaller than the first-best experimentation time for late completion under the separating contract. Intuitively, the shortened optimal experimentation time is to subsidise the welfare loss of the investor when the first completion is not observable. Such welfare loss stems from two aspects: the difference in the optimal stopping times in the early-completion and the late-completion; the asymmetric posterior belief updating where the investor holds higher posterior belief than the agent if the early-completion is not announced. Under the first-best, the late completion will be

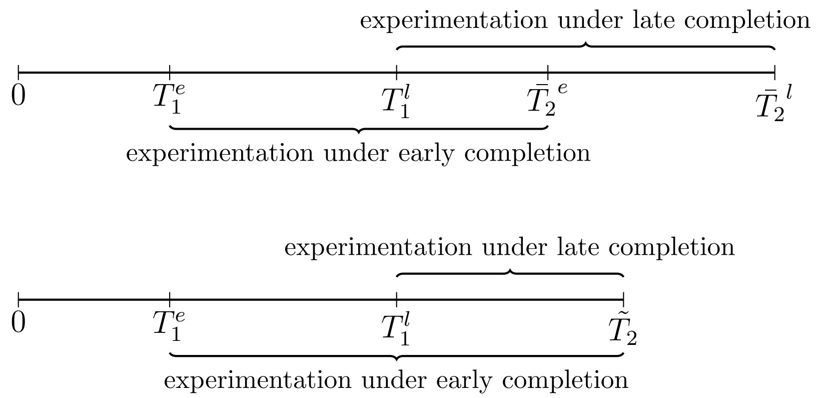

Figure 4.1: The optimal separating (above) and pooling (below) contracts

granted funding until $\bar{T}_2{}^l$, which is longer than the early-completion-first-best $\bar{T}_2{}^e$. Due to the restriction of an single funding deadline, $\tilde{T}_2$ is chosen such that the investor breaks-even in expectation no matter when the first stage is completed. Moreover, conditioning on the early completion of the first stage, the second experimentation stage starts at T_1^e, which is unknown to the investor. In this case, as long as the success does not come until T_1^l, the agent privately updates the posterior belief of the good state being $p_{T_1^l}$. Whereas, at T_1^l the investor forms the expected beliefs in the good state such that with α probability the belief is $p_{T_1^l}$ where the agent hides the early completion from her, and with $(1-\alpha)$ probability, the belief is p_0 where no experimentation has been conducted so far. Thus, the expected belief of the good state of the investor at T_1^l is $\Big(\alpha p_{T_1^l} + (1-\alpha)p_0\Big)$, such that $\Big(\alpha p_{T_1^l} + (1-\alpha)p_0\Big) > p_{T_1^l}$, $\forall \alpha \in (0,1)$. Thus, the investor is more optimistic than the agent as his posterior belief is distorted upward. Such asymmetric belief updating will not cease unless the success arrives and is observed by the investor at the experimentation stage.

Under the optimal pooling contract, the agent is not incentivised to report the early-completion as the contract is restricted to a single stopping time and bonus. As assumed that both parties have full commitment power once the contract is signed, even when the early-completion is revealed, the agent has the right to continue requesting funding from the investor until $\tilde{T}_2$[5].

[5]However, when renegotiation is allowed, the investor would have incentive to provide the agent additional benefit to reveal the early-completion of the first stage and stop at an earlier time, which would improve welfare for both parties.

Figure 4.1 shows the timings of the optimal separating and pooling contract. Under the separating contract, a longer experimentation time can be achieved if the first stage is completed late as $\tilde{T}_2 < \bar{T}_2^{\,l}$ is proven in Proposition 10.

Result 2 *The optimal separating contract dominates the optimal pooling contract where the agent obtains a longer experimentation time and captures the entire surplus in the former contract, and the investor breaks even in both contracts. Thus, the rational agent would always propose the separating contract.*

4.3.3 Optimal contract with non-verifiability

This section considers the problem of imperfect information where the completion of the first stage can neither be observed nor verified by the investor. Although the agent can not proceed to the experimentation stage until she finishes the development stage, she could claim an early completion even when she finishes late without being found out by the investor. By lying about the progress and taking the early contract, the agent would only get funded until $\bar{T}_2^{\,e}$ but pay less bonus to the investor in the event of success at the second stage under the previous optimal separating contract if no additional incentive constraint is imposed. Moreover, there will be no bonus transfer during T_1^e and T_1^l since the second experimentation stage only initiates at T_1^l.

Let $U^l(T^e, B^e)$ be the expected payoff of the agent when the late completion of the first stage is claimed to be early:

$$U^l(T_2^e, B^e) = \int_{T_1^l}^{T_2^e} p_0 e^{-(t-T_1^l)\lambda} \lambda (R - B^e) \mathrm{d}t.$$

In this case, the first-best experimentation time is not achieved since there is no funding support during T_2^e to T_2^l. However, the agent could still benefit from such deviation if the difference between expected bonus between the late and early cases is larger than the expected revenue loss in $[T_2^e, T_2^l]$. Thus, the optimal contract needs to eliminate such distortion when the progress is non-verifiable, realise the maximised profit of the agent and make sure non-negative expected payoff of the investor.

The optimisation programme of the agent is as follows:

$$\begin{aligned}
\max_{T_2^e, B^e, T_2^l, B^l} \quad & U_0(C) & \\
\text{subject to} \quad & U^e(T_2^e, B^e) \geq U^e(T_2^l, B^l) & (IC^e) \\
& U^l(T_2^l, B^l) \geq U^l(T_2^e, B^e) & (IC^l) \\
& V_0(C) \geq 0, & (IR)
\end{aligned}$$

where we add the IC^l constraint to eliminate the agent's incentive of lying about the progress.

Following the same logic dealing with the optimisation problem in Section 4.3.1, we first want to check if the first-best experimentation times and bonuses are attainable under a separating contract where the agent has incentives to reveal the progress truthfully. Based on the previous analysis, Proposition 9 indicates that a larger expected bonus should be transferred to the investor in the late-completion case to prevent the agent from hiding the early completion. However, the larger the difference between B^l and B^e, the stronger the incentive to mimic the early-completion for the agent. Thus, with the extra incentive compatibility constraint IC^l, there are additional restrictions on the relationship between the bonuses B^l and B^e that need to be checked.

