Dynamic Problems and Learning

Microtheoretic Applications to Two-Sided Markets, Team Production, and the Principal-Agent Problem

Inaugural-Dissertation
zur Erlangung des Grades
Doctor oeconomiae publicae (Dr. oec. publ.)
an der Ludwig-Maximilians-Universität München

2012

vorgelegt von

Piers Trepper

Referent:	Prof. Sven Rady, Ph.D.
Korreferent:	Prof. Peter Norman Sørensen, Ph.D.
Promotionsabschlussberatung:	7. November 2012

Datum der mündlichen Prüfung: 23. Oktober 2012

Namen der Berichterstatter: Sven Rady, Peter Norman Sørensen, Fabian Herweg

Acknowledgements

First and foremost, I would like to thank my supervisor, Sven Rady. Not only were his skills of analytical and economic understanding an inspiration, but also did his careful and provident management style allow me to engage in various paths of economic interests. In addition, his friendly and supporting nature made working with him as a co-author but also working for him in teaching and seminars a great pleasure.

I am also much obliged to Peter Norman Sørensen for being a resourceful advisor and a great host in Copenhagen, and of course also for agreeing to be on my thesis committee. Further, I am very grateful to Fabian Herweg for joining my committee and for his feedback especially on the second chapter of this dissertation, and to Martin Peitz for the productive cooperation on the first chapter.

Working at the University in Munich in the last years has been a great delight, which to a large extent was due to the very friendly and comforting atmosphere at the "Seminar für dynamische Modellierung". My special thanks for this go to our cheerful and caring secretary Manuela Beckstein, our student research assistant Niklas Garnadt, and to my colleagues Lukas Buchheim, Gilbert Spiegel, and Ludwig Straub.

I would of course also like to thank my family for their constant support in all the roads I have taken in recent years.

Last but not least I would like to thank Caspar, Henrike, Michael, Sebastian, Stefan, and Werner for being friends.

Piers Trepper

Contents

List of Figures

Preface

The story of this dissertation is one of uncertainty, learning, and information. A common theme of all three chapters is the dynamic nature of the problem a risk-neutral decision maker faces. On the one hand, I analyze optimal learning strategies, i.e., matters of information generation (Chapters 1 and 2). On the other hand, I also discuss the implications of learning and unraveling information on optimal decisions, i.e., matters of information processing (Chapter 3).

Our time is often characterized as the *information age*, reflecting the accelerated arrival and wider scope of news flows. During such times it is crucial to dynamically adapt to these streams of news by incorporating the relevant information into individual and firm decisions. Unfortunately not all information is unambiguous in its content, in other words there is uncertainty.

When facing uncertainty about the state of the world, the classical way of updating one's beliefs after receiving new information is the well known rule of weighting the likelihood of a signal under an assumed event against its likelihood given the counterevent which was introduced by Bayes (1763). This is a rational way of processing information and thus reacting to the world's imponderables, and indeed one is usually well advised to take all relevant pieces of information into account when making decisions. However, processing information is only one dimension of dealing with uncertainty. In a dynamic world today's actions necessarily impact tomorrow's learning. A strategic decision maker should seize this opportunity to actively generate valuable news. Since it might be very fruitful to have more precise evidence at an earlier stage, substantial investments into producing information can turn out very profitable. This challenge of intertemporal optimization with improving information was first acknowledged in seminal works by Prescott (1972) and Rothschild (1974). We consider two types of markets on which information generation and processing are of particular relevance, namely two-sided markets and job markets.

The theory of two-sided markets, introduced by Rochet and Tirole (2003), is a rather new field of economic research. It has shaped our understanding of how markets work that contain two interacting types of participants who exert an externality on each other. The interaction usually takes place on some type of platform that has to take the interests of the two sides into account. Standard examples are credit card companies that have to consider the interests of customers using the card as well as merchants accepting the card, shopping malls that care about the interplay of the variety of shops and the number of visitors, night clubs that want to allure male and female guests, and smart phone providers that want to attract users on the one side and application developers on the other side. In general, the platform will be uncertain about the extent of the market sides' externalities. When each side of the market profits from an increased number of participants on the opposite side, knowledge about this mutual impact can become key. Little has been said so far on how a platform provider that faces a two-sided market may learn about and make use of such information. This is quite surprising as information might be very valuable, especially when the respective market is young and growing and the platform has not had time to gain experience. The market for smart phone applications, for example, is expected to increase by a factor of eight from US$ 1.9bn in 2009 to US$ 15.7bn in 2013.[1]

The optimal handling of information has also gained significance on job markets and for employment decisions. The importance of employment decisions for firms and enterprises was already addressed by Coase (1937). In his essay on "The Nature of the Firm" he discusses that it might be beneficial for organizations to engage in longer relationships with employees instead of just hiring contractors for specific tasks. A possible rationale for this, mentioned by Knight as early as in 1921, is the efficient allocation of risk. In the

[1] See www.research2guidance.com/smartphone-application-market-to-reach-us15-65-billion-in-2013/, accessed June 9, 2012.

economic literature it is regularly assumed that for various reasons entrepreneurs have a higher capacity of risk tolerance than their employees and should, therefore, bear a more sizeable share of the uncertainty associated with business opportunities. This puts the entrepreneur in a position of having to act prudently in his employment decisions as, to some extent, he effectively lays his fate into the employees' hands unless he is able and willing to fully monitor the production process. When workers are heterogeneous it is, therefore, of utmost interest to the employer to gather information on workers' characteristics in order to find the most apt workers for his purposes.

The argument is reinforced by the notion that in our times skilled work has gained considerably in importance in the industrialized world and will do so even more in the future. The requirement of livelong learning is more and more emphasized and employment biographies nowadays are seldom single-tracked. Identifying and attracting talent as well as retaining and motivating skilled employees has turned into a major challenge for employers, especially as individuals tend to change their jobs and pursue new career goals more frequently. Hand in hand with the appreciation of skilled work the prominence of team work also has increased strongly in recent times. Certainly, organizing work in teams has several benefits, e.g., spill-over and motivational effects. But it also imposes further challenges on an employer who wants to identify talent: When work is accomplished in teams it is harder to distinguish individual contributions to team results, i.e., to infer individual worker characteristics. Therefore, the importance of optimally gathering and deducting information has considerably increased in employment relations.

This thesis addresses such issues of information acquisition and processing. Chapter 1 considers optimal strategic learning in a two-sided markets environment. Chapter 2 deals with the optimal information generation and employment decision when production is conducted in teams and the employer wants to learn about individual ability. The third chapter discusses the optimal reaction to information about ability in a principal-agent framework where positive news on ability of the agent at the same time are negative news on the costs of giving incentives. All three chapters of this dissertation are self-contained and include their own introductions and appendices such that they can be read independently.

The stakeholder in the first chapter, which is based on joint work with Martin Peitz and Sven Rady, is a platform provider that is exposed to uncertainty about the impact the two market sides it faces have on each other. The demand on each side is linear in the number of participants on the opposite side. The platform maximizes lifetime revenue. It can either choose prices or quantities, i.e., the number of participants on each side, in a continuous-time infinite-horizon setting. While at first sight the price setting appears to be the more appealing one the quantity setting turns out to be more tractable and offers clear intuitions. The trade-off at every instant of time is between maximizing current revenue and deviating from this in order to gain additional information on the size of the externalities. The amount of learning is increasing in the number of participants on each side.

A first conjecture for a platform provider that is interested in learning would be to increase quantities on both sides and thereby reduce prices. Indeed this holds true when both sides of the market are sufficiently similar with respect to the size of the externality exerted. However, we find that it need not be true that both prices decrease when market sides have a rather different impact on each other. While the price on one side of the market will always be lowered, the direction of the price adjustment on the other side can be ambiguous. The reason for this is that the platform provider has two instruments at hand, and can thus pursue two goals at once in the following way: For any optimal level of information there is a range of quantity choices all generating this amount of information. Within this set even the most patient platform provider, i.e., one that is solely interested in learning, will choose the combination yielding the highest expected revenue. Because of the two choice variables it might be optimal to increase participation on side A of the market starkly, especially if the impact of side A on side B is sizeable, while only moderately pushing demand on the other side B, especially if it only has a minor impact on demand on side A. In this case a patient platform provider recoups part of the rent side B gains from the higher amount of participants on side A by inducing a higher price compared to the myopic optimum on side B.

In addition, not only will the patient platform provider optimize current revenue, but also will even a fully myopic decision maker gather information over time and learn perfectly in the end, as positive quantities always generate information.

Our model of two-sided markets with a monopolistic platform builds on Armstrong (2006). We contribute to the literature by introducing uncertainty and learning and thereby extending this work. The structure of the learning process in a continuous-time infinite-horizon is related to Keller and Rady (1999), who consider a similar set-up in a one-sided market. Here, our main contribution is to provide the decision maker with a second instrument to learn, which allows for the simultaneous pursuit of two goals.

The second chapter considers an entrepreneur ("she") who faces uncertainty about the ability of her workers. While inferring such ability from a series of outputs from a single worker is a statistical problem already well understood, few insights have been gained so far on the optimal policy of an employer when production is carried out in teams. Today, hardly any job exists that does not at least in part include working in a team, and the notion of "capacity for teamwork" has become a standard clause in job postings. This puts employers into the difficult position of having to assess individual ability based on team outputs, i.e., from signals that are jammed in the sense that they contain joint information on several individuals. In general, this can lead to situations in which even after infinite time and absent any other shocks the employer does not fully learn about individual ability. This is in stark contrast to the individual worker case.

However, observing and judging from output is not the only thing the employer can do in a job market environment. In particular, at least in the long run an entrepreneur can usually decide on whom to employ and whom to potentially replace, as well as on who to pair in a team with whom. In terms of information generation this leads to the possibility of perfectly learning about an individual's ability by pairing him infinitely often with random partners. This way team output becomes a series of noisy observations on one individual's ability. But this procedure might become very costly, especially when it takes a long time to verify that an agent is of insufficient ability, as the employer foregoes some revenue in every period. It might thus be more attractive to employ a more strategic approach that takes current profits into account. Hence, once more the trade-off is one between current revenue and the acquisition of information: An entrepreneur in every period hires workers of unknown, either high or low ability that produce output in teams of two. Output is either high, medium, or low, depending on the combination of the ability levels of the employees. While output is fully revealing about overall ability levels, these levels are not immediately attributable to individual workers. Once the ability of an individual worker is fully known, though, he can serve as a perfect identifier for his team mate's ability. As this "identificational quality" is independent of the worker's actual level of ability, it could give rise to the incentive to reemploy workers of known low ability for reasons of information acquisition: When such a worker is reemployed and output is medium in the next period, the employer can be sure that the new team member is of high ability. However, it turns out that this will never be optimal as there always exists an alternative strategy that does not include reemployment of a known low ability worker and yields higher expected revenues.

The goal of the entrepreneur is to find a team of two high ability workers as quickly as possible. However, when the entrepreneur can only employ one team at a time it might be optimal to settle for a mixed team of a high and a low ability worker: When the employer cannot identify who is who in the mixed team, she might be reluctant to replace anyone for fear of replacing the high type. This argument will gain importance when high types are in low supply, benefits from a second high type are relatively low, or future income has a minor weight when compared to current income. In these situations the entrepreneur will not want to risk losing a high ability worker, will in effect not be able to generate any new information, and will thus quit learning.

The second part of the chapter considers a situation in which the entrepreneur has already found a team of two high ability workers, a "winning team", and expands the firm by hiring a second team of new workers with unknown ability. This opens a new strategic option to the employer who, in addition to reemploying and replacing workers, can now also decide who of the four individuals employed to match with whom. The winning team will always produce high output when working together, and will generate perfect information when working with the newly employed individuals. The latter will always be optimal when the relative benefit of a second high type added to a team is low for the following reason: In addition to the improved information, splitting the winning team ensures that there is a first, relatively valuable high type in each of the teams employed and thus this strategy is also optimal in terms of current revenue. Splitting the winning team can also be optimal in some situations when a second high type adds more to team output than a first high type. This will be the case when the future and thus information is relatively important. The availability of the winning team further crowds out the informational incentive of reemploying known low ability workers so that such policies again never turn out to be optimal. Moreover, the employer will almost surely end up with four high ability workers, as she can always identify individual abilities of a mixed team by rematching them with the winning team without losing any revenues.

One of the few papers discussing ability inference in teams is Meyer (1994). She considers an overlapping generation model with two teams each composed of a junior and a senior worker in which juniors might work part-time for each team and output is noisy. Similar to our set-up team composition influences the quality of the signals inferred from team output. While in the work of Meyer aggregate output is independent of team composition and the focus is on optimal inference of information, our model adds the trade-off between exploration and exploitation as well as endogeneity of the reemployment decision.

The employer in the second chapter is mainly concerned about optimal termination decisions and ability inference and wages are not explicitly considered. But an entrepreneur will often also have to take into account the costs of giving incentives when making optimal employment decisions. In the third chapter we analyze such a situation and return to a standard bilateral setting with one principal ("he") and one agent. This chapter is based on joint work with Caspar Siegert. The uncertainty again concerns the ability of the agent but in addition the principal faces a classical hidden action problem. The principal wants to optimally implement effort in a two-period model and gains information on the agent's commonly unknown ability. As in the second chapter the principal can again make employment decisions, but due to the moral hazard problem we have to explicitly account for the costs of implementing effort, namely wages and bonuses.

In addition to the well established effort-and-ability set-up we allow for wealth effects making a wealthy agent more costly to motivate due to the lower marginal utility of income. While higher wealth may in general also have a positive insurance effect due to a higher risk-bearing capacity of the agent, we impose the overall effect to be negative, a result that is produced by a wide range of utility functions under some mild technical assumptions. The fundamental trade-off then is between wealth and ability. A successful agent is inferred to be more able and thus more likely to be successful again in the second period. Yet, he will also receive a bonus payment that makes him richer and thus less attractive for reemployment.

In our set-up the principal has four employment policies at his disposal: He may either always or never continue employment, or he may continue conditional on the binary output level observed, i.e., continue after high output only or after low output only. At first glance one would expect the principal to reward success by reemployment and thus only continue after high output. However, the empirical literature comparing firm performance and CEO tenure has struggled to establish such a strong link (e.g., Warner et al., 1988, Hadlock and Lumer, 1997, and Huson et al., 2001). On average, a manager in the bottom decile is only two to six percentage points more likely to be forced out of his job than a manager from the top decile.

We offer a potential explanation why this might be the case, as in our model the other three policies turn out to be optimal under certain parameter constellations: An unsuccessful and hence arguably less able agent might be willing to make up for past failure. He does so by exerting more effort at lower incentives due to his lower wealth and his higher marginal utility of income. In consequence, this might render an unsuccessful agent more attractive than a new or a successful one. This will be the case whenever uncertainty about ability is low, such that the information inferred from first-period output is negligible, and if the changes in wealth levels are significant.

The principal will always opt for the most able agent at hand, i.e., only continue employment after high output for low and very high differences in profit levels in period two. In these cases it is optimal to implement similar effort levels for all types of agents, namely low levels of effort when stakes are low, and close to maximum effort when they are high. The costs of implementing effort will be small for all agents in the first case, while in the second case differences in the costs of giving incentives become negligible compared to the prize at stake, such that differences in ability are the relevant criterion for employment.

However, unsuccessful agents become more attractive the more differences in effort levels matter. Intermediate differences in profits might lead to situations in which it is profitable to give high incentives to poor agents but not to rich agents. The increased effort from the former might then outweigh the ability advantage of the latter, such that only reemploying unsuccessful agents becomes optimal, especially if differences in inferred ability are small.

In addition, when effort is binary we can establish that the unconditional policies of always or never continuing employment become optimal for profit levels in between the ones for which reemployment conditional on the first-period outcome is optimal.

Most of the existing literature on dynamic agency problems with unknown ability deliberately abstracts from wealth effects on the costs of incentives. It usually either assumes risk neutrality of the agent (e.g., Holmström, 1999) or constant absolute risk aversion where higher wealth affects marginal utility of income but also marginal disutility of effort in a way such that both effects offset each other (e.g., Gibbons and Murphy, 1992). Thiele and Wambach (1999) discuss conditions under which principals prefer poorer agents in a static setting and Spear and Wang (2005) consider a dynamic setting with wealth effects and limited liability. However, neither of the two models incorporates ability. The analysis of the interaction of wealth and ability thus is our main contribution to the literature.

Chapter 1

Experimentation in Two-Sided Markets*

1.1 Introduction

In many real-world markets, transactions are intermediated through platforms. This chapter studies a monopolistic platform in a two-sided market framework. The platform provider is uncertain about the size of the positive externality each side of the market is exerting on the other and, therefore, may want to experiment in order to learn about the externality parameters. Its aim is to maximize expected lifetime profit in a continuous-time infinite-horizon setting.

In every instant of time, the platform provider's actions determine its current profit as well as the amount of information received. Thus, there is a trade-off between maximizing current profit and extracting information that will increase future profits. The higher the rate at which future profits are discounted, the more important current profit becomes, up to the extreme of myopic behavior which completely ignores information acquisition. Reversely, the benefit of information increases if the discount rate decreases, up to the opposite extreme of no discounting when maximal weight is put on learning.

*This chapter is based on joint work with Martin Peitz and Sven Rady.

We consider two variants of the model, one in which the platform provider sets prices and learns from quantities, and one in which the platform provider selects quantities and learns from prices.[1] Prices take the form of membership or subscription fees. In both versions, we first compute the myopic benchmark, then investigate the optimal experimentation policy of a forward-looking platform provider, and finally consider the undiscounted limit in which experimentation is maximal. Our investigation of the optimal experimentation policy relies on an analysis of the first-order conditions associated with the platform provider's Bellman equation; we show that the second-order conditions for a maximum are always satisfied. In general, there are no closed-form solutions for the platform provider's value function and optimal policy. Turning to the undiscounted limit, by contrast, we are able to identify special cases of the model that yield a maximal experimentation policy in closed form.

In the price-setting version of the model, we first establish that the experimenting platform provider will charge a fee lower than the myopic benchmark on at least one side of the market. This immediately implies that if the two sides are approximately symmetric with respect to the participants' intrinsic platform value, the strength of the externality and the informativeness of observed quantities, the provider will charge fees lower than their myopically optimal counterparts on both sides of the market. In sufficiently asymmetric settings, however, the platform provider may find it optimal to charge a fee higher than the myopic benchmark on one side of the market. More precisely, we show that a price increase may occur on a side that exerts a low externality on the other side, yet itself benefits from a strong externality in the other direction. In such a situation, it is optimal to increase participation on the side that exerts the strong externality by lowering the fee there and to extract part of the additional surplus through a higher fee on the side that exerts the weak externality.

In the quantity-setting variant of the model, we obtain analogous results for expected prices. While the platform provider increases the quantity on both sides of the market

[1] The price-setting version of the model seems more widely applicable, but the quantity-setting version turns out to be more tractable.

relative to the myopic benchmark, this may entail an increase in the expected price on one side if the externality that this side exerts is much weaker than the externality it experiences.

Pricing implications in two-sided markets have received a lot of attention in industrial economics recently. In general, a market is said to be two-sided whenever potential participants care about the number of counterparts on the other side of the market — i.e., when each side exerts an externality on the other side, be it positive or negative. Potential interactions take place on some platform or by means of some vehicle, allowing the provider of such a platform or vehicle to charge participants for services and to manage usage on both sides.

Real world examples and applications of two-sided markets are manifold. Examples include payment systems (where card holders will want to hold a card if many merchants accept it, while merchants will be willing to accept cards that many customers hold), game consoles (players, software developers), smart phones (users, application developers), nightclubs and matching agencies (men, women), shopping malls, supermarkets, and department stores (where consumers are interested in a large variety of products, and producers in a large number of customers).