Proposition 11 *The first-best experimentation time can be achieved under the optimal separating contract* $C = \{\bar{T_2}^e, B^e; \bar{T}_2^l, B^l\}$, *where* B^l *and* B^e *should satisfy the following relationship:*

$$R\gamma^e < \int_{T_1^l}^{\bar{T}_2^l} p_0 e^{-(t-T_1^l)\lambda} \lambda B^l \mathrm{d}t - \int_{T_1^l}^{\bar{T}_2^e} p_0 e^{-(t-T_1^l)\lambda} \lambda B^e \mathrm{d}t \leq R\gamma^l,$$

where $\gamma^l = \int_{\bar{T}_2^e}^{\bar{T}_2^l} p^l_{\bar{T}_2^e} e^{-(t-\bar{T}_2^e)\lambda} \lambda \mathrm{d}t$, $\gamma^e = \int_{\bar{T_2}^e}^{\bar{T_2}^l} p^e_{\bar{T_2}^e} e^{-(t-\bar{T_2}^e)\lambda} \lambda \mathrm{d}t$, $p^e_{\bar{T_2}^e} = \frac{p_0 e^{-\lambda(\bar{T_2}^e - T_1^e)}}{p_0 e^{-\lambda(\bar{T_2}^e - T_1^e)} + 1 - p_0}$ *and* $p^l_{\bar{T_2}^e} = \frac{p_0 e^{-\lambda(\bar{T}_2^e - T_1^l)}}{p_0 e^{-\lambda(\bar{T}_2^e - T_1^l)} + 1 - p_0}$. *Moreover, the full-information bonuses satisfy the above inequality, which means the full-information equilibrium* $\{\bar{T_2}^e, \bar{B}^e; \bar{T}_2^l, \bar{B}^l\}$ *is achievable even when the completion of the first stage is non-verifiable.*

The above proposition indicates that the full information equilibrium contract could attain a complete separation between the agent with early or late completion of the first stage such that there is no incentive for the agent to deviate from this contract even with the unobervablity and non-verifiability of the project progress. The non-verifiablity adds an additional constraint on the bonuses where there is an upper bound of the difference between the late-bonus and the early-bonus. Intuitively, by lying about the early contract, the agent gives up the experimentation during $\bar{T}_2^e$ and $\bar{T}_2^l$, so $R\gamma^l$ denotes the loss of the agent. However, the difference between $\bar{B}^l$ and $\bar{B}^e$ is the gain of deviating from the late contract such that less bonuses is needed under the early contract. Thus, to deter the agent from lying, Proposition 11 indicates that the gain of deviating should be no greater than the loss.

4.4 Informative first-stage completion

This section extends the baseline financial contracting model to the case when the first stage contains additional information on the unknown state of the world such that the early completion indicates a higher prior belief of the good state than the late completion. In this case, the early completion will be granted a longer experimentation time due to the higher prior belief. When the innovation progress remains unobservable to the investor, the agent would have stronger incentive to mimic the early contract when she finishes late, whereas, the incentive of mimicking the late contract under the early-completion decreases. As the completion of the first stage contains information about the state of the world but is privately observed by the agent, it is not obvious whether the first-best is achievable or not under an optimal separating contract.

Let p_0^e be the prior belief of the good state when the first stage is completed early, and p_0^l corresponds to the late-completion-prior-belief where $p_0^e > p_0^l$ holds. The investor will hold either high or low prior beliefs of the project once the completion time of the first stage is announced by the innovator. The first-best stopping times can be solved by substituting p_0^e and p_0^l into the optimisation programme of the agent subject to the individual rationality constraint of the investor such that:

$$\bar{T_2}^e = T_1^e + \frac{1}{\lambda} \ln \left(\frac{p_0^e}{1 - p_0^e} \frac{\lambda R - i}{i} \right);$$
$$\bar{T_2}^l = T_1^l + \frac{1}{\lambda} \ln \left(\frac{p_0^l}{1 - p_0^l} \frac{\lambda R - i}{i} \right),$$

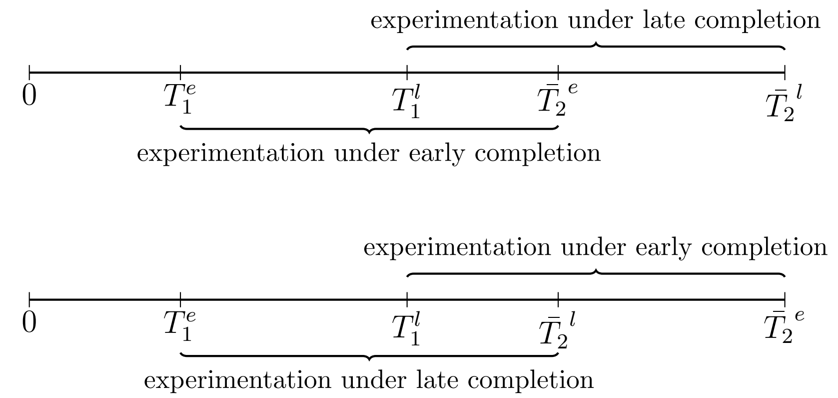

Figure 4.2: Two cases of the first-best stopping times when $\bar{T_2}^e < \bar{T_2}^l$ and $\bar{T_2}^e > \bar{T_2}^l$

which indicates that the early-completion requires longer experimentation than the late-completion as $\bar{T_2}^e - T_1^e > \bar{T_2}^l - T_1^l$. However, the relationship between $\bar{T_2}^e$ and $\bar{T_2}^l$ is undetermined as it depends on the value of exogenous parameters p_0^e, p_0^l, T_1^e and T_1^l, which means the first-best stopping time of the early-completion may fall behind or exceed the late-completion-stopping-time (as shown in Figure 4.2). The first-best bonuses could be solved by setting the individual participation constraint of the investor binding such that:

$$\bar{B}^e \int_{T_1^e}^{\bar{T}_2^e} p_0^e e^{-(t-T_1^e)\lambda}\lambda = \int_0^{T_1^e} i\mathrm{d}t + \int_{T_1^e}^{\bar{T_2}^e} (p_0^e e^{-(t-T_1^e)\lambda} + 1 - p_0^e) i\mathrm{d}t;$$
$$\bar{B}^l \int_{T_1^l}^{\bar{T}_2^l} p_0^l e^{-(t-T_1^l)\lambda}\lambda = \int_0^{T_1^l} i\mathrm{d}t + \int_{T_1^l}^{\bar{T}_2^l} (p_0^l e^{-(t-T_1^l)\lambda} + 1 - p_0^l) i\mathrm{d}t.$$

The relationship between the bonuses depends on the first-best experimentation times $\bar{T_2}^e$ and $\bar{T_2}^l$. As the bonuses are used to compensate the total investment cost of the investor, if $\bar{T_2}^e > \bar{T_2}^l$, the bonus on the early completion will be larger, otherwise, early-completion-bonus is smaller.