Seminal papers on two-sided markets are Rochet and Tirole (2003) and (2006) and Armstrong (2006). For a theoretical investigation of media platforms see, in particular, Anderson and Coate (2005). A general model of monopoly platforms is analyzed by Nocke et al. (2007). Empirical work includes Rysman (2004) and Kaiser and Wright (2006). For a selective survey, see Rysman (2009). None of the existing literature treats two-sided markets in a setting of uncertainty where it is unclear how strong the relevant externalities are, and where the platform provider might benefit from experimenting with prices or quantities in order to learn about the true state of the world. Relative to the existing literature on two-sided markets, our contribution is to introduce uncertainty and learning into the set-up proposed by Armstrong (2006). This allows us to analyze how the optimal price structure differs from the myopic benchmark and how it evolves over time. Our analysis suggests that markets characterized by indirect network effects of uncertain size provide incentives for the experimenting platform provider to initially lower at least one price. This provides a new rationale for price discounts in dynamic two-sided markets.[2]

The economics literature on optimal experimentation by a single Bayesian decision maker starts with the work of Prescott (1972) and Rothschild (1974); a brief overview of this literature can be found in Keller and Rady (1999). Our contribution here is to extend the analysis of optimal experimentation to two-sided markets and, building upon the infinite-horizon continuous-time model of Keller and Rady (1999), to provide a tractable framework for it. To the best of our knowledge, ours is the first experimentation model in which the decision maker has more than one instrument (i.e., two quantities or two prices) with which to trade off exploration versus exploitation. Because of this, even a platform provider primarily concerned about information acquisition can still pursue the secondary goal of current profit maximization: from all pairs of actions generating the same amount of information, the optimal policy selects the pair with the highest current profit.

The remainder of the chapter is structured as follows. Section 1.2 presents the model for the price-setting platform provider and characterizes the evolution of beliefs. Section 1.3 analyzes the directions of optimal experimentation, while Section 1.4 elaborates on the maximal experimentation policy. The optimal policy of a quantity-setting platform provider is analyzed in Section 1.5. Section 1.6 concludes. Technical proofs are relegated to the appendix.

1.2 The Model

We propose a two-sided market model following Armstrong (2006) to focus on participation decisions. For tractability reasons, we analyze a setting with linear demand functions

[2]An alternative explanation could be dynamic consumer behavior which might make a platform provider strive to build up a critical mass. We exclude this channel by assuming that participants can revise their participation decision in each period at no cost.

on both sides of the market. We refer to the two sides as A and B. Depending on the application, these may be buyers and sellers, advertisers and readers, or men and women. The novelty is to introduce uncertainty with respect to the size of the network effect. Arguably, such uncertainty is an important feature of network industries: a platform provider typically cannot perfectly foresee how strongly one side reacts to the number of users on the other side and has to infer this from market outcomes which noisily reveal the true state of the world.

1.2.1 The price-setting platform provider

In each period, there is a continuum of participants on both sides of the market. Invoking a uniform distribution over the value of the outside option (on a support that is sufficiently large such that aggregate demand is decreasing when positive) gives rise to linear demand functions. The platform provider can set membership fees (M_A, M_B), but no usage fee.[3] Suppose that the total mass of potential participants is such that demand n_i on side $i = A, B$ satisfies $dn_i/dM_i = -1$. The resulting masses of participants n_A and n_B are then characterized by the system of linear equations

$$n_A = u_0 + \tilde{u} n_B - M_A, \tag{1.1}$$

$$n_B = \pi_0 + \tilde{\pi} n_A - M_B, \tag{1.2}$$

where u_0 and π_0 are the intrinsic platform values, and $\tilde{u}$ and $\tilde{\pi}$ are externality parameters. For the sake of concreteness, we assume positive intrinsic values and positive externalities. While the intrinsic values are common knowledge, the externality parameters are known to market participants, but not to the platform provider.[4] The provider only knows that $(\tilde{u}, \tilde{\pi}) \in \{(\underline{u}, \underline{\pi}), (\overline{u}, \overline{\pi})\}$ with $0 < \underline{u} < \overline{u} < 1$ and $0 < \underline{\pi} < \overline{\pi} < 1$. We denote the probability

[3] Our notation closely follows Belleflamme and Peitz (2010).

[4] We impose this for the sake of tractability. If side A, say, does not know the strength of the externality it exerts on the other side either, it has to form a belief about it. This, in turn, has to be taken into account by the platform provider who then must form a belief about the true strength of the externalities as well as about the belief of side A. We leave the analysis of such a model for future work. In the present set-up, only the platform provider holds beliefs and learns.

that the platform provider initially assigns to the realization $(\overline{u}, \overline{\pi})$ by p_0 and assume that this prior belief is non-degenerate, i.e., $0 < p_0 < 1$.[5]

As $\tilde{u}\,\tilde{\pi} \neq 1$, the system (1.1)–(1.2) has a unique solution, given by

$$\begin{aligned} n_A(M_A, M_B, \tilde{u}, \tilde{\pi}) &= \frac{u_0 - M_A + \tilde{u}(\pi_0 - M_B)}{1 - \tilde{u}\tilde{\pi}}, \\ n_B(M_A, M_B, \tilde{u}, \tilde{\pi}) &= \frac{\pi_0 - M_B + \tilde{\pi}(u_0 - M_A)}{1 - \tilde{u}\tilde{\pi}}. \end{aligned}$$

This constitutes the unique Nash equilibrium of the anonymous game that potential participants play for given membership fees.

In every period $t \in [0, \infty[$, the platform provider sets prices (M_A^t, M_B^t) and then observes noisy signals of the quantities $n_A(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})$ and $n_B(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})$. More precisely, the provider observes the cumulative quantity processes N_A^t and N_B^t with increments given by

$$\begin{aligned} dN_A^t &= n_A(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})\, dt + \sigma_A dZ_A^t, \\ dN_B^t &= n_B(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})\, dt + \sigma_B dZ_B^t, \end{aligned}$$

where Z_B^t and Z_A^t are independent standard Brownian motions and the constants σ_A and σ_B are positive. Note that, using normally distributed shocks, we cannot restrict the observed quantities dN_A^t and dN_B^t to be positive. We will, however, only allow the platform provider to choose prices such that, in expectation, demand is non-negative. Later, when we use quantities as choice variables, we can explicitly rule out negativity.

The platform provider's revenue increment is

$$\begin{aligned} dR_t &= M_A^t\, dN_A^t + M_B^t\, dN_B^t \\ &= M_A^t \left[n_A(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})\, dt + \sigma_A dZ_A^t \right] + M_B^t \left[n_B(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})\, dt + \sigma_B dZ_B^t \right]. \end{aligned}$$

We normalize costs to zero. Hence, the platform provider's total expected profits (ex-

[5] The assumption that the externality parameters are perfectly positively correlated is clearly restrictive. Imperfect correlation leads to a much more complicated situation with two-dimensional beliefs. We will see that our results for the quantity-setting scenario carry over to perfect *negative* correlation.

pressed in per-period terms) are

$$\mathrm{E}^{p_0}\left[\int_0^\infty re^{-rt}dR_t\right],$$

where $r > 0$ is the discount rate. By the martingale property of the stochastic integral with respect to Brownian motion, this expectation reduces to

$$\mathrm{E}^{p_0}\left[\int_0^\infty re^{-rt}\left\{M_A^t\, n_A(M_A^t, M_B^t, \tilde{u}, \tilde{\pi}) + M_B^t\, n_B(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})\right\} dt\right].$$

Let p_t be the subjective probability at time t that the platform provider assigns to the realization $(\overline{u}, \overline{\pi})$. Invoking the law of iterated expectations, we can rewrite total expected profits as

$$\mathrm{E}^{p_0}\left[\int_0^\infty re^{-rt}R(M_A^t, M_B^t, p_t)\ dt\right], \tag{1.3}$$

where

$$R(M_A, M_B, p) = M_A\,\mathrm{E}^p\left[n_A(M_A, M_B, \tilde{u}, \tilde{\pi})\right] + M_B\,\mathrm{E}^p\left[n_B(M_A, M_B, \tilde{u}, \tilde{\pi})\right] \tag{1.4}$$

is the expected *current* revenue from charging the fees (M_A, M_B) given the posterior belief p.

1.2.2 The myopic benchmark

If the platform provider were myopic (corresponding to $r = \infty$), it would maximize expected current revenue at each instant. Under our parameter restrictions, this revenue is strictly concave in (M_A, M_B), so the myopically optimal fees,

$$(M_A^\mu(p), M_B^\mu(p)) = \arg\max_{M_A, M_B} R(M_A, M_B, p),$$

are well-defined.

To compute these fees, we write the expected quantities appearing on the right-hand side of (1.4) as

$$\begin{aligned}
\mathrm{E}^p\left[n_A(M_A, M_B, \tilde{u}, \tilde{\pi})\right] &= \ell_0(p)[u_0 - M_A] + \ell_A(p)[\pi_0 - M_B], \\
\mathrm{E}^p\left[n_B(M_A, M_B, \tilde{u}, \tilde{\pi})\right] &= \ell_0(p)[\pi_0 - M_B] + \ell_B(p)[u_0 - M_A],
\end{aligned}$$

where

$$\ell_0(p) = \frac{1-p}{1-\underline{\pi}\,\underline{u}} + \frac{p}{1-\overline{\pi}\,\overline{u}}$$

and

$$\begin{aligned}\ell_A(p) &= \frac{(1-p)\underline{u}}{1-\underline{\pi}\,\underline{u}} + \frac{p\,\overline{u}}{1-\overline{\pi}\,\overline{u}},\\ \ell_B(p) &= \frac{(1-p)\underline{\pi}}{1-\underline{\pi}\,\underline{u}} + \frac{p\,\overline{\pi}}{1-\overline{\pi}\,\overline{u}}\end{aligned}$$

measure the expected direct and indirect effects, respectively, of lowering M_A or M_B.

With the dependence on the belief p suppressed, the right-hand side of (1.4) now becomes

$$[\ell_0 u_0 + \ell_A \pi_0]M_A + [\ell_0 \pi_0 + \ell_B u_0]M_B - \ell_0 M_A^2 - [\ell_A + \ell_B]M_A M_B - \ell_0 M_B^2.$$

As $0 < \ell_i < \ell_0$ for $i = A, B$ and hence $0 < \ell_A + \ell_B < 2\ell_0$, this quadratic function is indeed strictly concave, and we obtain

$$M_A^\mu = u_0 - \frac{[2\ell_0^2 - (\ell_A+\ell_B)\ell_A]\,u_0 - (\ell_A - \ell_B)\ell_0\pi_0}{4\ell_0^2 - (\ell_A+\ell_B)^2}\ , \tag{1.5}$$

$$M_B^\mu = \pi_0 - \frac{[2\ell_0^2 - (\ell_A+\ell_B)\ell_B]\,\pi_0 - (\ell_B - \ell_A)\ell_0 u_0}{4\ell_0^2 - (\ell_A+\ell_B)^2}\ . \tag{1.6}$$

As is well known from the literature on two-sided markets, the myopically optimal fee on one side of the market depends on market characteristics on *both* sides. Independent of the values of the externality parameters $\underline{u}, \overline{u}, \underline{\pi}, \overline{\pi}$, the fee on either side is always increasing in the intrinsic platform value on that same side. Whether or not the fee on one side is increasing in the intrinsic platform value on the *other* side depends on the relative strength of the network effects on both sides. To be precise, the fee M_A^μ is increasing in π_0 if and only if $\ell_A - \ell_B > 0$. Broadly speaking, when the externality side A is experiencing is higher than the one it is exerting, it benefits from the higher attractiveness of the platform for participants on side B as the intrinsic platform value π_0 rises, and can thus be charged a higher price; in this sense, side A "subsidizes" side B.

Further, M_A^μ can only exceed the intrinsic platform value u_0 if ℓ_A exceeds ℓ_B by a sufficient amount, and vice versa for M_B^μ and π_0. Thus, at most one fee at a time can exceed the intrinsic platform value and both fees will be lower than the respective intrinsic platform values if the expected externalities are equal ($\ell_A = \ell_B$) or close together.

For future reference, we denote the myopically optimal revenue by

$$R^\mu(p) = \max_{M_A, M_B} R(M_A, M_B, p) = R(M_A^\mu(p), M_B^\mu(p), p), \tag{1.7}$$

and, suppressing the dependence on p and other variables, rewrite the expected current revenue as

$$R = R^\mu - \ell_0 \left[M_A - M_A^\mu\right]^2 - \left[\ell_A + \ell_B\right]\left[M_A - M_A^\mu\right]\left[M_B - M_B^\mu\right] - \ell_0 \left[M_B - M_B^\mu\right]^2. \tag{1.8}$$

1.2.3 The evolution of beliefs

The platform provider revises its beliefs over time. Writing $\overline{n}_A(M_A, M_B) = n_A(M_A, M_B, \overline{u}, \overline{\pi})$ and using analogous definitions for $\underline{n}_A$, $\overline{n}_B$ and $\underline{n}_B$, we define

$$S(M_A, M_B) = \left[\frac{\overline{n}_A(M_A, M_B) - \underline{n}_A(M_A, M_B)}{\sigma_A}\right]^2 + \left[\frac{\overline{n}_B(M_A, M_B) - \underline{n}_B(M_A, M_B)}{\sigma_B}\right]^2.$$

Lemma 1 *The beliefs of the price-setting platform provider evolve according to*

$$dp_t \sim N\Big(0, p_t^2(1 - p_t)^2 S(M_A^t, M_B^t)\, dt\Big). \tag{1.9}$$

Proof: See the appendix. □

In the expression for the infinitesimal variance of the change in beliefs, $S(M_A^t, M_B^t)$ measures the information content of the demand observations obtained after setting prices (it is the sum of the squared signal-to-noise ratios of these observations). The more informative the observations are, the more strongly the beliefs react to them.[6]

[6] If the platform provider were uncertain about the intrinsic platform values (u_0, π_0) instead of the externalities (u, π), the quantity of information would be independent of the fees charged. The platform provider would then trivially always set the myopically optimal fees.

We can gain more precise insights into the structure of the function S by noting that

$$\begin{aligned}
\overline{n}_A(M_A, M_B) - \underline{n}_A(M_A, M_B) &= d_0\,[u_0 - M_A] + d_A\,[\pi_0 - M_B], \\
\overline{n}_B(M_A, M_B) - \underline{n}_B(M_A, M_B) &= d_0\,[\pi_0 - M_B] + d_B\,[u_0 - M_A],
\end{aligned}$$

where $d_i = \ell_i(1) - \ell_i(0) > 0$ for $i = 0, A, B$, and computing

$$S(M_A, M_B) = s_A\,[M_A - u_0]^2 + 2s_{AB}\,[M_A - u_0]\,[M_B - \pi_0] + s_B\,[M_B - \pi_0]^2$$

with the constants

$$s_A = \frac{d_0^2}{\sigma_A^2} + \frac{d_B^2}{\sigma_B^2}\,, \quad s_B = \frac{d_A^2}{\sigma_A^2} + \frac{d_0^2}{\sigma_B^2}\,, \quad s_{AB} = \frac{d_0 d_A}{\sigma_A^2} + \frac{d_0 d_B}{\sigma_B^2}\,.$$

Since $s_A s_B - s_{AB}^2 = \sigma_A^{-2}\sigma_B^{-2}(d_0^2 - d_A d_B)^2$ and, as a simple computation reveals, $d_0^2 < d_A d_B$, we see that S is a strictly convex function which assumes its global minimum of zero at $(M_A, M_B) = (u_0, \pi_0)$.

The beliefs $p = 0$ and $p = 1$ are absorbing—if the platform provider is subjectively sure about the true state of the world, no further learning is possible. For a non-degenerate belief p to be invariant under the optimal learning dynamics, two conditions are necessary: the platform must charge the myopically optimal fees at this belief (since this belief will persist forever), and the information content of the resulting demand observations must be zero (so that this belief will indeed persist). Taken together, this requires $S(M_A^\mu(p), M_B^\mu(p)) = 0$ or, equivalently, $(M_A^\mu(p), M_B^\mu(p)) = (u_0, \pi_0)$ which is impossible since the latter fees generate an expected current revenue of zero and marginally lowering one of the fees would improve upon that. There are thus no potentially confounding actions in the sense of Easley and Kiefer (1988). By well-known results, this implies

Lemma 2 *Any optimal pricing policy induces complete learning in the long run: The platform provider's posterior belief p_t converges to the truth almost surely as $t \to \infty$.*

To determine how the information content of observed quantities changes with the fees charged, we look at the partial derivatives of S with respect to M_A and M_B. Figure 1.1 visualizes this in (M_A, M_B)-space. Along the line $M_B = \pi_0 - \frac{s_A}{s_{AB}}(M_A - u_0)$ we have

$\frac{\partial S}{\partial M_A} = 0$, and along the line $M_B = \pi_0 - \frac{s_{AB}}{s_B}(M_A - u_0)$ we have $\frac{\partial S}{\partial M_B} = 0$; as $s_A s_B - s_{AB}^2 > 0$, the former line is steeper than the latter. For either line, the respective partial derivative is positive above the line and negative below. Along the myopically optimal pricing policy, both $\overline{n}_A - \underline{n}_A$ and $\overline{n}_B - \underline{n}_B$ can be shown to be positive, which directly implies that both partial derivatives of S are negative. Thus, the myopically optimal fees lie below both lines in Figure 1.1.

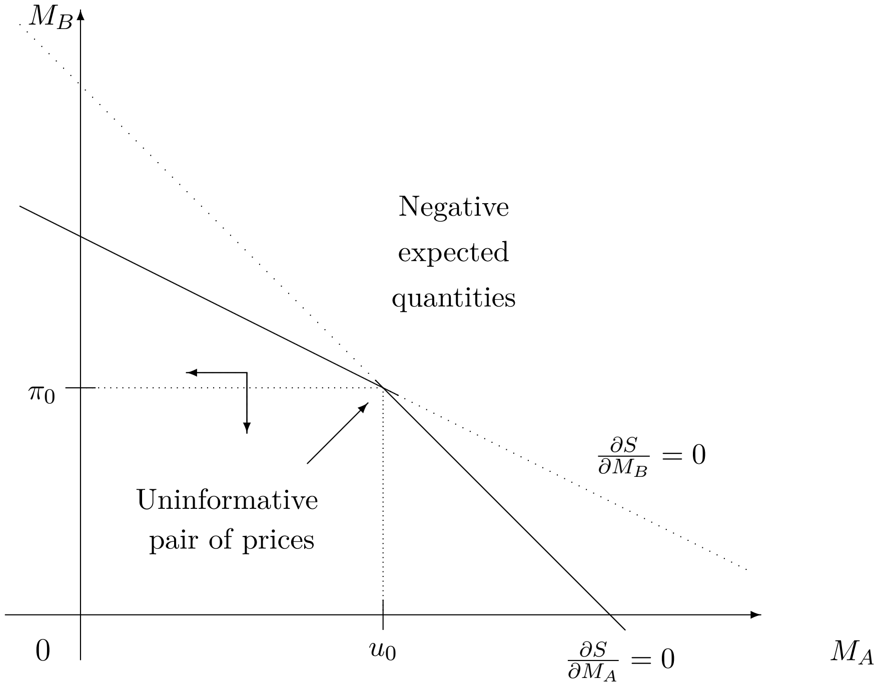

Figure 1.1: The directions of increasing information in the price plane.

In fact, this is true for *all* admissible price combinations.

Lemma 3 *Over the admissible range of prices, a price decrease on either side of the market increases the information content of observed quantities, whereas a price increase reduces it.*

Proof: The proof consists in showing that above either line of vanishing marginal information content in Figure 1.1, at least one of the implied expected quantities becomes negative. See the appendix for details. □

1.3 The Optimal Pricing Strategy

We are now ready to characterize the pricing strategy. In view of the objective function (1.3) and the law of motion (1.9), standard arguments yield the following Bellman equation for the platform provider's value function, v:

$$v(p) = \max_{M_A, M_B} \left\{ R(M_A, M_B, p) + \frac{p^2(1-p)^2}{2r} S(M_A, M_B)\, v''(p) \right\} . \qquad (1.10)$$

Arguing as in Keller and Rady (1999), one shows that v is strictly convex, twice continuously differentiable, and the unique solution to (1.10) subject to the condition that $v(p) = R^{\mu}(p)$ at $p = 0$ and 1, where the myopically optimal expected current revenue $R^{\mu}(p)$ has been defined in equation (1.7).

We can interpret the second term of the maximand in the Bellman equation as the value of information, given by the product of the shadow price of information, $p^2(1-p)^2 v''(p)/2r$, and the quantity of information, $S(M_A, M_B)$. For $p \in \{0, 1\}$, the value of information is zero, and the platform provider chooses the myopically optimal prices. For all other beliefs, the platform provider experiments, i.e., deviates from the myopic strategy so as to increase the information content of its demand observations.

The maximand in (1.10) is the sum of two quadratic functions, one of them strictly concave (expected current revenue), the other strictly convex (value of information). As the value function is bounded, so must be the maximum on the right-hand side of (1.10); and as admissible fees are unbounded below, the shadow price of information must actually be small enough for the combined quadratic function to be strictly concave (the precise argument is in the appendix).

This ensures that optimal fees are fully characterized by the (linear) first-order conditions for the maximization problem in (1.10). Using the representation of expected current revenues in (1.8), writing

$$V(p) = \frac{p^2(1-p)^2}{2r}\, v''(p)$$

for the shadow price of information, and suppressing the dependence on p, we compute the optimal pair of fees as

$$M_A^* = M_A^\mu + \frac{2V}{h(V)}\left\{2(\ell_0 - s_B V)S_A^\mu - (\ell_A + \ell_B - 2s_{AB}V)S_B^\mu\right\}, \tag{1.11}$$

$$M_B^* = M_B^\mu + \frac{2V}{h(V)}\left\{2(\ell_0 - s_A V)S_B^\mu - (\ell_A + \ell_B - 2s_{AB}V)S_A^\mu\right\}, \tag{1.12}$$

where

$$h(V) = 4(\ell_0 - s_A V)(\ell_0 - s_B V) - (\ell_A + \ell_B - 2s_{AB}V)^2$$

is the determinant of the Hessian matrix of the maximand in (1.10) and

$$\begin{aligned} S_A^\mu &= \frac{\partial S}{\partial M_A}(M_A^\mu, M_B^\mu) \;=\; s_A(M_A^\mu - u_0) + s_{AB}(M_B^\mu - \pi_0) \;<\; 0, \\ S_B^\mu &= \frac{\partial S}{\partial M_B}(M_A^\mu, M_B^\mu) \;=\; s_{AB}(M_A^\mu - u_0) + s_B(M_B^\mu - \pi_0) \;<\; 0 \end{aligned}$$

are the partial derivatives of the quantity of information S at the myopically optimal fees.[7] Strict concavity of the maximand in (1.10) means $\ell_0 - s_A V > 0$ and $h(V) > 0$, which in turn implies $\ell_0 - s_B V > 0$.

Our first result on the platform provider's optimal pricing strategy is

Proposition 1 *At any non-degenerate belief, the platform provider charges a fee lower than the myopic benchmark on at least one side of the market.*

Proof: Suppose that $M_A^* \geq M_A^\mu$. By (1.11), this implies $\ell_A + \ell_B - 2s_{AB}V > 0$ and

$$S_B^\mu \leq \frac{2(\ell_0 - s_B V)}{\ell_A + \ell_B - 2s_{AB}V}\, S_A^\mu .$$

As a consequence,

$$2(\ell_0 - s_A V)S_B^\mu - (\ell_A + \ell_B - 2s_{AB}V)S_A^\mu \leq \frac{h(V)}{\ell_A + \ell_B - 2s_{AB}V}\, S_A^\mu < 0,$$

and so $M_B^* < M_B^\mu$ by (1.12). In exactly the same way, $M_B^* \geq M_B^\mu$ implies $M_A^* < M_A^\mu$. □

[7] The argument why both of them are negative was given in Section 1.2.3.

The intuition for this result is clear. The purpose of deviating from the myopic optimum is to increase the information content of observed demands. As higher fees mean less information (see Lemma 3), at least one fee must be reduced relative to the myopic benchmark.

This has an obvious consequence for approximately symmetric setups.

Proposition 2 *For $(u_0, \underline{u}, \overline{u}, \sigma_A)$ sufficiently close to $(\pi_0, \underline{\pi}, \overline{\pi}, \sigma_B)$, the platform provider always sets both fees below their myopically optimal levels.*

Proof: For $(u_0, \underline{u}, \overline{u}, \sigma_A) = (\pi_0, \underline{\pi}, \overline{\pi}, \sigma_B)$, we have $M_A^\mu = M_B^\mu$ by (1.5)–(1.6), and $M_A^* - M_A^\mu = M_B^* - M_B^\mu \geq 0$ by (1.11)–(1.12) and Proposition 1, with a strict inequality, and the expression in curly brackets bounded away from 0, on the open unit interval. The result thus follows by continuous dependence of the value function and its second derivative on $(u_0, \underline{u}, \overline{u}, \sigma_A)$. □

The analysis of asymmetric settings is more complicated. A lower fee on one side of the market makes reducing the fee on the other side more attractive from an informational perspective (the cross-partial derivative of the quantity of information with respect to prices, s_{AB}, is positive), but less attractive as far as expected current revenue is concerned (its cross-partial derivative, $-(\ell_A + \ell_B)$, is negative). The overall effect is ambiguous.

A different way to see this is to think of the platform provider as following a two stage-procedure. At the first stage, it determines the combination of fees that maximizes current expected revenue subject to the constraint that a certain quantity of information be achieved. This amounts to identifying points of tangency between iso-information and iso-revenue curves in the (M_A, M_B)-plane. At the second stage, the provider then chooses the optimal quantity of information. Depending on the geometry of the iso-information and iso-revenue curves, this may lead it to charge a fee *higher* than in the myopic benchmark on one side of the market, as we shall see below.

To identify the directions of optimal experimentation in some asymmetric settings, we insert the expressions for S_A^μ and S_B^μ into (1.11)–(1.12) and collect the terms in $M_A^\mu - u_0$

and $M_B^\mu - \pi_0$, respectively:

$$\begin{aligned} M_A^* &= M_A^\mu + \frac{2V}{h(V)} \Big\{ \Big[2\ell_0 s_A - (\ell_A + \ell_B) s_{AB} - 2(s_A s_B - s_{AB}^2) V \Big] (M_A^\mu - u_0) \\ &\qquad + \left[2\ell_0 s_{AB} - (\ell_A + \ell_B) s_B \right] (M_B^\mu - \pi_0) \Big\}, \qquad (1.13) \\ M_B^* &= M_B^\mu + \frac{2V}{h(V)} \Big\{ \left[2\ell_0 s_{AB} - (\ell_A + \ell_B) s_A \right] (M_A^\mu - u_0) \\ &\qquad + \Big[2\ell_0 s_B - (\ell_A + \ell_B) s_{AB} - 2(s_A s_B - s_{AB}^2) V \Big] (M_B^\mu - \pi_0) \Big\}. \qquad (1.14) \end{aligned}$$

Proposition 3 *Let $s_A < s_{AB} < s_B$. Whenever both myopically optimal fees are lower than the respective intrinsic values, the platform provider lowers the fee on side B relative to the myopically optimal level.*

Proof: It is enough to show that in equation (1.14), the coefficients of $M_A^\mu - u_0$ and $M_B^\mu - \pi_0$ in the expression in curly brackets are positive. As $s_{AB} > s_A$ and $\ell_A + \ell_B < \ell_0$, this is obvious for the coefficient of $M_A^\mu - u_0$. Regarding the coefficient of $M_B^\mu - \pi_0$, we distinguish two cases. If $(\ell_A + \ell_B)/(2s_{AB}) < \ell_0/s_B$, the positivity of $s_A s_B - s_{AB}^2$ and the fact that $V < \ell_0/s_B$ imply that the coefficient of $M_B^\mu - \pi_0$ exceeds $2\ell_0 s_B - (\ell_A + \ell_B)s_{AB} - 2\ell_0(s_A s_B - s_{AB}^2)/s_B$, which is positive. If $(\ell_A + \ell_B)/(2s_{AB}) \geq \ell_0/s_B$, we have $V < (\ell_A + \ell_B)/(2s_{AB})$ and the coefficient of $M_B^\mu - \pi_0$ is no smaller than $2\ell_0 s_B - (\ell_A + \ell_B)s_{AB} - (\ell_A + \ell_B)(s_A s_B - s_{AB}^2)/s_{AB}$, which is again positive. □

The situation assumed in this proposition is one where the marginal informational benefit of lowering the fee is so much larger on side B than on side A that the platform provider will definitely lower the fee on side B. This situation arises naturally when the strength of the externality that side A exerts on side B is relatively well known, i.e., when $\overline{\pi}$ and $\underline{\pi}$ are relatively close to each other. More precisely, as $\overline{\pi}$ and $\underline{\pi}$ tend to a common value π, the ratios d_B/d_0, d_0/d_A, s_{AB}/s_B and s_A/s_{AB} all converge to π, which implies $s_A < s_{AB} < s_B$ for sufficiently small differences $\overline{\pi} - \underline{\pi}$.

The limiting case in which $\overline{\pi} = \underline{\pi} = \pi$ lends itself to a simple graphical illustration that will prove valuable when it comes to formulating a sufficient condition for the fee M_A to rise relative to the myopic benchmark. In fact, the identity $s_{AB}/s_B = s_A/s_{AB} = \pi$ implies that in the (M_A, M_B)-plane, the level curves of the function S are parallel straight lines with slope $-\pi$. These iso-information lines and the myopically optimal pricing policy are illustrated in Figure 1.2.

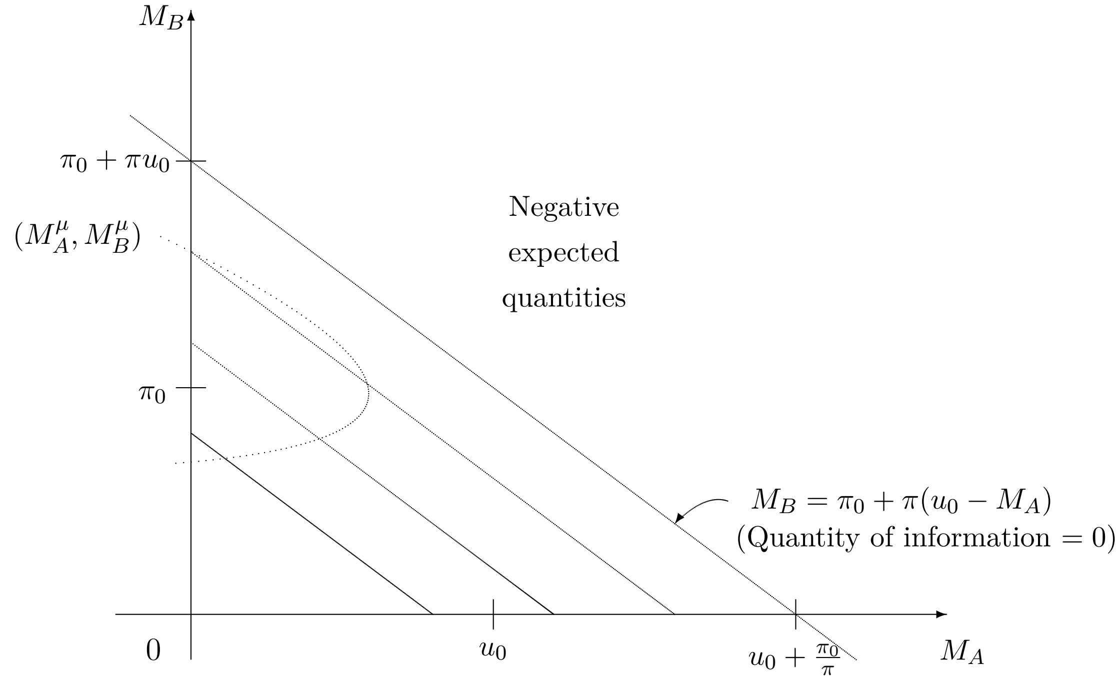

Figure 1.2: Iso-information lines and myopically optimal policy in the price plane when there is uncertainty about the externality parameter $\tilde{u}$ only ($\overline{\pi} = \underline{\pi} = \pi$).

The experimenting platform provider will deviate from the myopically optimal prices so as to reach an iso-information line that is closer to the origin in Figure 1.2. On any iso-information line, it will choose the fees that correspond to a point of tangency with an iso-revenue curve. As Figure 1.3 illustrates, the slope of the locus of tangency points between iso-information lines and iso-revenue curves (ellipses, to be precise) depends on parameters. In the left panel, this locus slopes upward – the optimal trade-off between information and current revenue induces a decrease in both fees for increased information. However, if the iso-information lines are rather flat (i.e., if π is small), it is optimal to decrease M_B but increase M_A as indicated by the locus of optimal fees in the right panel.

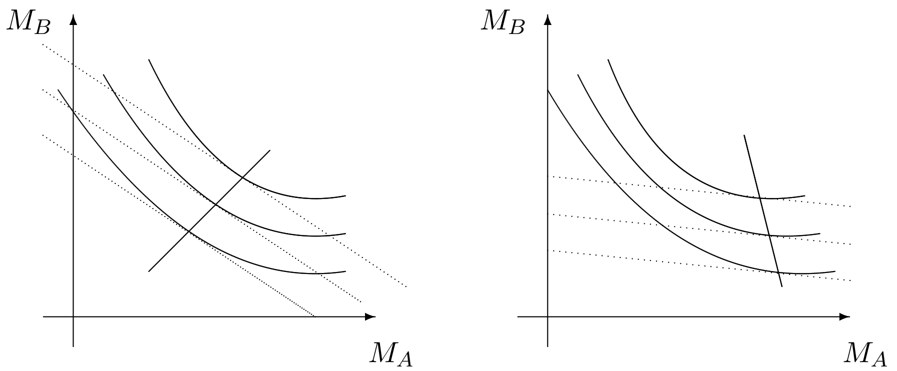

Figure 1.3: Two examples of iso-information lines (dotted) and iso-revenue curves (solid) for $\overline{\pi} = \underline{\pi} = \pi$. The solid line in each case indicates the locus of optimal fees.

This suggests that for $\overline{\pi}$ different from $\underline{\pi}$ but sufficiently small, we should also be able to see an optimal fee M_A^* that exceeds M_A^μ. Our next result bears this out.

Proposition 4 *For $\overline{\pi}$ sufficiently close to 0, the platform provider increases the fee on side A relative to the myopically optimal level.*

Proof: For $\overline{\pi} = \underline{\pi} = 0$, we have $\ell_0 = 1$ and $\ell_B = 0$, implying $d_0 = d_B = 0$ and $s_A = s_{AB} = 0$. By (1.6), moreover, $M_B^\mu - \pi_0$ is negative and bounded away from 0 on the unit interval. Now, the expression in curly brackets in (1.13) reduces to $-\ell_A s_B(M_B^\mu - \pi_0)$, which is positive and again bounded away from 0. The result thus follows by continuous dependence of the value function and its second derivative on $(\underline{\pi}, \overline{\pi})$. □

We can offer the following intuition for this result. When the externality that side A is exerting on side B is known to be very small, the platform provider learns most by lowering the fee on side B. Side A then benefits from higher participation on side B. Since participation on side A hardly affects participation on side B, the provider can safely extract part of the additional surplus given to side A by charging this side a higher fee.

1.4 Maximal Experimentation

In the previous section, we were able to analyze the *directions* of optimal experimentation without having to solve for the value function. To establish the precise *extent* of optimal experimentation, one could plug the fees (1.11)-(1.12) into the maximand in (1.10) and numerically solve the resulting second-order ordinary differential equation for the value function.

An alternative route to this differential equation is to write the Bellman equation in the form $0 = \max_{M_A,M_B}\{R - v + \frac{p^2(1-p)^2}{2r}\, S\, v''\}$ and to observe that the maximum remains zero, and the set of maximizers is unchanged, when we divide the maximand by the quantity of information, S.[8] Re-arranging then yields

$$\frac{p^2(1-p)^2}{2r}\, v''(p) = \min_{M_A,M_B} \frac{v(p) - R(M_A, M_B, p)}{S(M_A, M_B)} \; .$$

This in turn permits an alternative characterization of the optimal combination of fees as a function of the belief p and the associated value $v(p)$:

$$(M_A^*(p), M_B^*(p)) = \arg\min_{M_A,M_B} \frac{v(p) - R(M_A, M_B, p)}{S(M_A, M_B)} \; .$$

Arguing as in Keller and Rady (1999), one shows that the value $v(p)$ is decreasing in r at all p in the open unit interval, and that it converges to the *ex ante full-information pay-off*

$$\overline{R}(p) = p\, R^\mu(1) + (1-p) R^\mu(0)$$

as $r \downarrow 0$. This means that the optimal fees converge to

$$(\overline{M}_A(p), \overline{M}_B(p)) = \arg\min_{M_A,M_B} \frac{\overline{R}(p) - R(M_A, M_B, p)}{S(M_A, M_B)} \; , \tag{1.15}$$

which is the optimal policy of a platform provider maximizing its undiscounted transient payoff, that is, total expected revenue net of the full-information payoff that it will obtain in the long run; see Bolton and Harris (2000).

[8] As the admissible pair of fees (u_0, π_0) is clearly suboptimal (yielding zero revenue and zero information), the function S is indeed positive on the relevant domain.

Intuitively speaking, the lower the platform provider's discount rate, the greater is its incentive to learn, and the farther it might want to deviate from the myopic optimum. Experimentation is maximal when $r = 0$. Once we know the optimal strategy of the infinitely patient provider, therefore, we have fully characterized the range of experimentation in which an impatient provider will set his fees.

Studying the maximal experimentation strategy $(\overline{M}_A, \overline{M}_B)$ has the further advantage that it does not require computation of the value function for the maximization of transient payoffs.[9] While the system of first-order conditions for (1.15) in general does not permit explicit solutions, it is considerably easier to solve numerically than the differential equation for the value function under discounting. In the next subsection, we will take advantage of this to illustrate the maximal experimentation policy and the associated learning dynamics in a numerical example. Thereafter, we will briefly return to the limiting case $\overline{\pi} = 0$ which does permit a closed-form solution.

1.4.1 An example

We assume the following parameters: $u_0 = 0.4$, $\pi_0 = 0.1$, $\underline{u} = 0.1$, $\overline{u} = 0.9$, $\underline{\pi} = 0.1$, $\overline{\pi} = 0.2$, $\sigma_A = \sigma_B = 1$, $p_0 = 0.5$, and the "true" values are $(\overline{u}, \overline{\pi})$. These parameters translate into expected direct and indirect price effects of $\ell_0(p_0) = 1.11$, $\ell_A(p_0) = 0.60$, and $\ell_B(p_0) = 0.17$, respectively. In particular, the externality that side B is expected to exert on side A, $\ell_A(p_0)$, is assumed more than three times as large as the expected opposite externality, $\ell_B(p_0)$. Also note that $s_A = 0.06$, $s_{AB} = 0.24$, and $s_B = 1.04$, hence $s_A < s_{AB} < s_B$.

The optimal fees set by a myopic and an infinitely patient platform provider are depicted in Figure 1.4. It is straightforward to check that both myopically optimal fees are lower than the respective intrinsic values at all beliefs. In line with Proposition 3, the maximal experimentation policy reduces the fee on side B relative to the myopic benchmark at any

[9] This is crucial for the characterization of Markov perfect equilibria in Bolton and Harris (2000), for example.

non-degenerate belief.[10] The fee set on side A under the maximal experimentation policy is lower than the myopic benchmark at all beliefs below a threshold that approximately equals 0.275, and higher than the myopic benchmark at all beliefs above that threshold. Thus, in accordance with Propositions 2 and 4, the fee on side A is reduced when the externalities are of similar expected size (for beliefs close to 0, the large difference between $\overline{u}$ and $\overline{\pi}$ does not matter much), but is increased when $\ell_A(p)$, the expected strength of the externality that side A experiences, is sufficiently larger than $\ell_B(p)$, the expected strength of the externality that side A exerts.

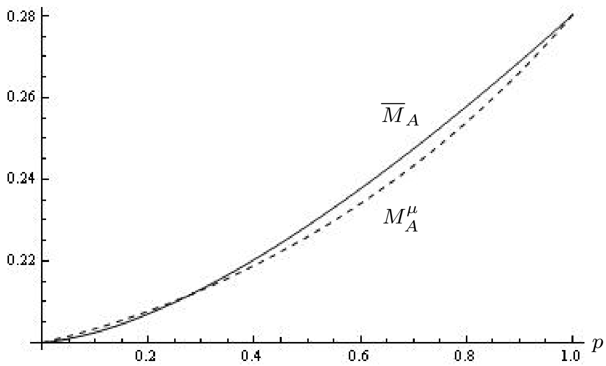

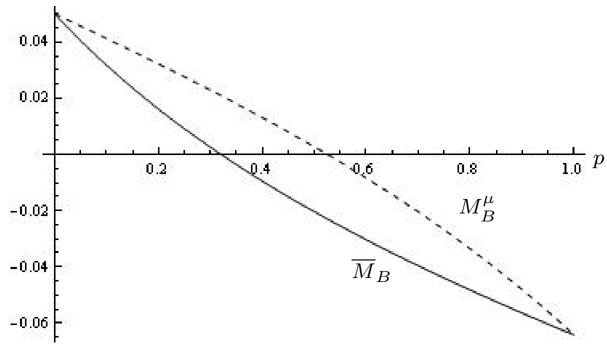

Figure 1.4: Optimal myopic fees (dashed line) and maximal experimentation fees (solid line) on market side A (left) and B (right) as a function of the belief.