Under imperfect information on the project progress, we want to explore whether the first-best outcome is achievable in the sense that $C = \{\bar{T_2}^e, \bar{B}^e; \bar{T}_2^l, \bar{B}^l\}$ could induce a complete separation between the early and late cases. Thus, we impose the same optimisation programme of the agent as in the previous section, which includes the IC^l, IC^e and IR constraint.

Proposition 12 *The first-best experimentation time can be achieved under the optimal separating contract $C = \{\bar{T_2}^e, B^e; \bar{T}_2^l, B^l\}$, where B^l and B^e should satisfy*

the following relationship:
When $\bar{T_2}^e \geq \bar{T_2}^l$,

$$\int_{T_1^e}^{\bar{T_2}^e} p_0^e e^{-(t-T_1^e)\lambda}\lambda B^e \mathrm{d}t - \int_{T_1^l}^{\bar{T_2}^l} p_{T_1^l}^e e^{-(t-T_1^l)\lambda}\lambda B^l \mathrm{d}t > \int_{\bar{T_2}^l}^{\bar{T_2}^e} e^{-(t-\bar{T_2}^l)\lambda} i \mathrm{d}t \leq \int_{\bar{T_2}^l}^{\bar{T_2}^e} p_{T_2^l}^e e^{-(t-T_2^l)\lambda}\lambda R \mathrm{d}t; \tag{4.5}$$

When $\bar{T_2}^e < \bar{T_2}^l$,

$$\int_{T_1^l}^{\bar{T}_2^l} p_0^l e^{-(t-T_1^l)\lambda}\lambda B^l \mathrm{d}t - \int_{T_1^l}^{\bar{T}_2^e} p_0^l e^{-(t-T_1^l)\lambda}\lambda B^e \mathrm{d}t \in (R\gamma^e, R\gamma^l], \tag{4.6}$$

which is the same as in Proposition 11. Moreover, the full-information bonuses satisfy in both scenarios the above inequality, which means the full-information equilibrium $\{\bar{T_2}^e, \bar{B}^e; \bar{T}_2^l, \bar{B}^l\}$ *is achievable.*

Proof. See appendix. ■

The above proposition indicates that the relationship between the optimal bonuses and experimentation times under the early-completion and late-completion depends largely on the difference between the prior beliefs p_0^e and p_0^l. If the early-completion conveys an exceptionally positive news on the state being good such that at T_1^l the agent with the early-completion is still more optimistic of the good state than the late-completion, i.e. $p_{T_1^l}^e > p_0^l$ holds, the optimal experimentation will last longer in the early-completion scenario. In this case, a larger bonus should be paid to compensate the investor as the agent would have stronger incentive to take the early contract. The inequality Equation 4.5 shows the range of the gap between B^e and B^l within which the agent would truthfully reveal the progress and the full separation could be obtained. Specifically, by taking the early contract, the late-completion-agent would potentially gain $\int_{\bar{T_2}^l}^{\bar{T_2}^e} e^{-(t-\bar{T_2}^l)\lambda} i dt$, which is the expected payoff in the extended experimentation time. Moreover, the potential benefits of taking the late-contract in the early-completion case would be upward distorting the posterior belief of investor and undertaking no bonus transfer before T_1^l when the late-contract is triggered if the experimentation success can somehow be hidden by the agent. Thus, to deter the mimicking of the late-contract, the bonus gap should be smaller than the benefit of continuing experimentation until $\bar{T_2}^e$, which is $\int_{\bar{T_2}^l}^{\bar{T_2}^e} p^e_{T_2{}^l} e^{-(t-T_2{}^l)\lambda} \lambda R \mathrm{d}t$.

In the case when $\bar{T_2}^e < \bar{T_2}^l$, the relationship of B^e and B^l follows from Proposition 11 as the experimentation stopping time in the late-completion exceeds it in the early-completion. However, the distance between the bonuses becomes smaller. Due to the information spillover effect, the early contract has a longer experimentation time than the late contract, so the agent would have stronger incentive to mimicking an early-completion. To deter them from mimicking, the early-completion bonus increases. Conversely, the bonuses on the late-completion is adjusted downward due the reducing incentive of lying. Although in the case when $\bar{T_2}^e < \bar{T_2}^l$, the bonus of the late-completion still dominates it of the early-completion, the difference between the two bonuses gets smaller.

4.5 Conclusion

In conclusion, this chapter has studied a multistage innovation financing problem in which the progress of the project is neither observable nor verifiable by the investor. The chapter has shown that the optimal long-term separating contract consists of differential maximum funding periods in the event of early and late completion of the first stage respectively and subsequent bonuses from the agent to the investor conditioning on success of the experimentation stage. It has been shown that the first-best experimentation time is attainable as long as the bonus of the late completion exceeds that of the early completion, and the difference between the bonuses should be confined within a specific range. In the extension, there exists spillover effect over the stages, bonuses should be adjusted in accordance with the agent's differentiate distortion incentives. In particular, when the deadline of the early-completion exceeds the late-completion, the bonus on the early-completion would increase in order to balance the increased incentive of mimicking the early-completion. This chapter extends the study of long-term innovation financing by including the imperfect information on project progress, and enriches the theoretical bandit game by incorporating the spillover effect in a multistage experimentation.

Although this chapter has not considered heterogeneous innovators, a useful extension for future research could be a model with two types of agents with differentiate learning abilities. Private information on the innovation progress can be an indication of agent's ability such that the agent who finishes the first-stage early is a high-type or fast-learner with a larger arrival rate, i.e. $\lambda^e > \lambda^l$. The result may vary from current chapter as λ is non-monotonic in the stopping time.

Appendix

4.A Proof of Proposition 9

Let $U^e(T_2^l, B^l)$ be the expected payoff of the agent when hiding the early completion from the investor:

$$U^e(T_2^l, B^l) = \int_{T_1^e}^{T_1^l} p_0 e^{-(t-T_1^e)\lambda} \lambda R \mathrm{d}t + \int_{T_1^l}^{{T_2}^l} p_{T_1^l} e^{-(t-T_1^l)\lambda} \lambda (R - B^l) \mathrm{d}t$$
$$= \int_{T_1^e}^{{T_2}^l} p_0 e^{-(t-T_1^e)\lambda} \lambda R \mathrm{d}t - \int_{T_1^l}^{{T_2}^l} p_{T_1^l} e^{-(t-T_1^l)\lambda} B^l \mathrm{d}t$$

According to IC^e, $U^e({T_2}^e, B^e) \geq U^e({T_2}^l, B^l)$:

$$\int_{T_1^e}^{{T_2}^e} p_0 e^{-(t-T_1^e)\lambda} \lambda (R - B^e) \mathrm{d}t \geq \int_{T_1^e}^{{T_2}^l} p_0 e^{-(t-T_1^e)\lambda} \lambda R \mathrm{d}t - \int_{T_1^l}^{{T_2}^l} p_{T_1^l} e^{-(t-T_1^l)\lambda} B^l \mathrm{d}t$$

Make some rearrangement:

$$\int_{T_1^l}^{{T_2}^l} p_{T_1^l} e^{-(t-T_1^l)\lambda} \lambda B^l \mathrm{d}t - \int_{T_1^e}^{{T_2}^e} p_0 e^{-(t-T_1^e)\lambda} \lambda B^e \mathrm{d}t \geq \int_{T_2^e}^{T_2^l} p_{{T_2}^e} e^{-(t-\bar{T_2}^e)\lambda} \lambda R \mathrm{d}t, \tag{4.7}$$

where $p_{T_1^l} = \frac{p_0 e^{-\lambda({T_1}^l - T_1^e)}}{p_0 e^{-\lambda({T_1}^l - T_1^e)} + 1 - p_0}$ and $p_{{T_2}^e} = \frac{p_0 e^{-\lambda({T_2}^e - T_1^e)}}{p_0 e^{-\lambda({T_2}^e - T_1^e)} + 1 - p_0}$. As $p_0 > p_{T_1^l}$, the above inequality indicates the following:

$$\int_{T_1^l}^{{T_2}^l} p_0 e^{-(t-T_1^l)\lambda} \lambda B^l \mathrm{d}t - \int_{T_1^e}^{{T_2}^e} p_0 e^{-(t-T_1^e)\lambda} \lambda B^e \mathrm{d}t > \int_{{T_2}^e}^{{T_2}^l} p_{\bar{T_2}^e} e^{-(t-{T_2}^e)\lambda} \lambda R \mathrm{d}t. \tag{4.8}$$

Thus, the first-best experimentation times $\bar{T_2}^l$ and $\bar{T_2}^e$ can be chosen as long as the above inequality on the bonuses holds.

Then we want to check whether the first-best bonuses satisfy the above inequality. Under perfect information, the bonuses are just used to compensate the total expected investment cost by the investor, which means:

$$\bar{B}^e \int_{T_1^e}^{\bar{T}_2^e} p_0 e^{-(t-T_1^e)\lambda}\lambda = \int_0^{T_1^e} i\mathrm{d}t + \int_{T_1^e}^{\bar{T}_2{}^e} (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0) i\mathrm{d}t$$
$$\bar{B}^l \int_{T_1^l}^{\bar{T}_2^l} p_0 e^{-(t-T_1^l)\lambda}\lambda = \int_0^{T_1^l} i\mathrm{d}t + \int_{T_1^l}^{\bar{T}_2^l} (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0) i\mathrm{d}t$$

As $\bar{T}_2{}^l - T_1^l = \bar{T}_2{}^e - T_1^e$, $\bar{B}^l \int_{T_1^l}^{\bar{T}_2^l} p_0 e^{-(t-T_1^l)\lambda}\lambda - \bar{B}^e \int_{T_1^e}^{\bar{T}_2^e} p_0 e^{-(t-T_1^e)\lambda}\lambda = (T_1^l - T_1^e)i$. Substitute the full information bonuses into the left-hand-side of inequality Equation 4.8 and check whether the inequality still holds, which is equivalent of comparing $(T_1^l - T_1^e)i$ and $\int_{\bar{T}_2{}^e}^{\bar{T}_2{}^l} p_{\bar{T}_2{}^e} e^{-(t-\bar{T}_2{}^e)\lambda}\lambda R\mathrm{d}t$:

$$(T_1^l - T_1^e)i = \int_{\bar{T}_2{}^e}^{\bar{T}_2{}^l} i\mathrm{d}t > \int_{\bar{T}_2{}^e}^{\bar{T}_2{}^l} p_{\bar{T}_2{}^e} e^{-(t-\bar{T}_2{}^e)\lambda}\lambda R\mathrm{d}t$$

As at T_2^e, the early agent would terminate the project as the expected marginal payoff becomes non-positive such that $p_{\bar{T}_2{}^e}\lambda R - i = 0$.

4.B Proof of Proposition 10

The expected payoff of the agent and the investor under the contract $C = \{B, T_2\}$ are as follows:

$$\begin{aligned}
U_0 =& \alpha \int_{T_1^e}^{T_2} p_0 e^{-(t-T_1^e)\lambda}\lambda(R-B)\mathrm{d}t + (1-\alpha)\int_{T_1^l}^{T_2} p_0 e^{-(t-T_1)\lambda}\lambda(R-B)\mathrm{d}t \\
V_0 =& \alpha \int_{T_1^e}^{T_2} \{p_0 e^{-(t-T_1^e)\lambda}\lambda B - (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0)i\}\mathrm{d}t \\
&+ (1-\alpha)\int_{T_1^l}^{T_2} \{(\alpha p_{T_1^l} + (1-\alpha)p_0)e^{-(t-T_1^l)\lambda}\lambda B \\
&- \Big((\alpha p_{T_1^l} + (1-\alpha)p_0)e^{-(t-T_1^l)\lambda} + 1 - (\alpha p_{T_1^l} + (1-\alpha)p_0)\Big) i\}\mathrm{d}t,
\end{aligned}$$

where the first and second terms of V_0 are the investor's expected payoff when the first stage finishes early and late respectively. Specifically, $(\alpha p_{T_1^l} + (1-\alpha)p_0)$ is the expected belief of the investor when informed of the late completion of the first stage at T_1^l. The agent would optimally choose an T_2 that maximise his expected profit and make sure the investor get at least zero expected payoff. Thus, the optimisation programme of the agent is as follows:

$$\max_{T_2, B} U_0(C), \quad \text{subject to } V_0(C) \geq 0$$

Let the individual rationality constraint of the investor bind and substitute the expected bonus into U_0:

$$U_0 = \alpha \int_{T_1^e}^{T_2} \{p_0 e^{-(t-T_1^e)\lambda}\lambda R - (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0)i\}\mathrm{d}t$$
$$+ (1-\alpha)\int_{T_1^l}^{T_2} p_0 e^{-(t-T_1^l)\lambda}\lambda R \mathrm{d}t - (1-\alpha)\frac{p_0}{\alpha p_{T_1^l} + (1-\alpha)p_0} i \int_{T_1^l}^{T_2}$$
$$\left((\alpha p_{T_1^l} + (1-\alpha)p_0)e^{-(t-T_1^l)\lambda} + 1 - (\alpha p_{T_1^l} + (1-\alpha)p_0)\right)\mathrm{d}t$$