Figure 1.5 illustrates that the infinitely patient provider learns faster – its beliefs converge more quickly to the true state.[11]

Figure 1.6 shows a sample path for the optimal fee on side A. At any given belief, the

[10]Note that for the given set of parameters, the optimal fee under full information on side B is actually negative, i.e., participants on side B receive a payment from the platform provider. Monetary payments to participants on one side may not always be feasible. However, as pointed out in the two-sided market literature, in-kind payments can often substitute for monetary payments.

[11]Simulations were carried out using Wolfram Mathematica 8. Normal shocks were generated by random draws from the normal distribution using the commands “RandomReal” and “NormalDistribution” with mean equal to 0 and variances equalling σ_A and σ_B respectively.

experimenting monopolist increases the fee relative to the myopic monopolist. In later periods, the increase is reinforced by the better information driving the fee towards the high optimum more rapidly. The evolution of optimal fees on side B is shown in Figure 1.7. Maximal experimentation fees are consistently below their myopic counterparts.

The expected per-period revenues depicted in Figure 1.8 show the advantages of each policy. While the myopic policy creates higher revenues in the very early periods, revenues in later periods are higher for the patient platform provider as its belief approaches the true state of the world more rapidly.

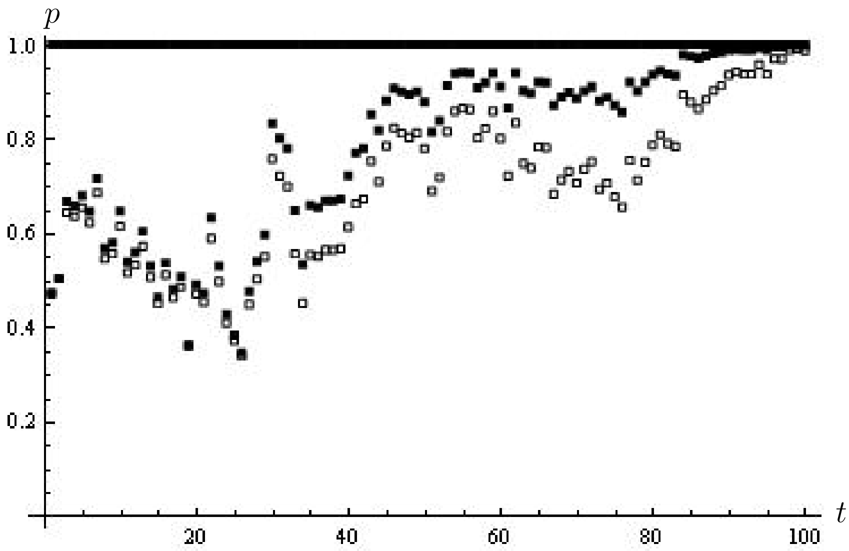

Figure 1.5: Evolution of beliefs for the myopic policy (white squares) and the infinitely patient policy (black squares), and true state (thick line).

1.4.2 A closed-form solution

We have seen in Proposition 4 above that, for vanishing externality parameter $\overline{\pi}$, the platform provider raises the fee on side A relative to the myopically optimal policy. The limiting case $\overline{\pi} = 0$ turns out to permit a closed-form solution for the maximal experi-

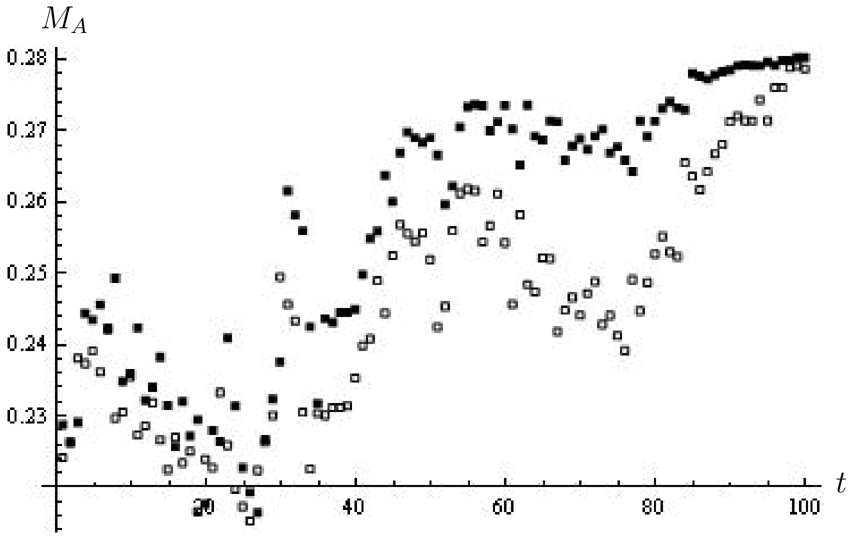

Figure 1.6: Evolution of fees on side A for the myopic policy (white squares) and the infinitely patient policy (black squares).

mentation policy.[12]

The myopically optimal fees in this case are

$$M_A^{\mu}(p) = \frac{2u_0 + u(p)\pi_0}{4 - u(p)^2},$$
$$M_B^{\mu}(p) = \frac{2\pi_0 - u(p)[u_0 + \pi_0 u(p)]}{4 - u(p)^2},$$

where $u(p) = \mathrm{E}^p[\tilde{u}] = p\overline{u} + (1-p)\underline{u}$. The myopic revenue is

$$R^{\mu}(p) = \frac{\pi_0^2 + u_0^2 + \pi_0 u_0 u(p)}{4 - u(p)^2}.$$

The quantity of information simplifies to $S(M_A, M_B) = \sigma_A^{-2}(\overline{u} - \underline{u})^2(\pi_0 - M_B)^2$, reflecting the fact that only the demand observed on side A is informative.

The minimum of $[\overline{R}(p) - R(M_A, M_B, p)]/(\pi_0 - M_B)^2$ is attained at

$$\overline{M}_A(p) = \frac{\pi_0 u_0 + 2\overline{R}(p)u(p)}{2\pi_0 + u_0 u(p)},$$
$$\overline{M}_B(p) = \pi_0 + \frac{u_0^2 - 4\overline{R}(p)}{2\pi_0 + u_0 u(p)}.$$

[12] As an example, consider readers whose utility of a magazine is independent of the amount of advertising.

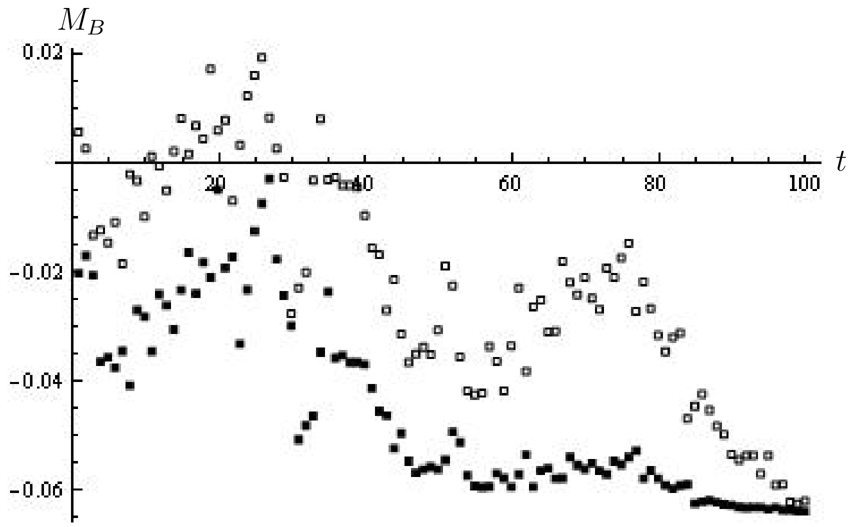

Figure 1.7: Evolution of fees on side B for the myopic policy (white squares) and the infinitely patient policy (black squares).

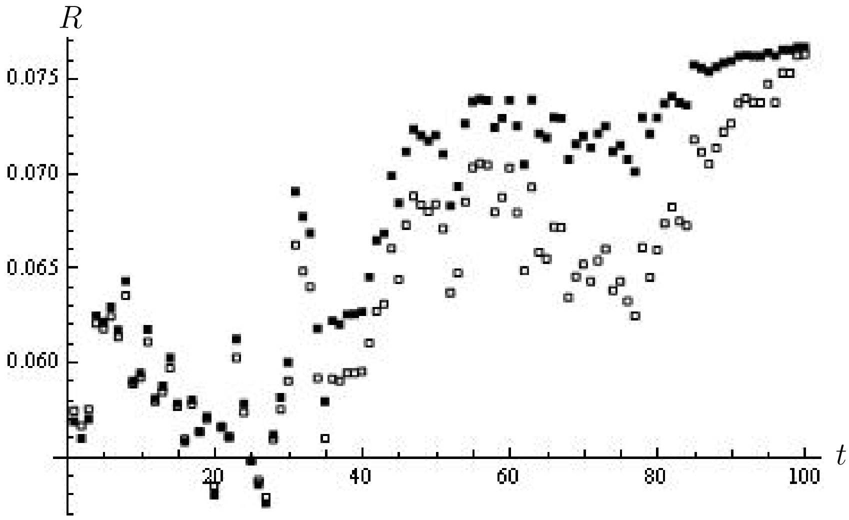

Figure 1.8: Evolution of expected per-period revenues for the myopic policy (white squares) and the infinitely patient policy (black squares).

Comparing these fees to the myopically optimal ones, we first see that

$$\overline{M}_B(p) - M_B^\mu(p) \quad = \quad \frac{4[R^\mu(p) - \overline{R}(p)]}{2\pi_0 + u_0 u(p)} \ .$$

As $\overline{R}(p) = p\, R^\mu(1) + (1-p)R^\mu(0)$ and R^μ is strictly convex, the right-hand side is negative for $0 < p < 1$. Thus, in line with Proposition 3, the infinitely patient platform provider will indeed decrease the fee that generates information.

On the other side of the market, we find

$$\overline{M}_A(p) - M_A^\mu(p) = \frac{2u(p)[\overline{R}(p) - R(p)]}{2\pi_0 + u_0 u(p)} = -\frac{u(p)}{2}[\overline{M}_B(p) - M_B^\mu(p)],$$

so for non-degenerate beliefs, there is a price increase relative to the myopic benchmark, as predicted by Proposition 4.

The expected quantity on side B clearly increases relative to the myopic optimum since the fee M_B goes down. Using the above expression for $\overline{M}_A(p) - M_A^\mu(p)$, one can additionally establish that the expected quantity on side A changes by $-\frac{u(p)}{2}[\overline{M}_B(p) - M_B^\mu(p)]$, which is again positive for non-degenerate beliefs. Hence, the platform provider also expects activity on this side to rise relative to the myopic optimum. Overall, therefore, optimal experimentation leads to uniform increases in expected quantities while price adjustments on the two sides go in opposite directions.

1.5 The Quantity-Setting Platform Provider

We now assume that the platform provider sets quantities. The quantity-setting assumption seems appropriate in real-world markets where capacity constraints matter. For instance, a shopping mall owner has to decide how much parking space and shop space to provide. If prices are market-clearing, this choice of capacities corresponds to quantity setting.

In standard monopoly, it does not matter (under certainty) whether a price or a quantity is chosen. In two-sided markets, setting quantities means that the platform directly controls the size of the externality, whereas a price setter does so only indirectly. This explains why the quantity-setting case is more tractable: there are no feedback effects to be taken into account when the quantity is changed on one side of the market. As we shall see below, this makes the information content of market observations additively separable across the two sides and implies unambiguous directions of experimentation.

Let the platform provider choose quantities $(n_A, n_B) \in \mathbb{R}_+^2$ and observe noisy signals of the prices

$$\begin{aligned} M_A(n_A, n_B, \tilde{u}) &= u_0 + \tilde{u} n_B - n_A, \\ M_B(n_A, n_B, \tilde{\pi}) &= \pi_0 + \tilde{\pi} n_A - n_B, \end{aligned}$$

where $\tilde{u} \in \{\underline{u}, \overline{u}\}$ and $\tilde{\pi} \in \{\underline{\pi}, \overline{\pi}\}$ with $0 < \underline{u} < \overline{u}$, $0 < \underline{\pi} < \overline{\pi}$ and $\overline{u} + \overline{\pi} < 2$. As we permit externality parameters exceeding 1, this is somewhat more general than what we assumed in the price-setting case.

We impose the natural restriction that the platform provider can only decide to sell non-negative quantities, while prices are not restricted. Negative prices are interpreted as subsidies to one side or (temporarily) both sides of the market, as discussed earlier. Note that the price on one side of the market does not depend on the externality parameter on the other side. However, as we assume perfect positive correlation between $\tilde{u}$ and $\tilde{\pi}$, any information gained on one side of the market immediately translates into a similar piece of information on the other side.[13]

As before, we write p for the subjective probability assigned to the realization $(\overline{u}, \overline{\pi})$. We maintain the assumption that costs are zero.

[13] Notably, all insights of this section carry over to the case of perfect *negative* correlation. Results only depend on expected externalities, exchanging the roles of $\overline{u}$ and $\underline{u}$ is unproblematic, therefore. As to Propositions 6–7 below, it suffices that signal-to-noise ratios coincide in absolute value.

1.5.1 Revenues and beliefs

In every period $t \in [0, \infty[$, the platform provider chooses quantities (n_A^t, n_B^t) and then observes the increments $M_A(n_A^t, n_B^t, \tilde{u}, \tilde{\pi})\, dt + \theta_A dW_A^t$ and $M_B(n_A^t, n_B^t, \tilde{u}, \tilde{\pi})\, dt + \theta_B dW_B^t$ of two cumulative price processes where W_A and W_B are independent standard Brownian motions and the constants θ_A and θ_B are positive. The resulting revenue increment at date t is

$$dR_t = n_A^t \left[M_A(n_A^t, n_B^t, \tilde{u})\, dt + \theta_A dW_A^t \right] + n_B^t \left[M_B(n_A^t, n_B^t, \tilde{\pi})\, dt + \theta_B dW_B^t \right].$$

With the notation

$$\begin{aligned} u(p) &= p\,\overline{u} + (1-p)\underline{u}\,, \\ \pi(p) &= p\,\overline{\pi} + (1-p)\underline{\pi} \end{aligned}$$

for the expected externalities, and

$$R(n_A, n_B, p) = n_A \left[u_0 + u(p)\, n_B - n_A \right] + n_B \left[\pi_0 + \pi(p)\, n_A - n_B \right]$$

for the expected per-period revenue, the platform provider's total expected payoff is

$$\mathrm{E}^{p_0} \left[\int_0^\infty r e^{-rt} R(n_A^t, n_B^t, p_t)\, dt \right].$$

The expected revenue R depends on the expected externalities only through the term $[u(p)+\pi(p)]n_A n_B$, so only the *sum* of the externalities matters here. As $|u(p)+\pi(p)| < 2$, moreover, R is strictly concave in (n_A, n_B). The myopically optimal quantities are

$$\begin{aligned} n_A^\mu(p) &= \frac{2u_0 + \pi_0[u(p)+\pi(p)]}{4 - [u(p)+\pi(p)]^2}\,, \\ n_B^\mu(p) &= \frac{2\pi_0 + u_0[u(p)+\pi(p)]}{4 - [u(p)+\pi(p)]^2}\,. \end{aligned}$$

They exhibit a symmetric structure with interchanged intrinsic platform values. If these platform values coincide, myopically optimal quantities are the same on both sides.

The corresponding expected prices for each group, however, depend on the specific externality the other group is exerting. They are given by

$$\begin{aligned} M_A^\mu(p) &= \frac{\pi_0[u(p)-\pi(p)]+u_0(2-\pi(p)[u(p)+\pi(p)])}{4-[u(p)+\pi(p)]^2}, \\ M_B^\mu(p) &= \frac{u_0[\pi(p)-u(p)]+\pi_0(2-u(p)[u(p)+\pi(p)])}{4-[u(p)+\pi(p)]^2}. \end{aligned}$$

The expected current revenue from the myopically optimal quantities is

$$R^\mu(p) = M_A^\mu(p)n_A^\mu(p) + M_B^\mu(p)n_B^\mu(p) = \frac{u_0^2+\pi_0^2+u_0\pi_0[u(p)+\pi(p)]}{4-[u(p)+\pi(p)]^2}.$$

To describe the law of motions of beliefs, we define the strictly convex function

$$\Sigma(n_A, n_B) = \rho_A n_A^2 + \rho_B n_B^2,$$

where the constants

$$\rho_A = \left(\frac{\overline{u}-\underline{u}}{\theta_A}\right)^2, \quad \rho_B = \left(\frac{\overline{\pi}-\underline{\pi}}{\theta_B}\right)^2$$

are the squares of the marginal signal-to-noise ratios.

Lemma 4 *The beliefs of the quantity-setting platform provider evolve according to*

$$dp_t \sim N\Big(0, p_t^2(1-p_t)^2\,\Sigma(n_A^t, n_B^t)\,dt\Big)$$

Proof: The proof is similar to the price-setting case and therefore omitted. □

Complete learning in the long-run follows from the same arguments as in the price-setting scenario (see Lemma 2 above). As Σ is increasing in both n_A and n_B, moreover, we obviously have

Lemma 5 *For the quantity-setting platform provider, a quantity increase on either side of the market increases the information content of observed prices, whereas a quantity decrease reduces it.*

Finally, we note that the marginal informational impact of adjusting the quantity on one side of the market does *not* depend on the quantity set on the other.

1.5.2 Optimal quantities

Under discounting at rate $r > 0$, the Bellman equation is

$$v(p) = \max_{n_A, n_B} \left\{ R(n_A, n_B, p) + \frac{p^2(1-p)^2}{2r} \, \Sigma(n_A, n_B) \, v''(p) \right\} \; .$$

The maximand is again the sum of a strictly concave quadratic function and a strictly convex one. A simpler version of the argument given in the price-setting case shows that the shadow price of information, $V(p) = p^2(1-p)^2 v''(p)/2r$, is again sufficiently small to make the combined quadratic function strictly concave at all beliefs (we omit the details).

Solving the first-order conditions for optimal quantities and suppressing the dependence on the belief p, we obtain

$$\begin{aligned} n_A^* &= n_A^\mu + \frac{2V}{\chi(V)} \left\{ 2(1-\rho_B V)\rho_A n_A^\mu + 2(u+\pi)\rho_B n_B^\mu \right\}, \\ n_B^* &= n_B^\mu + \frac{2V}{\chi(V)} \left\{ 2(1-\rho_A V)\rho_B n_B^\mu + 2(u+\pi)\rho_A n_A^\mu \right\}, \end{aligned}$$

where

$$\chi(V) = 4(1-\rho_A V)(1-\rho_B V) - (u+\pi)^2$$

is the determinant of the Hessian matrix of $R + V\Sigma$. Strict concavity of this function means $1 - \rho_A V > 0$ and $\chi(V) > 0$, which in turn implies $1 - \rho_B V > 0$. As an immediate consequence, we get

Proposition 5 *At any non-degenerate belief, the quantity-setting platform provider chooses quantities above the myopic benchmark on both sides of the market.*

The intuition behind this result is simple. As the information content of observed prices is increasing in quantities, the optimal deviation from the myopic benchmark must entail a higher quantity on at least one side of the market. This raises the marginal revenue on the other side of the market without affecting the marginal informational benefit of adjusting the quantity there. It is optimal, therefore, to set a quantity above the myopic level on that side as well.

In the price-setting scenario, by contrast, lowering the fee on one side of the market has an ambiguous effect on the incentives to lower the fee on the other side because the cross-partial derivative of the quantity of information with respect to these fees is not zero and has the opposite sign to the respective derivative of expected current revenue.