Taking the first-order condition with respect to T_2:

$$\frac{dU_0}{dT_2} = \alpha\left(p_0 e^{-(T_2-T_1^e)\lambda}\lambda R - (p_0 e^{-(T_2-T_1^e)\lambda} + 1 - p_0)i\right) + (1-\alpha)p_0 e^{-(T_2-T_1^l)\lambda}\lambda R$$
$$- (1-\alpha)i\frac{p_0}{\alpha p_{T_1^l} + (1-\alpha)p_0}[(\alpha p_{T_1^l} + (1-\alpha)p_0)e^{-(T_2-T_1^l)\lambda} + 1 - (\alpha p_{T_1^l} + (1-\alpha)p_0)]$$

$$\Rightarrow \tilde{T}_2 = \frac{1}{\lambda}\ln\frac{p_0}{\alpha(1-p_0) + (1-\alpha)(\frac{p_0}{\alpha p_{T_1^l} + (1-\alpha)p_0} - p_0)}\frac{\lambda R - i}{i}[\alpha e^{T_1^e\lambda} + (1-\alpha)e^{T_1^l\lambda}]$$
$$\leq \bar{T_2}^l = T_1^l + \frac{1}{\lambda}\ln\left(\frac{p_0}{1-p_0}\frac{\lambda R - i}{i}\right)$$

We then prove the fact that under this pooling contract, the agent always hide the early completion. When the agent choose to reveal the early completion, he will get funded from T_1^e to $\tilde{T}_2$ and the investor is supposed to earn zero, such that:

$$U(Revealing) = \int_{T_1^e}^{T_2} p_0 e^{-(t-T_1^e)\lambda}\lambda(R-B)\mathrm{d}t$$
$$V(Early) = \int_{T_1^e}^{T_2} \{p_0 e^{-(t-T_1^e)\lambda}\lambda B - (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0)i$$
$$= 0$$

Thus, the expected payoff of the agent is: $U(Revealing) = \int_{T_1^e}^{T_2}\{p_0 e^{-(t-T_1^e)\lambda}\lambda R - (p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0)i\}\mathrm{d}t$.

Whereas, if the agent choose to hide, the expected payoffs are as follows:

$$
\begin{aligned}
U(Hiding) =& \int_{T_1^e}^{T_1^l} p_0 e^{-(t-T_1^l)\lambda}\lambda R - -(p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0) i \mathrm{d}t \\
&+ \int_{T_1^l}^{T_2} p_{T_1^l} e^{-(t-T_1^l)\lambda}\lambda (R-B) \mathrm{d}t \\
V(Late) =& \int_{T_1^l}^{T_2} \{(\alpha p_{T_1^l} + (1-\alpha)p_0) e^{-(t-T_1^l)\lambda}\lambda B \\
&- \Big((\alpha p_{T_1^l} + (1-\alpha)p_0) e^{-(t-T_1^l)\lambda} + 1 - (\alpha p_{T_1^l} + (1-\alpha)p_0)\Big) i\} \mathrm{d}t \\
=& 0
\end{aligned}
$$

Substitute for expected bonus:

$$
\begin{aligned}
U(Hiding) =& \int_{T_1^e}^{T_1^l} p_0 e^{-(t-T_1^l)\lambda}\lambda R - -(p_0 e^{-(t-T_1^e)\lambda} + 1 - p_0) i \mathrm{d}t \\
&+ \int_{T_1^l}^{T_2} p_{T_1^l} e^{-(t-T_1^l)\lambda}\lambda R \mathrm{d}t - \frac{p_{T_1^l}}{\alpha p_{T_1^l} + (1-\alpha)p_0} i \int_{T_1^l}^{T_2} \\
&(\alpha p_{T_1^l} + (1-\alpha)p_0) e^{-(t-T_1^l)\lambda} + 1 - (\alpha p_{T_1^l} + (1-\alpha)p_0) \mathrm{d}t
\end{aligned}
$$

Comparing $U(Hiding)$ and $U(Revealing)$, the difference is the expected investment cost from T_1^l: $\frac{p_{T_1^l}}{\alpha p_{T_1^l} + (1-\alpha)p_0} i \int_{T_1^l}^{T_2} (\alpha p_{T_1^l} + (1-\alpha)p_0) e^{-(t-T_1^l)\lambda} + 1 - (\alpha p_{T_1^l} + (1-\alpha)p_0) \mathrm{d}t$ versus $i \int_{T_1^l}^{T_2} (p_{T_1^l} e^{-(t-T_1^e)\lambda} + 1 - p_{T_1^l}) \mathrm{d}t$, where the former is smaller. Hence, when the first stage is completed early, the agent always chooses to hide under the pooling contract.

4.C Proof of Proposition 11

According to IC^l:

$$
\begin{aligned}
U^l(T^l, B^l) \geq& U^l(T^e, B^e) \\
\int_{T_1^l}^{T_2^l} p_0 e^{-(t-T_1^l)\lambda}\lambda (R - B^l) \mathrm{d}t \geq& \int_{T_1^l}^{T_2^e} p_0 e^{-(t-T_1^l)\lambda}\lambda (R - B^e) \mathrm{d}t.
\end{aligned}
$$

The integrated value of R during T_1^l and T_2^e could be cancelled, and the above inequality becomes the following:

$$
\int_{T_1^l}^{T_2^l} p_0 e^{-(t-T_1^l)\lambda}\lambda B^l \mathrm{d}t - \int_{T_1^l}^{T_2^e} p_0 e^{-(t-T_1^l)\lambda}\lambda B^e \mathrm{d}t \leq \int_{T_2^e}^{T_2^l} p^l_{T_2^e} e^{-(t-T_2^e)\lambda}\lambda R \mathrm{d}t,
$$

where $p^l_{T_2^e} = \frac{p_0 e^{-\lambda(T_2^e - T_1^l)}}{p_0 e^{-\lambda(T_2^e - T_1^l)} + 1 - p_0}$.

From IC^e:

$$\int_{T_1^l}^{T_2^l} p_{T_1^l} e^{-(t-T_1^l)\lambda} \lambda B^l \mathrm{d}t - \int_{T_1^e}^{T_2^e} p_0 e^{-(t-T_1^e)\lambda} \lambda B^e \mathrm{d}t \geq \int_{T_2^e}^{T_2^l} p^e_{T_2{}^e} e^{-(t-T_2^e)\lambda} \lambda R \mathrm{d}t,$$

where $p^e_{T_2{}^e} = \frac{p_0 e^{-\lambda(T_2{}^e - T_1^e)}}{p_0 e^{-\lambda(T_2{}^e - T_1^e)} + 1 - p_0}$.

As $p_0 > p_{T_1^l}$, $\int_{T_1^l}^{T_2^l} p_0 e^{-(t-T_1^l)\lambda} \lambda B^l \mathrm{d}t > \int_{T_1^l}^{T_2^l} p_{T_1^l} e^{-(t-T_1^l)\lambda} \lambda B^l \mathrm{d}t$. Moreover, as $T_1^e < T_1^l$,
$\int_{T_1^l}^{T_2^e} p_0 e^{-(t-T_1^l)\lambda} \lambda B^e \mathrm{d}t < \int_{T_1^e}^{T_2^e} p_0 e^{-(t-T_1^e)\lambda} \lambda B^e \mathrm{d}t$, so the left hand side of IC^l is greater than the left hand side of IC^e, which means $IC^l(left) > IC^e(left)$ holds:

$$\int_{T_1^l}^{T_2^l} p_0 e^{-(t-T_1^l)\lambda} \lambda B^l \mathrm{d}t - \int_{T_1^l}^{T_2^e} p_0 e^{-(t-T_1^l)\lambda} \lambda B^e \mathrm{d}t > \int_{T_1^l}^{T_2^l} p_{T_1^l} e^{-(t-T_1^l)\lambda} \lambda B^l \mathrm{d}t - \int_{T_1^e}^{T_2^e} p_0 e^{-(t-T_1^e)\lambda} \lambda B^e \mathrm{d}t$$

Thus, $IC^l(right) \geq IC^l(left) > IC^e(left) \geq IC^e(right)$. $IC^l(right) > IC^e(right)$ indicates the following:

$$\int_{T_2^e}^{T_2^l} p^l_{T_2^e} e^{-(t-T_2^e)\lambda} \lambda R \mathrm{d}t > \int_{T_2^e}^{T_2^l} p^e_{T_2^e} e^{-(t-T_2^e)\lambda} \lambda R \mathrm{d}t \tag{4.9}$$

As $p^l_{T_2^e} > p^e_{T_2{}^e}$, the above inequality holds. Hence, the first-best experimentation time can be achieved as long as B^l and B^e satisfies the relationship:

$$R\gamma^e < \int_{T_1^l}^{T_2^l} p_0 e^{-(t-T_1^l)\lambda} \lambda B^l \mathrm{d}t - \int_{T_1^l}^{T_2^e} p_0 e^{-(t-T_1^l)\lambda} \lambda B^e \mathrm{d}t \leq R\gamma^l \tag{4.10}$$

where $\gamma^l = \int_{T_2^e}^{T_2^l} p^l_{T_2^e} e^{-(t-T_2^e)\lambda} \lambda \mathrm{d}t$ and $\gamma^e = \int_{T_2{}^e}^{T_2{}^l} p^e_{T_2{}^e} e^{-(t-T_2{}^e)\lambda} \lambda \mathrm{d}t$. Thus, the first-best experimentation times $\bar{T_2}^l$ and $\bar{T_2}^e$ can be chosen as long as the above inequality of the bonuses holds.

Then we want to check whether the bonuses under full information satisfy the above relationship, which is equivalent of checking whether the difference between expected $\bar{B}^l$ and $\bar{B}^e$ is within the interval. From Proposition 9, we know that:

$$\bar{B}^l \int_{T_1^l}^{\bar{T_2^l}} p_0 e^{-(t-T_1^l)\lambda} \lambda - \bar{B}^e \int_{T_1^e}^{\bar{T_2^e}} p_0 e^{-(t-T_1^e)\lambda} \lambda = (T_1^l - T_1^e) i = \int_{\bar{T_2}^e}^{\bar{T_2}^l} i \mathrm{d}t > R\gamma^e$$

Thus, we just need to check whether the upper bound is satisfied. Make some rearrangement of the inequality Equation 4.10:

$$\int_{T_1^l}^{\bar{T_2}^l} p_0 e^{-(t-T_1^l)\lambda}\lambda B^l \mathrm{d}t - \int_{T_1^l}^{\bar{T_2}^e} p_0 e^{-(t-T_1^l)\lambda}\lambda B^e \mathrm{d}t > \int_{T_1^l}^{\bar{T_2}^l} p_0 e^{-(t-T_1^l)\lambda}\lambda B^l \mathrm{d}t - \int_{T_1^e}^{\bar{T_2}^e} p_0 e^{-(t-T_1^l)\lambda}\lambda B^e \mathrm{d}t$$

Substituting the $\bar{B}^e$ and $\bar{B}^l$, we check the relationship between $(T_1^l - T_1^e)i$ and $R\gamma^l$:

$$\int_{T_1^l}^{\bar{T_2}^l} p_0 e^{-(t-T_1^l)\lambda}\lambda \bar{B}^l \mathrm{d}t - \int_{T_1^l}^{\bar{T_2}^e} p_0 e^{-(t-T_1^l)\lambda}\lambda \bar{B}^e \mathrm{d}t = (T_1^l - T_1^e)i = \int_{\bar{T_2}^e}^{\bar{T_2}^l} i \mathrm{d}t < \int_{T_2^e}^{T_2^l} p^l_{T_2^e} e^{-(t-T_2^e)\lambda}\lambda R \mathrm{d}t$$

As the late agent would not terminate the project until $\bar{T_2}^l$, so $p^l_{T_2^e}\lambda R - i > 0$ holds.

4.D Proof of Proposition 12

When $p_0^e > p_0^l$ and $\bar{T_2}^e > \bar{T_2}^l$. To stop the late completion from mimicking the early, IC^l is imposed, which indicates the following:

$$\int_{T_1^l}^{T_2^e} p_0^l e^{-(t-T_1^l)\lambda}\lambda B^e \mathrm{d}t - \int_{T_1^l}^{T_2^l} p_0^l e^{-(t-T_1^l)\lambda}\lambda B^l \mathrm{d}t \geq \int_{T_2^l}^{T_2^e} p^l_{T_2^l} e^{-(t-T_2^l)\lambda}\lambda R \mathrm{d}t,$$

where $p^l_{T_2^l} = \frac{i}{\lambda R}$.