1.5.3 Maximal experimentation

The maximal experimentation strategy is given by

$$(\overline{n}_A(p), \overline{n}_B(p)) = \arg \min_{n_A, n_B} \frac{\overline{R}(p) - R(n_A, n_B, p)}{\Sigma(n_A, n_B)} \ .$$

where $\overline{R}(p) = p\,R^\mu(1) + (1-p)R^\mu(0)$ is once more the expected full-information payoff. In general, the associated first-order conditions involve third-order polynomials in n_A and n_B. Owing to the simpler structure of the quantity-setting scenario, however, it is easier to obtain closed-form solutions than in the price-setting case, for example by assuming symmetric signal-to-noise ratios.[14]

Proposition 6 *Suppose that $\rho_A = \rho_B$ and $u_0 \neq \pi_0$. Then the quantities set by an infinitely patient platform provider are*

$$\begin{aligned}
\overline{n}_A(p) &= \frac{1}{2(u_0^2 - \pi_0^2)[u(p)+\pi(p)]} \Bigg\{ \pi_0(\pi_0^2 + u_0^2) + 4\overline{R}(p) u_0 [u(p)+\pi(p)] \\
&\qquad - \pi_0 \sqrt{(u_0^2 - \pi_0^2)^2 + (2u_0\pi_0 + 4\overline{R}(p)[u(p)+\pi(p)])^2} \Bigg\}, \\
\overline{n}_B(p) &= \frac{1}{2(u_0^2 - \pi_0^2)[u(p)+\pi(p)]} \Bigg\{ (-u_0(\pi_0^2 + u_0^2) - 4\overline{R}(p) \pi_0 [u(p)+\pi(p)] \\
&\qquad + u_0 \sqrt{(u_0^2 - \pi_0^2)^2 + (2u_0\pi_0 + 4\overline{R}(p)[u(p)+\pi(p)])^2} \Bigg\}.
\end{aligned}$$

[14] With quantities as the choice variables, it is less interesting to consider the limiting case of no uncertainty about the externality on one side of the market. If $\underline{\pi} = \overline{\pi}$, for instance, any deviation from the expected price on side B must be attributed to noise and is, thus, uninformative. The platform can then only experiment on side A, and only by adjusting the quantity n_B. This situation is isomorphic to the one analyzed in Keller and Rady (1999).

Proof: See the appendix. □

The reason why these quantities do not depend on the common marginal signal-to-noise ratio is simple. For $\rho_A = \rho_B = \rho$, the information content of observed prices simplifies to $\Sigma(n_A, n_B) = \rho\,[n_A^2 + n_B^2]$, so the maximal experimentation strategy minimizes $(\overline{R}(p) - R(p))/(n_A^2 + n_B^2)$. Note that for $\pi_0 > u_0$, both numerator and denominator of $\overline{n}_A(p)$ and $\overline{n}_B(p)$ are negative, so the quantities remain positive. The knife-edge case $u_0 = \pi_0$ will be covered later.

The expected prices $\overline{M}_A(p)$ and $\overline{M}_B(p)$ given the quantities $\overline{n}_A(p)$ and $\overline{n}_B(p)$ are straightforward to calculate. Comparing them with the myopic optimum confirms what we have already seen in the price-setting model: even if the externality parameters $\overline{u}$ and $\overline{\pi}$ are both smaller than 1, there are parameter constellations such that, on one side of the market, the expected price for the patient platform provider is higher than the myopic benchmark, as exemplarily shown for $\overline{M}_A$ in Figure 1.9.

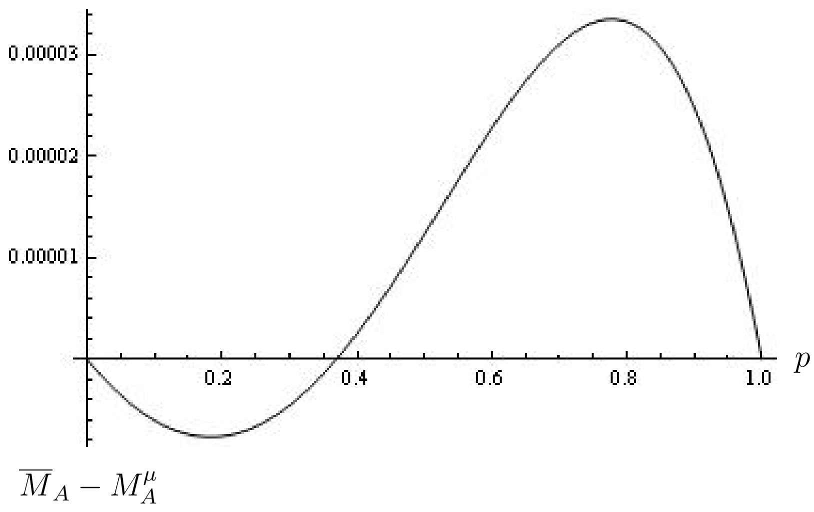

Figure 1.9: Difference between the expected prices induced by a myopic and an infinitely patient platform provider as a function of beliefs for $u_0 = 0.1$, $\pi_0 = 0.7$, $\underline{u} = 0.8$, $\overline{u} = 0.9$, $\underline{\pi} = 0.1$, $\overline{\pi} = 0.2$, $\theta_A = \theta_B = 1$.

The intuition behind this finding is the one we already gave for Proposition 4 in the price-setting case. If side A is expected to exert a relatively weak externality, i.e., if $\pi(p)$ is small relative to $u(p)$, the platform provider optimally learns by strongly increasing the number of participants on side B, and recoups part of the resulting surplus by inducing a higher than myopically optimal price on side A.

Maintaining symmetric signal-to-noise ratios, we further assume now that the intrinsic value of the platform is the same for all users, i.e., $u_0 = \pi_0$. This admittedly rather strong assumption seems appropriate in a number of examples, such as night clubs and matching agencies.[15] It simplifies the expressions for the optimal quantities considerably.

Proposition 7 *Suppose that $\rho_A = \rho_B$ and $u_0 = \pi_0 = c_0$. Then the optimal policy of an infinitely patient quantity-setting platform provider is symmetric across market sides and linear in the current belief:*

$$\overline{n}_A(p) = \overline{n}_B(p) = \overline{n}(p) = \frac{\overline{R}(p)}{c_0} = c_0 \left[\frac{p}{2 - (\overline{u} + \overline{\pi})} + \frac{1-p}{2 - (\underline{u} + \underline{\pi})} \right].$$

Proof: See the appendix. □

The intuition for the symmetry of the optimal quantities is as follows. With identical intrinsic platform values, the myopically optimal quantities are symmetric. With identical signal-to-noise ratios, moreover, the incentive to deviate from the myopic optimum is the same in both quantity dimensions.

The linearity of the maximal experimentation policy makes it easy to visualize the range of quantity experimentation; see Figure 1.10. It is the area bounded below by the myopic policy and above by the line joining the quantities that are optimal under full information.

Expected prices need not be symmetric. They are

$$\begin{aligned} \overline{M}_A(p) &= c_0 + [u(p) - 1]\, \overline{n}(p), \\ \overline{M}_B(p) &= c_0 + [\pi(p) - 1]\, \overline{n}(p). \end{aligned}$$

[15] It is clearly less appropriate in other examples, such as merchants and customers in the credit card market.

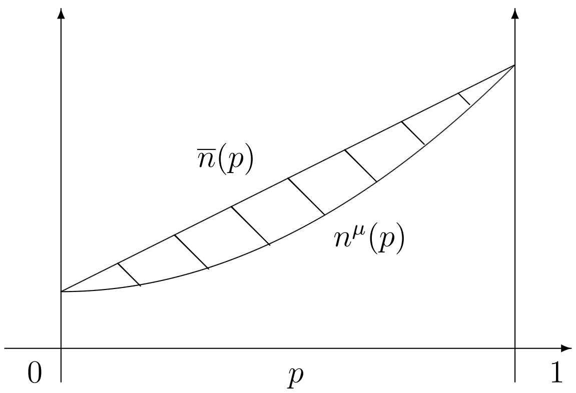

Figure 1.10: Range of quantity experimentation for symmetric signal-to-noise ratios and symmetric intrinsic values.

As $u(p) + \pi(p) < 2$, either both expected prices are lower than the intrinsic platform value, or one is lower and the other one higher. The ordering of expected prices depends on the size of the externalities and on the current belief, and may change with beliefs. Let $\underline{u} < \underline{\pi} < \overline{\pi} < \overline{u}$, for example. For high values of p, then, $u(p)$ will exceed $\pi(p)$ and side A will have to pay a higher price in expectation than side B, while for low values of p the reverse is true.

As to the comparison with the myopic benchmark, we have

Corollary 1 *Under the assumptions of Proposition 7, the expected price induced by an infinitely patient quantity-setting platform provider exceeds its myopically optimal counterpart on a given side of the market if and only if the expected externality that this side experiences is greater than* 1*.*

PROOF: See the appendix. □

The optimal and the myopic expected prices coincide at the beliefs 0 and 1 or if the expected externality equals 1. As $u(p) + \pi(p) < 2$, this of course implies that at any time at most one expected price can exceed the myopically optimal level. It also implies that for the "standard" case of both externalities smaller than 1, both expected prices will decrease relative to the myopic benchmark.

1.6 Conclusion

We have studied optimal behavior of a monopolistic platform provider in a two-sided market with uncertainty about the strength of interaction between the two sides. The platform provider either chooses prices or quantities (i.e., participation levels). The demand externalities considered are linear on both sides. Fees are charged for participation in the market, but not per transaction. In this respect, our setting follows the monopoly setting analyzed in Armstrong (2006).

When the platform provider faces uncertainty about the size of the externality and wants to maximize its expected lifetime profits, it faces the basic trade-off between the conflicting aims of maximizing current payoff and maximizing the information content of the signals it observes. We have characterized the optimal policies depending on how much weight the platform provider assigns to future profit. If it does not put any weight on the future ($r = \infty$), it chooses the myopically optimal actions given its current belief. As there is no potentially confounding action, even the myopic platform provider continuously accumulates information about the true state of the world and will, in the limit as time tends to infinity, almost surely learn the true state.

If the platform provider puts some weight on the future, it will deviate from the shortsighted policy and invest in learning. The upper bound on such experimentation is given by the optimal policy of an infinitely patient platform provider ($r = 0$).

The effect of experimentation on (expected) prices is ambiguous. Depending on the parameter constellation, either both prices will be lower than in the myopic benchmark or one price will be above and one price below the myopically optimal prices. The price on one side of the market will go up if the externality this side is exerting is weak while the externality it is experiencing is strong. The higher price recoups part of the surplus created by the higher participation on the other side of the market.

Our analysis concerns an unrestricted monopoly platform. Future work may want to look at markets with multiple differentiated platforms. As a starting point, it would be interesting to analyze duopoly experimentation in a two-sided market in which there is single-homing on both sides and full observability of actions and outcomes. In such a duopoly, a participant acquired by one platform provider is a participant lost for the competitor. Owing to indirect network effects, this makes demand more sensitive to price changes than demand in the monopoly setting with a fixed outside option that has been analyzed in this chapter. Therefore, one may conjecture that gaining information about the size of the network effect becomes more important. As has been pointed out in the literature on duopoly experimentation (e.g., Mirman et al. 1994, Harrington Jr. 1995, Keller and Rady 2003), however, the public information generated by market signals may have a negative value, in which case the duopolists have an incentive to generate *less* information than in the myopic equilibrium.

Suppose, for instance, that market participation is perfectly price-inelastic, as is the case in the Hotelling-type model introduced by Armstrong (2006). Then, learning does not increase future equilibrium profits in expectation because profits are linear in beliefs. Since deviations from the myopic best-response are costly, we conjecture that patient platform operators do not behave differently from infinitely impatient ones, and learn only passively. The duopoly setting merits further, more general investigation, and it would be interesting to understand the effect of the degree of differentiation on experimentation in a two-sided market.

Another interesting extension is to consider a market for two (or more) goods that are complements. Specifically, suppose that demands are linked through positive network effects. Here we have in mind a situation in which a monopolist sells a product (or technologically related products) to two distinct and distinguishable consumer groups (i.e., the monopolist can practice third-degree price discrimination). If consumers in each group care directly or indirectly about the sum of the total number of buyers in both groups (e.g., because a larger production volume increases average product quality through learning-by-doing), we can rewrite this as a demand system with indirect network effects. Thus, our analysis can possibly be extended to capture experimentation in markets with complementary goods.

A1 Appendix

Proof of Lemma 1

Given a pair of prices (M_A, M_B), the observed quantity increments are

$$\begin{pmatrix} dN_A \\ dN_B \end{pmatrix} = \begin{pmatrix} \tilde{n}_A \\ \tilde{n}_B \end{pmatrix} dt + \begin{pmatrix} \sigma_A & 0 \\ 0 & \sigma_B \end{pmatrix} \begin{pmatrix} dZ_A \\ dZ_B \end{pmatrix}$$

with $\tilde{n}_A = n_A(M_A, M_B, \tilde{u}, \tilde{\pi})$ and $\tilde{n}_B = n_B(M_A, M_B, \tilde{u}, \tilde{\pi})$.

Given the subjective probability p currently assigned to the state $(\overline{u}, \overline{\pi})$, the vector of expected demands is

$$\begin{pmatrix} \mathrm{E}^p\left[\tilde{n}_A\right] \\ \mathrm{E}^p\left[\tilde{n}_A\right] \end{pmatrix} = p \begin{pmatrix} \overline{n}_A \\ \overline{n}_B \end{pmatrix} + (1-p) \begin{pmatrix} \underline{n}_A \\ \underline{n}_B \end{pmatrix}$$

with $\overline{n}_A = n_A(M_A, M_B, \overline{u}, \overline{\pi})$ etc.

According to Liptser and Shiryayev (1977), the infinitesimal change in beliefs is given by

$$dp = p \begin{pmatrix} \overline{n}_A - \mathrm{E}^p\left[\tilde{n}_A\right] \\ \overline{n}_B - \mathrm{E}^p\left[\tilde{n}_B\right] \end{pmatrix} \begin{pmatrix} \sigma_A^{-1} & 0 \\ 0 & \sigma_B^{-1} \end{pmatrix} \begin{pmatrix} d\overline{Z}_A \\ d\overline{Z}_B \end{pmatrix}$$

where

$$\begin{pmatrix} d\overline{Z}_A \\ d\overline{Z}_B \end{pmatrix} = \begin{pmatrix} \sigma_A^{-1} & 0 \\ 0 & \sigma_B^{-1} \end{pmatrix} \begin{pmatrix} dN_A - \mathrm{E}^p\left[\tilde{n}_A\right] \\ dN_B - \mathrm{E}^p\left[\tilde{n}_B\right] \end{pmatrix}$$

is the increment of a standard two-dimensional Brownian motion relative to the platform provider's information filtration.

Simplifying the expression for dp, we obtain

$$dp = p(1-p)(\overline{n}_A - \underline{n}_A)\sigma_A^{-1} d\overline{Z}_A + p(1-p)(\overline{n}_B - \underline{n}_B)\sigma_B^{-1} d\overline{Z}_B.$$

As $d\overline{Z}_A$ and $d\overline{Z}_B$ are normally distributed with mean zero and variance dt, and the infinitesimal covariance $< d\overline{Z}_A, d\overline{Z}_B >$ is zero, the change in beliefs dp is normally distributed with mean zero and variance $p^2(1-p)^2(\overline{n}_A - \underline{n}_A)^2\sigma_A^{-2}\, dt + p^2(1-p)^2(\overline{n}_B - \underline{n}_B)^2\sigma_B^{-2}\, dt$. $\square$

Proof of Lemma 3

We wish to show that in the region where the information content of quantities is increasing in a fee, the expected quantity on at least one side of the market must be negative.

For a partial derivative of S to be positive, at least one of the differences $\overline{n}_A - \underline{n}_A$ or $\overline{n}_B - \underline{n}_B$ has to be negative. This in turn is equivalent to at least one of the following inequalities holding:

$$M_B > \pi_0 + \frac{d_0}{d_A}(u_0 - M_A), \tag{A1.1}$$
$$M_B > \pi_0 + \frac{d_B}{d_0}(u_0 - M_A). \tag{A1.2}$$

For the two expected demands to be non-negative, it is necessary that both $\overline{n}_A$ and $\overline{n}_B$ be non-negative. This requires the following inequalities to hold:

$$M_B \leq \pi_0 + \frac{1}{\overline{u}}(u_0 - M_A), \tag{A1.3}$$
$$M_B \leq \pi_0 + \overline{\pi}(u_0 - M_A). \tag{A1.4}$$

Comparing the coefficients of $u_0 - M_A$ on the right-hand sides of these four inequalities, we see that for $M_A > u_0$, (A1.3) contradicts both (A1.1) and (A1.2), while (A1.4) does so for $M_A < u_0$. For $M_A = u_0$ the contradiction is obvious. $\square$

Strict concavity of the maximand in the Bellman equation

Fixing a belief p and a shadow price of information $V = p^2(1-p)^2 v''(p)/2r$, we write the maximand in the Bellman equation (1.10) as $R(M_A, M_B, p) + V\,S(M_A, M_B)$ and compute its Hessian, suppressing the variable p from now on:

$$\mathcal{H}(V) = \begin{pmatrix} -2\ell_0 & -(\ell_A + \ell_B) \\ -(\ell_A + \ell_B) & -2\ell_0 \end{pmatrix} + V \begin{pmatrix} 2s_A & 2s_{AB} \\ 2s_{AB} & 2s_B \end{pmatrix}.$$

Its determinant is

$$h(V) = 4(\ell_0 - s_A V)(\ell_0 - s_B V) - (\ell_A + \ell_B - 2s_{AB}V)^2.$$

For global strict concavity of $R+VS$, we wish to show that $\ell_0 - s_A V > 0$ and $h(V) > 0$.

Since the value function, and hence the maximum of $R + VS$, is bounded, the latter is bounded from above along any ray $\{(M_A, M_B) : M_A = u_0 - x,\ M_B = \pi_0 - \beta x,\ x \geq 0\}$ with $\beta \geq 0$ (note that these fees are all admissible). As

$$R(u_0 - x, \pi_0 - \beta x) + VS(u_0 - x, \pi_0 - \beta x) = \Big\{u_0\left[\ell_0 + \ell_A \beta\right] + \pi_0\left[\ell_0 \beta + \ell_B\right]\Big\}\, x - q(\beta) x^2$$

with the quadratic function

$$q(\beta) = \ell_0 - s_A V + (\ell_A + \ell_B - 2s_{AB}V)\beta + (\ell_0 - s_B V)\beta^2,$$

this implies that q is positive on $[0, \infty[$. Setting $\beta = 0$ yields $\ell_0 - s_A V > 0$.

Next, let $V > (\ell_A + \ell_B)/2s_{AB}$, so that $q'(0) < 0$. As a consequence, $\ell_0 - s_B V > 0$ since q would become negative at high β otherwise. Moreover, q assumes its minimum at

$$\beta^* = \frac{2s_{AB}V - \ell_A - \ell_B}{2(\ell_0 - s_B V)} > 0.$$

This minimum equals

$$q(\beta^*) = \ell_0 - s_A V - \frac{(2s_{AB}V - \ell_A - \ell_B)^2}{4(\ell_0 - s_B V)} = \frac{h(V)}{4(\ell_0 - s_B V)}\,,$$

implying $h(V) > 0$ and concavity of $R + VS$.

As V multiplies the strictly convex function S, concavity of $R + VS$ now also follows for shadow prices $V \leq (\ell_A + \ell_B)/2s_{AB}$. $\square$

Proof of Propositions 6 and 7

For arbitrary ρ_A and ρ_B, the first-order conditions for the fees $\overline{n}_A(p)$ and $\overline{n}_B(p)$ can be written as

$$\begin{aligned}
&(u_0 + [u(p) + \pi(p)]n_B - 2n_A)(\rho_A n_A^2 + \rho_B n_B^2) \\
&\quad + \ 2\rho_A n_A \Big[\overline{R}(p) - (u_0 + u(p)n_B - n_A)\, n_A - (\pi_0 + \pi(p)n_A - n_B)\, n_B\Big] \ = \ 0, \\
&(\pi_0 + [u(p) + \pi(p)]n_A - 2n_B)(\rho_A n_A^2 + \rho_B n_B^2) \\
&\quad + \ 2\rho_B n_B \Big[\overline{R}(p) - (u_0 + u(p)n_B - n_A)\, n_A - (\pi_0 + \pi(p)n_A - n_B)\, n_B\Big] \ = \ 0.
\end{aligned}$$

For $\rho_A = \rho_B$, this system simplifies to

$$\begin{aligned}(u_0 + [u(p) + \pi(p)]n_B)(n_B^2 - n_A^2) + 2(\overline{R}(p) - n_B \pi_0)\, n_A &= 0 \\ (\pi_0 + [u(p) + \pi(p)]n_A)(n_A^2 - n_B^2) + 2(\overline{R}(p) - n_A u_0)\, n_B &= 0.\end{aligned}$$

For $u_0 \neq \pi_0$, the pair of quantities stated in Proposition 6 constitutes the unique solution to these equations. For $u_0 = \pi_0 = c_0$, setting both quantities equal to $\overline{R}(p)/c_0$ solves the system. □

Proof of Corollary 1

For $u_0 = \pi_0 = c_0$, the myopically optimal expected price on side A simplifies to

$$M_A^\mu(p) = \frac{c_0[1 - \pi(p)]}{2 - [u(p) + \pi(p)]}\ ,$$

so the price difference $\overline{M}_A(p) - M_A^\mu(p)$ has the same sign as

$$1 + (u(p) - 1)\left[\frac{p}{2 - (\overline{u} + \overline{\pi})} + \frac{1 - p}{2 - (\underline{u} + \underline{\pi})}\right] - \frac{1 - \pi(p)}{2 - [u(p) + \pi(p)]}\ .$$

Multiplying with $2 - [u(p) + \pi(p)]$ and simplifying, we see that this in turn has the same sign as

$$(u(p) - 1)\left\{(2 - [u(p) + \pi(p)])\left[\frac{p}{2 - (\overline{u} + \overline{\pi})} + \frac{1 - p}{2 - (\underline{u} + \underline{\pi})}\right] - 1\right\}.$$

The expression in curly brackets is strictly concave in p; as it vanishes at $p = 0$ and $p = 1$, it is positive for $0 < p < 1$. The proof for side B is analogous. □

Chapter 2

Never Change a Winning Team?

2.1 Introduction

One of the disadvantages of team production is that individual contributions to team output, be it effort or ability, usually are not fully observable. But individual characteristics often are of high interest for employment, salary or promotion decisions as well as for efficient allocation of resources and optimal provision of incentives. Therefore an employer might not only be interested in maximizing current team output, but also in inferring information on individuals. While it is often too costly to fully observe a team during production, employers can usually decide on the team composition as well as on whom they would like to keep and whom to replace.

This chapter focusses on the possibility of inferring ability from team output. While individual ability is unobservable, changing team composition will provide a useful tool to learn about ability on a long term basis.

Team production has gained importance in various areas of economic and non-economic life. In many companies it has become best practice to establish project teams of experts from various fields to perform certain tasks in order to benefit from the complementarities of the different specialists. Law and consulting firms usually send teams with varying experience levels to their clients to enforce learning effects especially for less experienced team members. In sports, teams can be mandatory according to the rules of the game and one often observes successes of collectives against collections of stars owing to match specific productivity. Even in individualist's sports like swimming or athletics a lot of training is done in teams to enhance motivation. Last but not least in research the number of co-authored papers has highly increased in the last three decades: Kim et al. (2009) report an increase of co-authored papers in 41 top academic journals from roughly 40% in the 1970's to over 70% in 2004.

We start by considering an employer hiring one team of two workers with unknown ability which might either be high or low. The employer after each period of production updates her belief on each worker's ability and can decide which worker(s) to keep and which to lay off. Output is increasing in both workers' abilities, thus the goal is to find two employees of high ability, although the employer might settle for one worker of high ability if those are in short supply, the relative benefit of a second high type is small, or there is a low weight on future output.

A worker of known low ability still can serve as a perfect identifier for his co-worker's ability and thus might temporarily be kept for this reason. For high discount factors, i.e., when the future matters more and identification becomes more important, keeping a "known fool", a worker of known low ability, might be advantageous to the employer when keeping everything else fixed. However, it turns out that for any policy keeping a known fool there always exists a policy not keeping this worker — and potentially differing in other characteristics — that does better in terms of expected revenue. Results are robust to sub- and supermodular production, i.e., to the relative benefit a first and a second high ability worker adds. In addition, for the majority of parameter constellations (essentially for high enough discount factors) the employer will end up with a team of two high ability workers, to which we refer as a "winning team".

The second part of the chapter considers an entrepreneur that already has found such a winning team and now enlarges the firm by hiring a second team. This now introduces the additional strategic option to split the winning team and rematch it with the new workers. If production is (sub-)modular this turns out to always be the dominant strategy, as in this case the information and revenue effect go in the same direction: Splitting the winning team will resolve all uncertainty and at the same time generate higher expected profits. For supermodular production safe high output of the winning team might become too valuable to forego for information, thus policies not or not immediately splitting the winning team will become optimal. Still, the result that it is never optimal to keep a known low ability workers persists. Further, in contrast to the one-team setting the employer will almost surely perfectly learn all workers' abilities for all parameter constellations in this set-up and thus end up with four workers of high ability with probability one.

Regarding existing literature, one of the first papers to address the problem of determining individual contribution to team output is Alchian and Demsetz (1972) who identify this inseparability as a key factor to the question of why firms exist as raised, e.g., by Coase (1937). Alchian and Demsetz already hinge on the possibility of revising team composition to infer capabilities of team members but do not further elaborate on this. Dynamic learning in team production has been considered by Meyer (1994). She provides a two-period overlapping generations model including uncertainty, where the employer can decide how to match two teams of senior and junior employees when juniors can split their capacity. However, in her setting production is modular, hence overall output is fixed and thus there is no trade-off between income and information. Her results thus favor the decision generating most information. Further, in contrast to Meyer (1994) this chapter endogenizes the decision if and when to replace an employee in an infinite horizon set-up. Bar-Isaac (2007) considers a situation where agents with differing experience can decide whether to team up with each other. This can serve as a commitment device to induce effort. Breton et al. (2003) also consider endogenous team choice when agents might differ in type or experience/reputation, deriving conditions when young workers might prefer to match intra- or intergenerationally. Che and Yoo (2001) emphasize the reputation effect generated by the repeated collaboration of team members and derive that team work can be optimal under the aspect of organizational design as it can create implicit incentives from the employees' repeated interaction. The question of how to optimally compose teams as treated in the second part of the chapter is also related to the literature on optimal organizational design, see for, e.g., Sah and Stiglitz (1986) and Bolton and Dewatripont (1994). An introduction into optimal decision problems when organizations consist of teams can be found in Marschak and Radner (1972).

In terms of modeling, the setting is close to Anderson and Smith (2010) who also consider team production and ability inference. While their focus is on optimal assortative matching of the complete continuum of workers, this chapter introduces a scarcity on the number of teams hired from the pool and thus emphasizes the strategic employment decision within the optimal sampling problem, as any acquired information is lost when an employee is laid off.

In our model we explicitly abstract from optimal effort implementation. The strand of literature on moral hazard problems and incentives in team production is already extensive, starting with Holmström (1982). Jeon (1996) for example considers an extension of Holmström's (1999) managerial incentive problem with career concerns to teams of two managers and shows that the result of managers exerting the expected level of effort in equilibrium is sustainable. For a discussion of how to infer and induce effort in multiple agent production see also, e.g., McAfee and McMillan (1991) or Mookherjee (1984).

The remainder of the chapter proceeds as follows: Section 2.2 introduces the set-up and describes the optimal hiring policy of an entrepreneur employing one team. Section 2.3 extends the setting to a growing firm employing a known good team and hiring a second team, thus offering team composition as an additional decision variable. Section 2.4 discusses finite time-horizons, firing costs, and additional uncertainty as extensions to the one-team setting, while the last section concludes.

2.2 Set-up: Employing One Team

Consider a continuum of workers of finite measure. Workers only differ by their exogenous, unobservable ability. A share p of workers in the continuum is of high ability denoted by

H and the remaining share $1-p$ is of low ability denoted by L.

A risk-neutral entrepreneur (to which we will also refer as employer or decision maker) wants to hire exactly two workers for production according to a deterministic production function that is symmetric and increasing in both workers' abilities. Thus output might be low (Y_L), medium (Y_M), or high (Y_H). The entrepreneur can only observe the output from production and wants to maximize her expected lifetime utility in a discrete-time, infinite-horizon setting in which the future is discounted by $\beta < 1$:

$$\max E\Big[\sum_{t=0}^{\infty} \beta^t Y_t\Big], \; Y_t \in \{Y_L, Y_M, Y_H\},$$

and the maximization is over the employment policies which will be specified in detail in the next section.

We normalize costs to zero and do not consider wages. A justification for this is that one might interpret Y as the return to capital after wages, i.e., the production surplus. If ability is unknown to the workers as well or not verifiable one can think of the workers being paid their average marginal contribution after output is realized. Another interpretation is to see ability as firm-specific productivity and allocate full bargaining power to the firm, such that workers are paid fixed compensations equalling their outside opportunity. Also, as already mentioned above, we do not explicitly consider effort.

2.2.1 Beliefs and strategies

Although the entrepreneur cannot observe individual ability she can, after observing output, decide in any period whether to keep the workers or whether to replace one or both of the workers by someone from the pool of unemployed.

We will argue in the following that future choices of an optimizing entrepreneur are already fully determined by her first-period decisions and that optimal choices can thus be identified with optimal first-period choices. In order to derive this we first describe how beliefs evolve.

Before the first round of production prior beliefs of having found a high ability worker correspond to the pool composition and thus equal p. Due to the continuity of the pool its composition does not change by countable numbers of workers being drawn from or rejoining the pool. In addition, a worker has a probability of zero of being re-hired after once having been laid off. Therefore, we do not consider such a situation with a possibly different prior. The assumption of homogeneous priors for all agents is made for calculational convenience and does not impose a real restriction, as we will discuss in more detail at the end of this section.

In any period, if output is high (low), it is fully informative. The employer will know with probability 1 from such an output that both workers are of high (low) ability. If output is medium she adjusts her beliefs using Bayesian updating. Note that after a realization of medium output the entrepreneur knows that exactly one of the two workers is of high ability. Posterior probabilities will thus add up to one for the two workers. This means that if priors are identical, for example when both workers are drawn from the pool, posteriors will move to 1/2 after medium output.

The evolution of beliefs is as follows: If both workers are kept, beliefs remain constant, as output for a pair of given workers is deterministic. Hence, repeated observation does not generate new information about the workers' ability. If both workers are replaced, two new workers are drawn from the pool and the belief about each of them again equals the prior p. What if one worker is kept and one worker is replaced? Let us denote the updated belief on the worker kept by q_t, where the subscript denotes the number of rounds for which the worker has already been employed. Especially, this means $q_0 = p$. If the newly constituted team produces high or low output, again all uncertainty vanishes. If output is medium, beliefs for the new worker are updated to

$$\hat{p}(q_t) = \frac{(1-q_t)p}{(1-q_t)p + (1-p)q_t}$$

and to

$$q_{t+1} = \frac{(1-p)q_t}{(1-q_t)p + (1-p)q_t}$$

for the worker kept from the previous round. Note that the posterior beliefs maintain the ordering of the prior beliefs, i.e., $q_{t+1} > (<)\{=\}\hat{p}(q_t)$ if and only if $q_t > (<)\{=\}p$. Thus if the employer ex ante assigns a higher belief to the new worker to be of high ability than she assigns to the worker kept, she will also do so ex post, and vice versa if the maintained worker has a higher prior of being the high type. In addition, as posteriors sum up to one, the higher of the two posteriors will always be (weakly) above 1/2, and the lower (weakly) below 1/2.

The next question to raise is when the employer will keep exactly one worker after observing medium output in the first period ($t = 0$). We can establish the following lemma:

Lemma 1 *In the one-team setting, the employer keeps exactly one worker after medium output in the first period only if $p < 1/2$. Moreover, if it is optimal to keep exactly one worker after medium output in the first period it is optimal to do so in all future periods until there is a fully revealing output. The retained worker always is the one from the first period.*

Proof: Keeping a worker after medium output is similar to drawing a worker from a set with a share of high types equal to 1/2. Hence, the employer will prefer drawing one worker from this set instead of from the pool of unemployed if and only if the chance of finding a high ability worker in this set is higher than finding one in the pool, i.e., if and only if $p < 1/2$. Thus $p < 1/2$ is a necessary condition for the employer to reemploy exactly one worker after medium output. It is, however, not sufficient, as the employer might still decide to keep the other worker as well. We will see later that this might indeed happen depending on the parameters.

By the evolution of beliefs derived above, if production is medium again in the second period the employer will assign a higher probability to the worker already employed in the first period of being the high type. Hence, if she decides to again keep only one of the two workers she will always keep the one who already worked for her in the first period.

Moreover, if she decides to keep exactly one worker after medium output in the first period, she will consequently do the same in all future periods until she observes a fully revealing output: If it was optimal in the first period to keep one worker with a chance of

50% of being a high type, it must also be optimal to keep a worker with an even higher chance of being of high ability. Hence, the entrepreneur will not replace both workers in later periods. Conversely, if it was optimal after the first period to lay off one worker who with a chance of 50% was a high type, it cannot be optimal to keep a worker with a lower chance of being type H later. Hence, the decision maker neither will keep both workers in later periods.

The reasoning why the entrepreneur always reemploys the same worker is as follows: The belief on a retained worker to be the high type in period $t+1$ of his employment when output was medium in the previous t periods, can be stated as

$$q_{t+1} \quad = \quad \frac{1}{1+\left(\frac{p}{1-p}\right)^t} \; .$$

Posteriors will thus approach one, if $p < 1/2$, remain constant for $p = 1/2$ and converge to zero for $p > 1/2$. Thus, if the employer decides to keep exactly one worker after medium output, she will keep this worker after Y_M in all future periods, become more and more convinced that this worker is of high ability and may only release him if he proves to be an L-type for sure, i.e., after observing low output. □

This result on repeatedly sticking to keeping one worker after medium level production is rather important for the further discussion of strategies. In general, a strategy has to describe the entrepreneur's employment decision after any history of outputs and beliefs. The last result allows us to fully characterize the behavior of an optimizing entrepreneur by her decision to keep two, one or no workers after the first period.

Lemma 2 *In the one-team setting, all future actions of an optimizing decision maker are determined by the choices in the first period.*

Proof: To proof the lemma we will in turn look at the implications of keeping both, none or one worker on future decisions.

If the employer decides to keep both workers, no further information is generated and thus the decision in future periods remains unchanged. This is independent of the actual level of output in the first period.

If she decides to replace both workers, all information is deleted and she faces the same decision problem in the second period as in the first period. Thus, she will replicate her strategy from this period.

Now, if she decides to keep exactly one worker we have to distinct between the output levels. If output is fully revealing, i.e., either high or low, all future output will also be fully revealing. Then the entrepreneur will sample one new worker from the pool in all future periods until she has got two high types. Hence, if she for example has a known low type and decides to keep him, she keeps him until she observes medium output and then keeps the thus identified high type and samples until she observes high output, i.e., identifies a second high type.

Finally, as derived in Lemma 1, if the employer decides to keep exactly one worker after medium output, she keeps this worker until his true type is revealed (which might take infinitely long). If the worker is of high ability, she will keep him for sure. If he is of low ability, she might as well keep him for identification or replace him with someone from the pool, depending on her policy on known low types that is already determined by her first period policy after low output. □

Note that the lemma can easily be extended to the case of individual priors: Suppose each worker has an individual ability signal yielding a prior possibly different from the average pool prior p. Posteriors after medium output will again add up to one and the decision maker will still only keep a worker if the posterior is higher than the pool prior. The only difference is that after observing Y_M the decision maker will not stick to a worker with posterior $q > p$ if she by chance sampled a worker with prior $r > q$. Although the last part of Lemma 1 does not apply when priors differ, it is straightforward that the last step of the proof of Lemma 2 only requires that exactly one worker is kept until output is fully revealing and does not hinge on the identity of the worker kept.

Following Lemma 2 we can equate strategies of an optimizing employer with first-period choices. A first-period strategy S is a triple (x, y, z) where $x, y, z \in \{0, 1, 2\}$ denotes the number of workers kept after high, medium or low output in the first period. Hence, a tuple $S = (x, y, z)$ describes the strategy to keep x workers if output is Y_H, y workers if output equals Y_M, and z workers after observing low output Y_L. We only consider pure strategies. Thus there are 27 possible strategies. We denote the expected lifetime utility from employing such a strategy S in the first period and behaving optimally in all subsequent periods by $u[S]$.

Dominant strategies

Clearly, when output is high there is no incentive to replace a worker: Both workers are of known high ability and there is neither scope for improvement in terms of output produced nor in terms of information acquired. Therefore, it is a dominant strategy to always keep both workers after high output. Conversely, after low output the employer knows that both workers are of low ability for sure. Thus if she kept both workers, she would receive the minimum output in the next period for sure without gaining any additional information. Replacing one or both of the workers will result in at least the same output with the opportunity for medium or high output. Hence, it always is a profitable deviation to replace at least one of the workers if both are of known low ability. Keeping two known low ability workers thus is strictly dominated.

This reduces the strategy space to six strategies: three in which known fools are replaced, $(2, 0, 0)$,$(2, 1, 0)$, and $(2, 2, 0)$, and three where a known fool is kept for the purpose of identification, namely $(2, 0, 1)$, $(2, 1, 1)$, and $(2, 2, 1)$.

2.2.2 Lifetime utilities and optimal policies

For the six relevant strategies we can calculate the lifetime utilities for the entrepreneur. The utilities consist of the expected first-period output as well as the continuation value from the different policies. If the entrepreneur for example uses strategy $(2, 0, 0)$, she receives high output for ever if she draws two high types, as she will keep these workers.

When observing medium or low output she gets the current benefit from this (Y_M and Y_L, respectively) and starts the search all over again in the next period. This occurs with probability $1-p^2$ in every period, hence the discount factor in front of the equation. The other lifetime utilities are derived in a similar way and we can state them as

$$\begin{aligned}
u\left[(2,0,0)\right] &= \frac{1}{1-\beta(1-p^2)}(p^2u_H + 2p(1-p)Y_M + (1-p)^2Y_L) \\
u\left[(2,1,0)\right] &= \frac{1}{1-\beta\frac{(1-p)^2}{1-\beta p}} * \\
&\quad \left(\frac{p^2}{1-\beta(1-p)}u_H + \frac{(2-\beta)p(1-p)}{(1-\beta p)(1-\beta(1-p))}Y_M + \frac{(1-p)^2}{(1-\beta p)}Y_L\right) \\
u\left[(2,2,0)\right] &= \frac{1}{1-\beta(1-p)^2}(p^2u_H + 2p(1-p)u_M + (1-p)^2Y_L) \\
u\left[(2,0,1)\right] &= \frac{1}{1-\beta 2p(1-p)}(p^2u_H + 2p(1-p)Y_M + (1-p)^2u_L) \\
u\left[(2,1,1)\right] &= \frac{p^2}{1-\beta(1-p)}u_H + \frac{(2-\beta)p(1-p)}{(1-\beta p)(1-\beta(1-p))}Y_M + \frac{(1-p)^2}{(1-\beta p)}u_L \\
u\left[(2,2,1)\right] &= p^2u_H + 2p(1-p)u_M + (1-p)^2u_L
\end{aligned}$$

where

$$\begin{aligned}
u_H &= \frac{Y_H}{1-\beta}, \\
u_M &= \frac{Y_M}{1-\beta}
\end{aligned}$$

are the lifetime utilities from keeping a pair with high and medium output, respectively.

$$u_L = \frac{Y_L}{1-\beta(1-p)} + \frac{\beta p Y_M}{(1-\beta(1-p))^2} + \frac{\beta^2p^2Y_H}{(1-\beta(1-p))^2(1-\beta)}$$

is the expected lifetime utility from keeping one known low type for identification, replacing it with a high type after finding one and then further searching on for a second high type.

The optimal policy for each parameter constellation (p,β) will in general depend on whether production is submodular, supermodular or both. In order to simplify the comparison of optimal policies we normalize Y_H to 1 and Y_L to 0. Thus the knife-edge case of modularity is characterized by $Y_M = (Y_H + Y_L)/2 = 1/2$, and for $Y_M < (>) 1/2$ production is supermodular (submodular). From the lifetime utilities it is straightforward to calculate the optimal policy for any parameter constellation $(p,\beta,Y_M) \in (0,1)^3$.