Additionally, we impose IC^e, which indiactes the following:

$$\int_{T_1^e}^{T_2{}^e} p_0^e e^{-(t-T_1^e)\lambda}\lambda B^e \mathrm{d}t - \int_{T_1^l}^{T_2{}^l} p^e_{T_1^l} e^{-(t-T_1^l)\lambda}\lambda B^l \mathrm{d}t \leq \int_{T_2^l}^{T_2^e} p^e_{T_2{}^l} e^{-(t-T_2{}^l)\lambda}\lambda R \mathrm{d}t$$

where $p^e_{T_1^l} = \frac{p_0^e e^{-\lambda(\bar{T_1}^l - T_1^e)}}{p_0^e e^{-\lambda(T_1{}^l - T_1^e)} + 1 - p_0^e}$.

As the left hand side of IC^e is greater than the left hand side of IC^l, the following holds:

$$\int_{T_1^e}^{T_2{}^e} p_0^e e^{-(t-T_1^e)\lambda}\lambda B^e \mathrm{d}t - \int_{T_1^l}^{T_2{}^l} p^e_{T_1^l} e^{-(t-T_1^l)\lambda}\lambda B^l \mathrm{d}t > \int_{T_2^l}^{T_2^e} e^{-(t-T_2^l)\lambda} i \mathrm{d}t$$

Thus, the first-best bonuses and experimentation times are achievable as long as the differences between the expected bonuses are within the range $(\int_{T_2^l}^{T_2^e} e^{-(t-T_2^l)\lambda} i\mathrm{d}t, \int_{T_2^l}^{T_2^e} p^e_{T_2{}^l} e^{-(t-T_2{}^l)\lambda}\lambda R\mathrm{d}t]$.

Now we turn to the case when $p_0^e > p_0^l$ and $\bar{T_2}^e < \bar{T_2}^l$. The agent with an early completion still have strong incentive to mimic the late due to extended experimentation time although it carries a higher prior belief.

IC^e indicates the following:

$$\int_{T_1^l}^{T_2{}^l} p^e_{T_1^l} e^{-(t-T_1^l)\lambda}\lambda B^l \mathrm{d}t - \int_{T_1^e}^{T_2{}^e} p_0^e e^{-(t-T_1^e)\lambda}\lambda B^e \mathrm{d}t \geq \int_{T_2^e}^{T_2^l} p^e_{T_2{}^e} e^{-(t-T_2{}^e)\lambda}\lambda R\mathrm{d}t,$$

where $p^e_{T_1^l} = \frac{p_0^e e^{-\lambda(\bar{T_1}{}^l - T_1^e)}}{p_0^e e^{-\lambda(T_1{}^l - T_1^e)} + 1 - p_0^e}$ and $p^e_{T_2{}^e} = \frac{p_0^e e^{-\lambda(T_2{}^e - T_1^e)}}{p_0^e e^{-\lambda(T_2{}^e - T_1^e)} + 1 - p_0^e}$.

IC^l indicates the following:

$$\int_{T_1^l}^{T_2^l} p_0^l e^{-(t-T_1^l)\lambda}\lambda B^l \mathrm{d}t - \int_{T_1^l}^{T_2^e} p_0^l e^{-(t-T_1^l)\lambda}\lambda B^e \mathrm{d}t \leq \int_{T_2^e}^{T_2^l} p^l_{T_2^e} e^{-(t-T_2^e)\lambda}\lambda R\mathrm{d}t,$$

where $p^l_{T_2^e} = \frac{p_0^l e^{-\lambda(T_2^e - T_1^l)}}{p_0^l e^{-\lambda(T_2^e - T_1^l)} + 1 - p_0^l}$. When $p^e_{T_1^l} \leq p_0^l$, the first-best experimentation times are achievable. The expected bonuses gap is the same as Proposition 11. When $p^e_{T_1^l} > p_0^l$, it contradicts with the assumption that $\bar{T_2}^e < \bar{T_2}^l$.

Bibliography

Akcigit, Ufuk and Qingmin Liu (2015), "The role of information in innovation and competition." *Journal of the European Economic Association*.

Akerlof, George A. (1970), "The market for "lemons": Quality uncertainty and the market mechanism." *The Quarterly Journal of Economics*, 84, 488–500.

Arrow, Kenneth (1962), "Economic welfare and the allocation of resources for invention." In *The rate and direction of inventive activity: Economic and social factors*, 609–626, Princeton University Press.

Baker, George, Robert Gibbons, and Kevin J Murphy (2002), "Relational contracts and the theory of the firm." *Quarterly Journal of economics*, 39–84.

Bentley MacLeod, W (2003), "Optimal contracting with subjective evaluation." *The American Economic Review*, 93, 216–240.

Bergemann, Dirk and Ulrich Hege (1998), "Venture capital financing, moral hazard, and learning." *Journal of Banking & Finance*, 22, 703–735.

Bergemann, Dirk and Ulrich Hege (2005), "The financing of innovation: learning and stopping." *RAND Journal of Economics*, 719–752.

Bergemann, Dirk, Ulrich Hege, and Liang Peng (2010), "Venture capital and sequential investments." *Available at SSRN https://dx.doi.org/10.2139/ssrn.1572103*.

Bobtcheff, Catherine and Raphaël Levy (2014), "More haste, less speed? signaling through investment timing." Technical report, Working paper.

Bolton, Patrick and Christopher Harris (1999), "Strategic experimentation." *Econometrica*, 67, 349–374.

Bonatti, Alessandro and Johannes Horner (2011), "Collaborating." *American Economic Review*, 101, 632–63.

Bougheas, Spiros (2004), "Internal vs external financing of *R&D*." *Small Business Economics*, 22, 11–17.

Bouvard, Matthieu (2014), "Real option financing under asymmetric information." *Review of Financial Studies*, 27, 180–210.

Brealey, Richard, Hayne E Leland, and David H Pyle (1977), "Informational asymmetries, financial structure, and financial intermediation." *The journal of Finance*, 32, 371–387.

Cho, In-Koo and David M Kreps (1987), "Signaling games and stable equilibria." *The Quarterly Journal of Economics*, 179–221.

Dahiya, Sandeep and Korok Ray (2012), "Staged investments in entrepreneurial financing." *Journal of Corporate Finance*, 18, 1193–1216.