In the following we will derive the main proposition of this section, namely that it is never optimal to reemploy workers of known low ability. For this let us compare the policies for the modular case to gain some intuition on what will drive employment decisions. We will start by first pairwise comparing policies that differ with respect to reemployment after medium output only and in the next step look at the differences of policies when varying the employment decision after low output. Combining these results we will be able to further characterize the overall optimal policies for all parameter constellations.

Figure 2.1 compares the three policies not keeping a worker of known low ability. The line in each panel shows the parameter constellations for which the compared policies yield identical utility under modularity. The strategies denoted in each region indicate dominance of the respective police over its competitor. A positive (negative) sign denotes that the policy's dominance in a region is reinforced if production is supermodular (submodular). The first panel in Figure 2.1 confirms the argument made earlier: When facing the decision whether to keep one or no worker after medium output the entrepreneur will prefer to replace both workers if and only if the chance of finding a high ability worker in the pool is above 1/2, will be indifferent for $p = 1/2$, and will prefer to keep one worker if $p < 1/2$. As in the discussion above we did not make assumptions on the production function, this will be true even under the premise of sub- and supermodularity.

Comparing the policy of keeping both workers with keeping neither after medium output (the middle frame of Figure 2.1), one can see that keeping both workers is better for low p and the lower β. This is for the following reason: If the chance of finding a high ability worker in the pool is rather low, the entrepreneur will be reluctant to lay off one, and might prefer keeping one low type to losing the high type. The lower the discount factor the less willing the entrepreneur will be to sacrifice safe medium output for the potential chance of high output in the distant future which might only come after a long series of low output if the share of high ability workers in the pool is low. This is reinforced if production is submodular, as indicated by the minus sign: Note that introducing submodularity in the production function can be interpreted as relatively decreasing the additional value of finding a second high type. If the additional benefit from high output versus medium

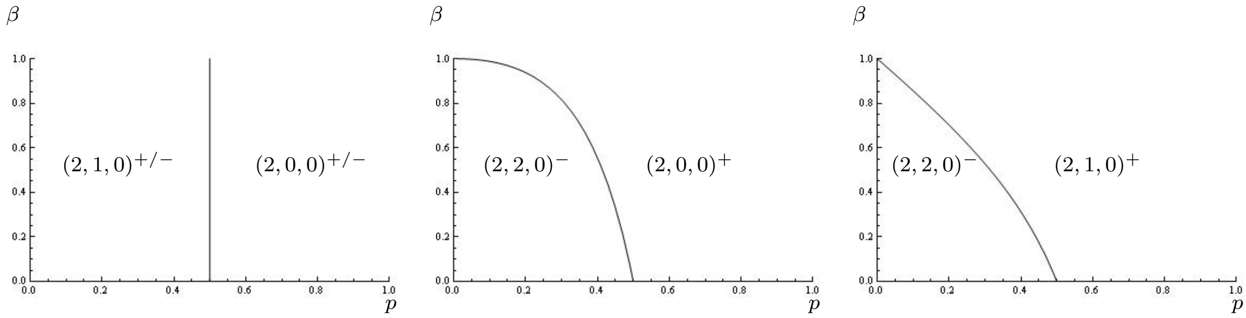

Figure 2.1: Comparing utility from policies replacing a worker of known low ability. Lines indicate parameter constellations (p, β) for which respective strategies generate equal payoffs given modularity. $+$ $(-)$ indicates that a policy will remain dominant in the respective region under super-(sub-)modularity.

output is small, the incentives not to sacrifice safe medium output for unlikely high output strengthen, thus the area for which keeping two workers after medium output is optimal will be even larger for submodular production functions and smaller for supermodular production functions, i.e. when the additional benefit of a second high type is relatively high.

The same line of argument qualitatively explains the drivers behind the decision whether to keep one or two workers after medium output as depicted in the third part of Figure 2.1.

Combining these results, one can see that within these three policies, replacing both workers after medium output is optimal for beliefs above 1/2, keeping both workers is best practice for very low beliefs and sufficiently low discount factors, and keeping exactly one worker is optimal for medium beliefs below 1/2, as depicted on the left hand side of Figure 2.2. The panel on the right hand side provides the outcome from comparing the three policies keeping a known low ability worker. The results and intuitions are qualitatively very similar to the case of not keeping known fools.

Of course the next question to raise is when to keep a known fool. Figure 2.3 compares

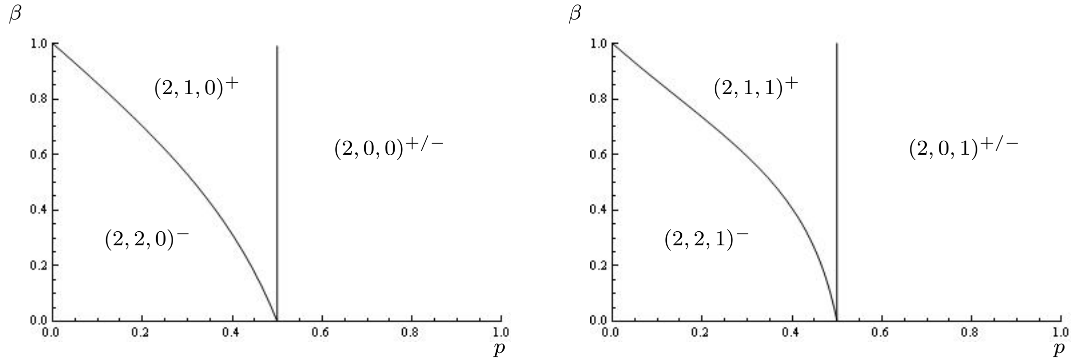

Figure 2.2: Comparing utility from keeping and replacing a worker of known low ability. Policies differ for medium output. Lines indicate parameter constellations (p, β) for which the respective strategies generate equal payoffs given modularity.

policies keeping known low ability workers with ones replacing them but being identical otherwise.

All three comparisons indicate that it is better not to keep a known low ability worker when discount factors are low, i.e., when future profits are worth less and current revenue becomes relatively more important than identification. The cost of keeping a known low ability worker basically consists of foregoing the chance of finding an additional high ability worker in the next period with probability p, while when firing a worker of known low ability the information collected so far is destroyed. Thus the more patient the employer, the more likely she is to sacrifice the short term gamble in exchange for knowledge that almost certainly will pay off in the long run.

In terms of the share p of high ability workers in the pool the picture is inconclusive. If the employer decides to lay off two workers after medium output in the first period, i.e., when comparing profits from $(2, 0, 0)$ and $(2, 0, 1)$ (the leftmost panel of Figure 2.3), it is better to keep a known low ability worker for identification if chances to find a high ability worker are low. Interestingly, this is independent of the size of Y_M: For the locus of parameters for which both policies yield equal utility, $\beta = 1/(2(1 - p))$, the impact of high, medium and low output on the utility from each strategy is exactly equal. Thus

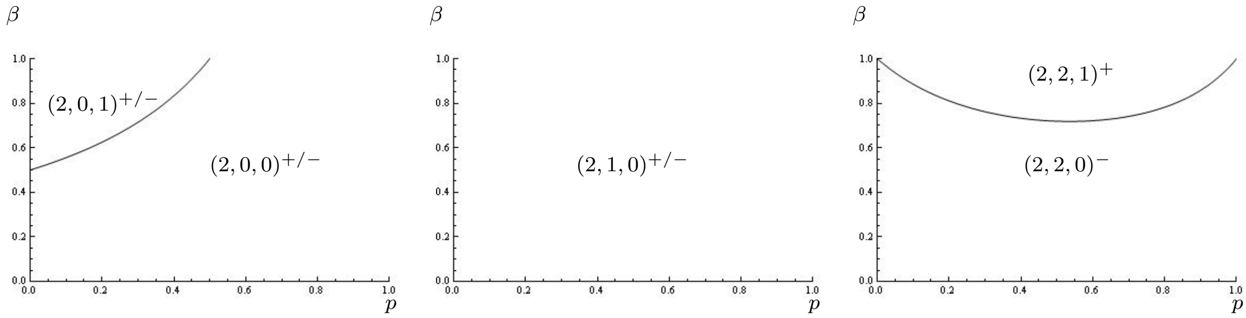

Figure 2.3: Comparing utility from keeping and replacing a worker of known low ability. The policy after high and medium output is fixed in each frame. Lines indicate parameter constellations (p, β) for which the respective strategies generate equal payoffs given modularity.

changing the level of Y_M has the same effect on both lifetime utilities and the result remains unaffected.

From the middle panel we can derive that $(2, 1, 0)$ always dominates $(2, 1, 1)$, independent of the level of modularity, as will formally be shown in the proof of Proposition 1. The reason is that $(2, 1, 0)$ always increases the probability of high output by more than it decreases the probability of medium output relative to $(2, 1, 1)$. Thus expected profit will by higher under the former policy independent of the pool composition p and the discount factor β.

When keeping both workers after medium output (the rightmost panel), the interval of discount factors for which keeping a known fool is optimal notably peaks for values of p close by $1/2$: These two policies only differ in the continuation value after low output, while for other policies the value from redrawing after medium output affects expectations. Therefore the decision on the trade-off between information and current revenue is least distorted. For $p \approx 1/2$ uncertainty about new workers is highest and identification is most valuable, rendering $(2, 2, 1)$ better than $(2, 2, 0)$ for a greater range of discount factors. In addition, keeping a fool here can even be superior to not keeping him when the outlook for finding a high ability worker is good. However, this is only of minor interest as we have already seen that for high values of p it is suboptimal to keep two workers after medium output anyway.

Now, let us establish which policy overall is best for each parameter combination: Graphically, if we take the right panel of Figure 2.2 and compare it to the results in Figure 2.3, we see that, for the three policies keeping a known low ability worker, in each region where one policy dominates the other two it is itself dominated by the corresponding policy not keeping the known low ability worker. For example, $(2,0,1)$ dominates $(2,1,1)$ and $(2,2,0)$ for values of $p > 1/2$ but is itself dominated by $(2,0,0)$ in this region. This yields our first result which notably does not depend on modularity:

Proposition 1 *In the one-team setting, it is never optimal to employ a known fool.*

Proof: See the appendix. □

The reasoning for why is it never optimal to employ a known fool is that whenever keeping a known fool is better than an otherwise similar policy not keeping a known fool the situation is such that the decision maker can do even better by changing another aspect of the policy:

- **(2,0,1)**: When the decision maker does not keep any worker after medium output, keeping a known fool is better than not keeping a known fool ($u\left[(2,0,1)\right] \geq u\left[(2,0,0)\right]$) if chances to find a high type are low ($p \leq 1/2$), as it is then valuable to immediately identify an H-type once found. But for the same reason it is then beneficial to keep (at least) one worker after medium output ($u\left[(2,1,1)\right] \geq u\left[(2,0,1)\right]$), as the chance of keeping the H-type from the team and finding a second one in the pool is higher than finding two high ability workers in the pool for $p \leq 1/2$. This even is completely independent of the modularity of the production function as shown in the appendix.
- **(2,1,1)**: When exactly one worker is kept after medium output, $(2,1,0)$ always does better than $(2,1,1)$: Using $(2,1,0)$ increases the expected lifetime utility from high

output by $C \cdot p$ vis-a-vis $(2, 1, 1)$, where C is a positive constant, as when replacing both low types the chance to find two high types increases. Under $(2, 1, 1)$ the difference in probability for medium output versus $(2, 1, 0)$ is $C \cdot (2p + \beta(1-p)^2 - 1)$. It is straightforward to check that this is smaller than $C \cdot p$. Thus a potential gain from getting medium output more often under $(2, 1, 1)$ is always outweighed by the higher chance to receive high output under $(2, 1, 0)$.

- **(2,2,1)**: $(2, 2, 1)$ does better than $(2, 1, 1)$ if medium output is relatively high (namely above a certain bound depending on p and β, see appendix), i.e., when the benefits of finding a second high type are relatively small compared to the loss of potentially losing the first high type. But then it is even better not to keep a known fool at all, i.e., use $(2, 2, 0)$, as chances to find a first, relatively valuable high type increase when replacing both workers after low output.

When production is not modular, the picture of optimal policies does not completely change: Figure 2.4 depicts the optimal policies for all combinations of p and β in a supermodular case ($Y_M = 0.2$, left panel), the modular set-up (center panel), and a submodular case ($Y_M = 0.8$, right panel).

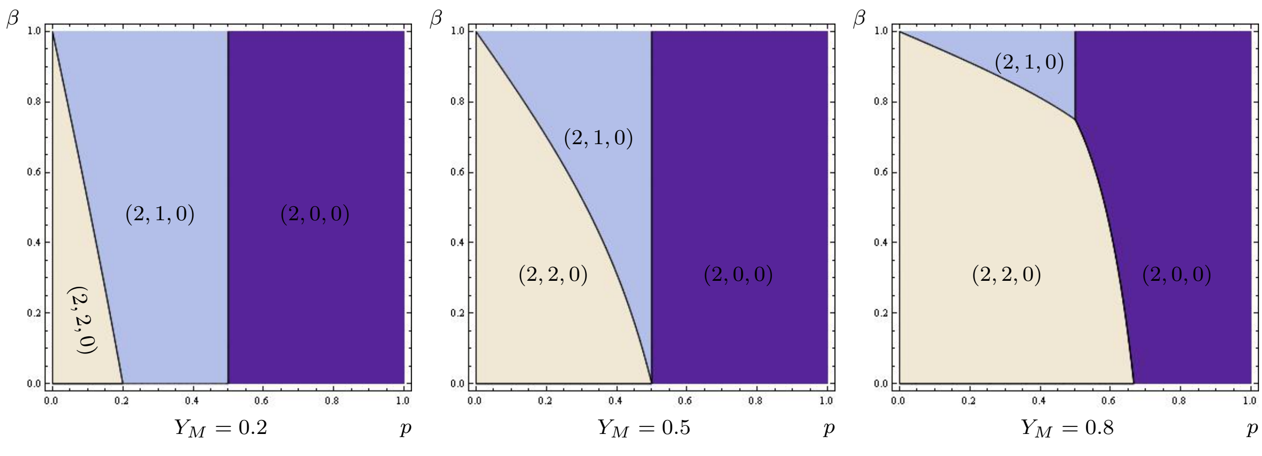

Figure 2.4: Globally optimal policies for varying levels of modularity Y_M.

For submodular production, i.e., when relative gains from high output versus medium are small and it is thus more important to find a first high ability worker than a second one, incentives to sacrifice safe medium output for possible high output (with the risk of a potentially long intermediate period of low output) vanish — the employer will become more and more reluctant to let a possibly high type go and thus is more eager to lock in medium output forever and not learn about ability any more. Therefore the area in which $(2, 2, 0)$ is optimal increases with the level of submodularity. If the future becomes less important (low β) the entrepreneur will even retain all workers after medium output when chances to find a high type are favorable ($p > 1/2$).

In this section we have fully characterized optimal policies for an employer hiring one team and being concerned about current revenue as well as learning about ability. While it is never optimal to employ a known fool, the other aspects of the optimal employment strategy depend on the shape of the production function, the pool composition and the discount factor. In the next part we will extend the entrepreneur's action space by allowing for employment of an additional team.

2.3 Introducing the Winning Team

Until now we only considered an entrepreneur employing one team. Especially in the services industry we often observe firms starting with a small number of employees and then hiring further employees once business is successfully established.[1] Expansion of firms might be possible due to accumulated capital, established business relations, economic growth or any other possible source of increased capital stock. Typically it is the most successful firms that are more likely to expand their business. This on the one hand is due to the capital effects mentioned, but also connected to the reputation effect. A firm with a team of known experts that performed well in the past is more likely to be hired again by a client, even if it is not necessarily the case that the same team of experts is

[1]The Boston Consulting Group for example started its business in Germany in 1975 with "a handful of employees" and very little invested capital, while it now employs more than 500 people (and obviously at lot more capital) in its Munich branch alone. See *www.bcg.de/bcg_deutschland/geschichte/grundung.aspx*, accessed June 5, 2012.

sent to the same client again.

It is therefore a natural extension to assume that an employer after having found a team generating high revenues, i.e., a team of two high ability workers, might want to expand her business by employing a second team. Thus she will again hire two individuals from the pool. But if she now has a team of two known high types (a “winning team”) at hand, she has an additional strategic option before letting teams produce in every period. Apart from deciding whom to keep and whom to replace in the new team, she can as well decide whether to let the winning team and the new team produce on their own or whether to split up the winning team and re-pair them with the members of the new team. Re-pairing offers the advantage of immediately identifying the ability of the new team members at the cost of forgoing certain high output from the winning team in the re-matching period. It will thus decisively depend on the modularity properties of the production function whether the entrepreneur should “change the winning team” or not. In this section we will discuss the entrepreneur’s optimal strategy for such a situation.

2.3.1 Beliefs and strategies

Suppose the entrepreneur now has such a winning team (WT) at hand. Thus she maximizes the sum of the discounted outputs of both teams. Whenever she decides to split this team up, she will always perfectly know the types of all employees after the next round of production. If she decides not to rematch team members the situation is similar in terms of beliefs about the new team as in the one-team scenario. But which consequences does the second team have on the strategy space? In general, the entrepreneur can decide to split the winning team before production in any period. The order of events is depicted in Figure 2.5.[2]

Correspondingly to the one-team case the optimal decisions of a rational decision maker are already pinned down in the early stages. We have the following Lemma:

[2]In $t = 0$ step (1) and (2) collapse to drawing the first pair of new workers.

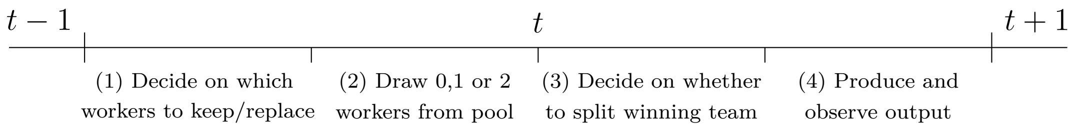

Figure 2.5: Order of decisions and events in any period.

Lemma 3 *In the two-team setting, all future actions of an optimizing decision maker are determined by the choices in the first two periods.*

Proof: To avoid tedious case distinctions we assume that the entrepreneur always keeps the winning team united when she is indifferent about matches, i.e., when she has at least three known high types. Before the first round of production the employer has two employees of known high ability and two of unknown ability. If she decides to split the winning team she will immediately learn the ability of the two unknowns. As above she will always keep known high types. Hence, if both teams produce high output she will keep all workers and future output is deterministic. If, after splitting, one team produces medium output and the other one high output, a thus identified low type will be replaced and the identified high type will be kept along with the winning team, as with three high types it is always better to sample a new, possibly high ability worker from the pool than to keep a known low type. If both teams produce medium output the employer might either want to replace both identified low types and start with the same situation as in the first period again (and thus will again split the winning team because she is in the same situation as before). Or she might keep one of the two low types for identification until another high type is found which in turn is then used in search for a fourth high type. Thus if she decides to split the winning team in the first round, her future strategy is pinned down by her decision in the second round. We denote the policy to change the winning team before the first period of production and replace low types by $CWT-RL$ and the policy to change the winning team and keep one known low type for identification by $CWT-KL$.