Danzon, Patricia M, Sean Nicholson, and Nuno Sousa Pereira (2005), "Productivity in pharmaceutical–biotechnology R&D: the role of experience and alliances." *Journal of health economics*, 24, 317–339.

David, Eric, Amit Mehta, Troy Norris, Navjot Singh, and Tony Tramontin (2010), "New frontiers in pharma R&D investment." URL `https://www.mckinsey.com/industries/pharmaceuticals-and-medical-products/our-insights/new-frontiers-in-pharma-r-and-38d-investment`.

Demarzo, Peter M and Yuliy Sannikov (2016), "Learning, termination, and payout policy in dynamic incentive contracts." *The Review of Economic Studies*, 84, 182–236.

Drugov, Mikhail and Rocco Macchiavello (2014), "Financing experimentation." *American Economic Journal: Microeconomics*, 6, 315–349.

Frank, Murray Z and Vidhan K Goyal (2011), "Trade-off and pecking order theories of debt." *Handbook of Empirical Corporate Finance: Empirical Corporate Finance*, 2, 135.

Galende, Jesús and Juan Manuel de la Fuente (2003), "Internal factors determining a firms innovative behaviour." *Research Policy*, 32, 715–736.

Gomes, Renato, Daniel Gottlieb, and Lucas Maestri (2013), "Experimentation and project selection: Screening and learning." Technical report, Working paper.

Green, Brett and Curtis R Taylor (2016), "Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects." *American Economic Review*, 106, 3660–3699.

Grenadier, Steven R and Neng Wang (2005), "Investment timing, agency, and information." *Journal of Financial Economics*, 75, 493–533.

Halac, Marina, Navin Kartik, and Qingmin Liu (2016a), "Optimal contracts for experimentation." *The Review of Economic Studies*, 83, 1040–1091.

Halac, Marina, Navin Kartik, and Qingmin Liu (2016b), "Contests for experimentation." Technical report, Columbia Business School working paper.

Hall, Bronwyn H and Josh Lerner (2010), "The financing of r&d and innovation." *Handbook of the Economics of Innovation*, 1, 609–639.

He, Zhiguo, Bin Wei, Jianfeng Yu, and Feng Gao (2017), "Optimal long-term contracting with learning." *The Review of Financial Studies*, 30, 2006–2065.

Heidhues, Paul, Sven Rady, and Philipp Strack (2015), "Strategic experimentation with private payoffs." *Journal of Economic Theory*, 159, 531–551.

Himmelberg, Charles P and Bruce C Petersen (1994), "R & d and internal finance: A panel study of small firms in high-tech industries." *The Review of Economics and Statistics*, 38–51.

Holmstrom, Bengt (1982), "Moral hazard in teams." *The Bell Journal of Economics*, 324–340.

Horner, Johannes and Larry Samuelson (2013), "Incentives for experimenting agents." *The RAND Journal of Economics*, 44, 632–663.

Kaplan, Steven N and Per Strömberg (2003), "Financial contracting theory meets the real world: An empirical analysis of venture capital contracts." *The Review of Economic Studies*, 70, 281–315.

Keller, Godfrey and Sven Rady (2010), "Strategic experimentation with poisson bandits." *Theoretical Economics*, 5, 275–311.

Keller, Godfrey, Sven Rady, and Martin Cripps (2005), "Strategic experimentation with exponential bandits." *Econometrica*, 73, 39–68.

Laffont, Jean-Jacques and David Martimort (2009), *The theory of incentives: the principal-agent model*. Princeton university press.

Lerner, Josh, Hilary Shane, and Alexander Tsai (2003), "Do equity financing cycles matter? evidence from biotechnology alliances." *Journal of Financial Economics*, 67, 411–446.

Levin, Jonathan (2003), "Relational incentive contracts." *The American Economic Review*, 93, 835–857.

MacLeod, W Bentley and James M Malcomson (1989), "Implicit contracts, incentive compatibility, and involuntary unemployment." *Econometrica: Journal of the Econometric Society*, 447–480.

Manso, Gustavo (2011), "Motivating innovation." *The Journal of Finance*, 66, 1823–1860.

Maskin, Eric and Jean Tirole (1992), "The principal-agent relationship with an informed principal, ii: Common values." *Econometrica: Journal of the Econometric Society*, 1–42.

Moroni, Sofia (2014), "Experimentation in organizations." Technical report, Yale University.

Moroni, Sofia (2016), "Experimentation in organizations." Technical report, Department of Economics, University of Pittsburgh.

Myers, Stewart C (1984), "The capital structure puzzle." *The journal of finance*, 39, 574–592.

Myers, Stewart C and Nicholas S Majluf (1984), "Corporate financing and investment decisions when firms have information that investors do not have." *Journal of financial economics*, 13, 187–221.

Neher, Darwin V (1999), "Staged financing: an agency perspective." *The Review of Economic Studies*, 66, 255–274.

Nelson, Richard R. (1959), "The economics of invention: A survey of the literature." *The Journal of Business*, 32, 101–127.

Rayo, Luis (2007), "Relational incentives and moral hazard in teams." *The Review of Economic Studies*, 74, 937–963.

Rothschild, Michael and Joseph Stiglitz (1976), *Equilibrium in competitive insurance markets: An essay on the economics of imperfect information*. Springer.

Sannikov, Yuliy (2014), "Moral hazard and long-run incentives." Technical report, mimeo, Princeton University.

Spence, A Michael (1979), "Investment strategy and growth in a new market." *The Bell Journal of Economics*, 1–19.

Spence, Michael (1973), "Job market signaling." *The quarterly journal of Economics*, 355–374.

Succurro, Marianna and Giuseppina Damiana Costanzo (2016), "Finance and r&d investment: A panel study of italian manufacturing firms." *International Journal of Economics and Finance*, 8, 95.

The Faculty Of Pharmaceutical Medicine (2017), "Pharmaceutical medicine: Innovating science and innovating healthcare." URL `https://www.fpm.org.uk/policypublications/innovationintosuccess`.

Tirole, Jean (2006), *The Theory of Corporate Finance*. Princeton university press.

Wilson, Charles (1977), "A model of insurance markets with incomplete information." *Journal of Economic theory*, 16, 167–207.

Wolf, Christoph (2017), "Informative milestones in experimentation." Technical report, Department of Economics, University of Mannheim.

Yu, Jianfeng, Bin Wei, and Zhiguo He (2012), "Optimal long-term contracting with learning." In *2012 Meeting Papers*, 221, Society for Economic Dynamics.