If the entrepreneur decides not to split before production in the first round, then after

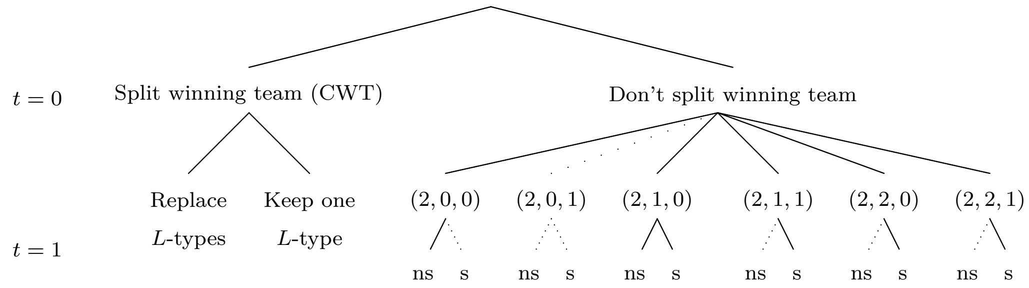

Figure 2.6: Possible strategies in two-team setting. The employer either decides to change the winning team in the first period and then whether or not to keep a known low type (two most left branches). Or she matches the new workers in $t = 0$ and possibly changes the winning team in the second period. Dotted lines mark dominated strategies.

observing output from this round she in the second round has the same set of strategies with respect to reemployment for the new team as in the one-team setting. But after having decided whom to keep and whom to replace (and after potentially hiring new workers) in this second round she can decide whether to now split the winning team before production. In general, this would imply 2^3 additional possibilities for each of the six first period strategies identified in the previous section (namely whether or not to split after each level of output). Yet, splitting the winning team after once having observed outcomes can only be of informational value when there is uncertainty about a worker's ability, i.e., after medium output. From this binary decision of splitting after medium output the number of possibly optimal policies thus doubles to twelve in the second round. Figure 2.6 depicts the decision tree of the employer leading to these twelve policies she might apply in the first two periods in addition to the two policies arising from splitting the winning team in the first period. This already characterizes the full set of relevant policies for a rational decision maker, as we will explain now.

Within this set of policies we can again identify some dominated actions. If the employer decides to dismiss both workers from the new team after medium output, it does not make sense to split the winning team for the second period of production if it did not already pay off to split it for the first period. Thus, splitting the winning team when both new workers are dismissed after medium output is dominated. The decision problem then is reduced to the one of a single team as discussed in the previous section and hence indeed fully described by the first period's optimal decision.

Conversely, if the employer keeps both workers after medium output she will always split the winning team: Output in the next period will be $Y_H + Y_M$ independent of the matching. By splitting she will always gain full information on the new team members' abilities and then keep the identified high type and replace the low type. Note that in contrast to the one-team setting the employer thus will now almost surely end up with four high ability workers even when she keeps two workers after observing medium output in the first period.

If the employer decides to keep one worker after medium output, splitting will once again provide full information on all workers. Once more high types will be kept and low types might be employed for identification which is again pinned down by her employment strategy in the first round. The most tricky case is when the entrepreneur decides to keep one worker and not to change team compositions in the next round. In this case she will also not change team composition after medium output in any future period and again always keep the same worker. The reason for this is that beliefs then evolve as in Section 2.2, i.e., after every period of medium production she will become more certain that the kept worker is of high ability, as again keeping a worker is only rational if $p < 1/2$. In other words, uncertainty of the kept worker's type and hence value of information will be highest after medium production in the first period (remember that the posterior belief is one half in this period). Hence if it did not pay off to split after the first period, it will also not pay off to split in later periods when she has better information on the kept worker's type. Thus, any optimal policy is already fully characterized by the entrepreneur's decisions in the first two rounds. We will again identify strategies with sub-strategies for the first two periods.

Finally, comparing policies not splitting the winning team is similar to the one-team case, thus as seen before all policies keeping a known fool and not splitting are dominated. Policies that can be excluded as dominated from the previous discussion are marked by dotted lines in Figure 2.6. □

2.3.2 Optimal employment policies

In addition to the utilities from policies splitting the winning team before the first round of production, we define the utility from any other policy as $\tilde{u}\left[((x,y,z),c)\right]$ where x, y, and z are as before and $c \in \{s, ns\}$ denotes the decision whether or not to split the winning team after medium output in the first round (s = split, ns = no split). The utilities from the eight policies depicted in Figure 2.6 as not dominated are provided in the following. u_H is the lifetime utility from eternal joint production of a team of two high types, i.e., as defined in the previous section.

$$\begin{aligned}
\tilde{u}_{CWT-RL} &= 2\underbrace{(u_H - (1-p)\frac{Y_H - Y_M}{1-\beta(1-p)})}_{=:\tilde{u}_H} \\
\tilde{u}_{CWT-KL} &= p^2 2u_H + 2p(1-p)(u_H + Y_M + \beta\tilde{u}_H) \\
&\quad +(1-p)^2\Big[2Y_M + \beta(u_H + \underbrace{\frac{1}{1-\beta(1-p)}(p(Y_M+\beta\tilde{u}_H) + (1-p)Y_L))}_{=:\tilde{u}_L}\Big] \\
\tilde{u}\left[((2,1,0),s)\right] &= \frac{1}{1-\beta^2 p(1-p)^2 - \beta(1-p)^2} \\
&\quad * \Big[p^2 2u_H \\
&\qquad +2p(1-p)\left[Y_H + Y_M \right.\\
&\qquad\qquad \left. +0.5\beta(p2u_H + u_H + Y_M + \beta\tilde{u}_H + (1-p)2Y_M)\right] \\
&\qquad +(1-p)^2(Y_H+Y_L)\Big]
\end{aligned}$$

$$\begin{aligned}
\tilde{u}\left[((2,1,1),s)\right] &= p^2 2u_H \\
&\quad +2p(1-p)\Big[Y_H + Y_M \\
&\qquad\qquad +0.5\beta\left[p2u_H + u_H + Y_M + \beta\tilde{u}_H + \right. \\
&\qquad\qquad\qquad \left.(1-p)(2Y_M + \beta(u_H + \tilde{u}_L))\right]\Big] \\
&\quad +(1-p)^2(u_H + Y_L + \beta\tilde{u}_L) \\
\tilde{u}\left[((2,2,0),s)\right] &= u_H + \frac{1}{1-\beta(1-p)^2}\Big[p^2 u_H + 2p(1-p)((1+\beta)Y_M + \beta^2\tilde{u}_H) \\
&\qquad\qquad +(1-p)^2 Y_L\Big] \\
\tilde{u}\left[((2,2,1),s)\right] &= u_H + p^2 u_H + 2p(1-p)((1+\beta)Y_M + \beta^2\tilde{u}_H) + (1-p)^2(Y_L + \beta\tilde{u}_L) \\
\tilde{u}\left[(d,ns)\right] &= u_H + u(d)
\end{aligned}$$

where $d \in \{(2,0,0),(2,1,0)\}$. $\tilde{u}_H$ is the expected payoff when the entrepreneur already has identified one high type and samples until she finds a second high type, i.e., the expected lifetime benefit of high output net of the loss in output before the second high type is found, multiplied by the expected discounted arrival time of a second high type:

$$\tilde{u}_H \quad = \quad u_H - (Y_H - Y_M)\frac{1-p}{1-\beta(1-p)} \; .$$

Correspondingly, $\tilde{u}_L$ is the expected payoff from sampling given a known low type and thus equals u_L from the previous section. Utilities again consist of the immediate pay-off from first-period production and the values from continuing with the respective strategies after observing first-period output.

(Sub-)Modularity

Now suppose production is (sub-)modular. In this case the entrepreneur will always split the winning team for the following reasons: As stated before, in any period the entrepreneur on the one hand wants to maximize current revenue, but on the other hand wants to infer as much information as possible. When facing the decision whether to split the winning team after initially having drawn two workers the consideration of the employer is as follows: If both workers drawn are of high ability, output will be $2Y_H$ no matter how teams are composed, and all uncertainty is resolved. If exactly one of the two workers drawn is of high ability, output will be high plus medium, again independent of team composition. Hence, splitting will generate full information without current costs. If both drawn workers are of low ability, the benefit from splitting the winning team is non-negative as $2Y_M \geq Y_H + Y_L$ due to (sub-)modularity. In addition, changing the winning team again generates full information. Thus

Lemma 4 (Always change a winning team) *If production is (sub-)modular, it is always optimal to split up the winning team.*

Hence, to identify the optimal policy in the (sub-)modular case we only have to consider the two leftmost branches of the tree in Figure 2.6.

Further, if the employer decides to change the winning team before the first period and to keep a known fool after production (the second arm from the left in Figure 2.6), she has to decide whom to team up with whom in the second period. Reuniting the winning team generates expected revenue of $Y_H + pY_M + (1-p)Y_L$, while keeping it split generates $Y_M + pY_H + (1-p)Y_M$. Due to (sub-)modularity splitting again is optimal (note that information generation is not an issue, as when the abilities of three workers are known, information on the fourth will always be complete after production). But this means letting a known high type produce with a known low type. Clearly, the entrepreneur could improve on this by also replacing the known low type: The known low type does not provide any informational advantage and expected output will also increase by replacement. Hence, we have

Lemma 5 (Never keep a known fool) *If production is (sub-)modular and the entrepreneur changes the winning team, it is never optimal to keep a known fool.*

Combining Lemma 4 and 5 we have thus established

Proposition 2 *If production is (sub-)modular, it is a dominant strategy in the two-team setting to always change the winning team and replace known fools.*

Economically, this result is not too surprising, as under submodularity splitting has positive effects on revenue and on information — i.e., there is no trade-off between the competing goals of the entrepreneur. Let us therefore turn to the more interesting opposite case of supermodularity.

Supermodularity

If production is supermodular, the decision maker will be more reluctant to split up the winning team as the relative costs of doing so increase. To further characterize optimal behavior given supermodularity, we have to describe more precisely the degree of supermodularity. We again normalize Y_L to equal 0 and set Y_H equal to 1. Then production is supermodular if and only if $Y_M \in (0, 1/2)$, and the closer Y_M is to 0 the higher the relative benefit from high output, thus the "more supermodular" production is. The optimal policy will thus depend on the extent of supermodularity in addition to the size of the discount factor β and the pool composition parameter p. It is straightforward to calculate the optimal policy for any parameter constellation by just plugging in the numbers. Figure 2.7 visualizes the optimal policy for all (p, β) combinations given increasing values for Y_M from 0.1 to 0.4.

Again it turns out that policies keeping known low ability workers always are dominated. This we state as

Lemma 6 *In the two-team setting, it is never optimal to keep a worker of known low ability.*

Proof: The proof again involves showing that for any set of parameters each policy employing a worker of known low ability is dominated by a policy not employing such a worker. For details see the appendix. $\square$

For low shares of H-types in the pool the best policy is to keep both workers after medium output, because having found a high ability worker, the entrepreneur will not risk to let

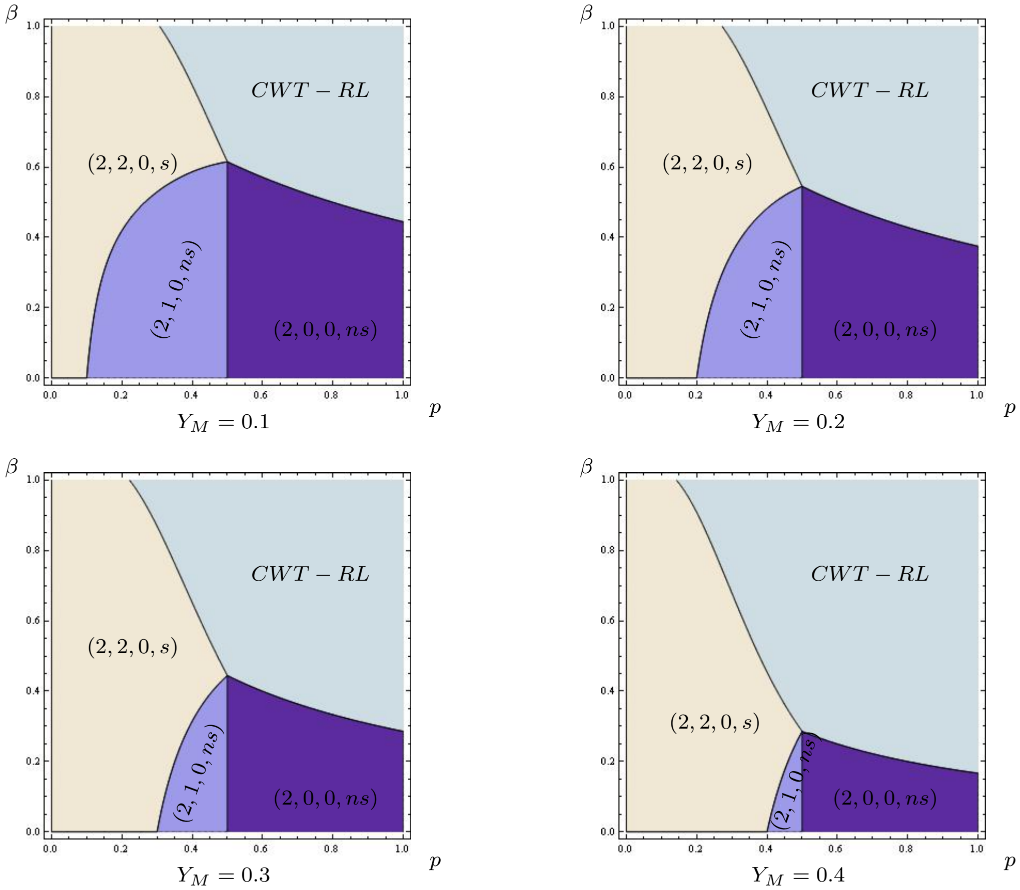

Figure 2.7: Optimal policies for all (p, β)-combinations given values of Y_M equal to 0.1, 0.2, 0.3, and 0.4, respectively.

him go again when p is low. Because of the winning team the H-type can be identified immediately by splitting in the next period (and the simultaneously identified low type will be replaced afterwards). The "less supermodular" production is, i.e., the higher the gain from high versus medium output, the more likely the entrepreneur is to keep both workers after medium output.

For high discount factors and sufficiently high levels of p immediately splitting the winning team is best practice: When the future matters a lot the loss from splitting the winning team in the first few periods until having found a third high ability worker weighs relatively less. This strengthens with increasing value of medium output Y_M as then potential current losses from getting Y_M instead of Y_H decrease.

When the near future is more important than the long term to the employer (low β) and the share of high ability workers is high, it is optimal not to split the winning team: chances are good to find two high ability workers from the pool anyway and waiting for this to happen is not too costly when β is low. These incentives increase with falling Y_M, i.e., when the potential loss from splitting the winning team is high. As argued above if only known high ability workers are kept it does not matter whether the employer splits after observing this, hence, not splitting is (weakly) dominant.

Following the same reasoning as in the one-team setting, for beliefs below $1/2$ it is rational to keep a worker after medium output (i.e., employ $((2, 1, 0,), ns)$) as the chance of keeping the high ability worker is higher than that of finding two high ability workers in the pool. The reason why not splitting is optimal in this case is again due to the low discount factor: By splitting, the decision maker trades off early high output for additional information that might pay off in the distant future only. For low values of β she will thus prefer not to split over splitting when employing $(2, 1, 0)$.

Altogether, for high discount factors the decision maker will split the winning team in the first or second period, as information is more valuable than current revenue. When the present is more important, she is more willing to "gamble" on the future, especially if odds are good to find high ability workers (p is high).

An interesting difference between the one-team setting and the two-team setting is that while the entrepreneur is willing to settle with a mixed team for certain parameter constellations when employing only one team, in the two-team setting she almost certainly ends up with two winning teams: Even if high types are in low supply such that the entrepreneur is reluctant to lay off a potential high ability worker and thus reemploys both workers after medium output, the option to split the winning team in the next period at no cost will lead to certain identification of the high type and subsequent replacement of the low type.

In this section we derived optimal policies when the employer has the additional option of inferring ability by changing team composition. While not changing the winning team can only be optimal under supermodular production, the result on always replacing known fools remains. In the following we will consider the robustness of our results in the one- and two-team setting.

2.4 Robustness

2.4.1 Finite time

While enterprises might in principle last for ever, working life in general is finite. One might therefore ask how optimal decisions change when considering a finite framework. The assumption of an infinite time horizon bears the advantage of time invariance of the employer's continuation decisions. When we do not impose this assumption, the optimal decision will crucially depend on the number of periods remaining, and comparing the pay-offs from the different policies for arbitrary finite time horizons T is more involved than in the infinite setting as we cannot assume the entrepreneur to stick to one policy in every period. However, we can conjecture that the result of not reemploying known fools will be reinforced in finite time: Assuming the lifetime of a generation of workers to equal T and no overlap between generations, the entrepreneur faces the problem of maximizing profit over the finite time horizon. This makes employing a known fool even less attractive as the number of periods over which the entrepreneur can profit from identifying high types earlier necessarily becomes shorter. Consider, e.g., the one-team setting. In the extreme case of $T = 2$ this becomes most evident, as keeping a known low type after the first period will reduce expected second-period profit for sure while the information gained by full identification in the second period is of no value, as all workers leave after this period. Following this consideration the entrepreneur will always switch to replacing known fools in the last two periods, no matter what she did before.

It is straightforward to calculate that under modularity even for $T = 3$ all policies replacing known fools dominate their respective counterparts keeping known fools. Thus this is even stronger than in the infinite-horizon setting, where it was combinations of replacement-policies dominating the keeping-policies only. The alluring induction argument that this could hold true for all finite T, however, does not go through. Comparing, e.g., $(2,0,1)$ and $(2,0,0)$, already for $T = 4$ keeping the known fool (when found in the first period) can be more attractive for values of p below 0.217.[3]

The case of overlapping generations can at least partly be covered by our two-team setting. Remember the entrepreneur is almost sure to end up with two winning teams. If the lifetime of the first winning team is sufficiently long, the entrepreneur will be in the same situation as in the beginning when this team leaves the work force.

2.4.2 Firing costs

Returning to the one-team setting, assume that whenever the employer decides not to keep workers a small but positive amount c for each worker replaced is incurred in the next period.[4] These costs can be interpreted as a settlement fee for a worker replaced or as expenses for the search of a substitute when there are labor market frictions. For the sake of simplicity we restrict the analysis to modular production in this section, i.e., $Y_M = (Y_H + Y_L)/2$. The firing costs make replacement of workers less attractive. When costs are sufficiently high relative to the discount factor, it might thus even become optimal to keep two known low ability workers, i.e $(2,2,2)$ might become an optimal policy. However, $(2,0,2)$ and $(2,1,2)$ will remain dominated as it cannot pay off not to keep a worker after medium output when firing costs are so high that even two known fools are kept. Indeed it turns out that in this setting an employer might want to keep known low ability workers. The optimal policies for costs c of 0.1 and 0.2 are depicted in Figure 2.8.

[3]In the finite case we abstract from discounting, i.e., $\beta = 1$. The argument of this section is reinforced when additionally considering discount factors as they make future profits less valuable, and hence information gathering less attractive.

[4]The timing of costs is postponed for calculational convenience only. It is straightforward to derive equivalent results when firing costs accrue in the same period.

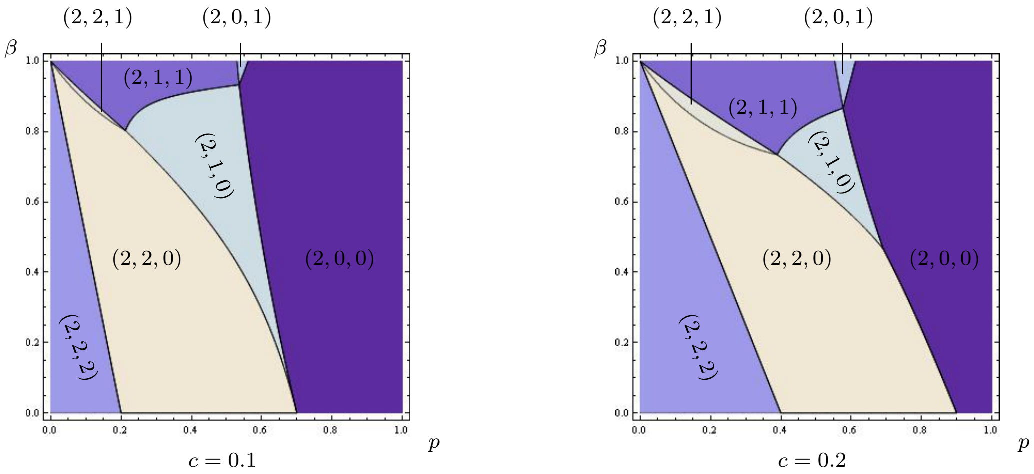

Figure 2.8: Optimal policies for all (p, β)-combinations given firing costs c of 0.1 and 0.2, respectively.

In this scenario policies of keeping one known fool for identification might indeed be optimal especially when discount factors are high, so that identification of high types is very valuable. It should be noted that this is so because keeping a fool in this situation is the lesser evil: The replacement costs accruing for sure might outweigh the potential benefit from finding two high types rather than one in the next period. The effect is the more pronounced the less likely it is to find high types and the higher the fine for quitting employment. Nonetheless, in these situations the entrepreneur does indeed capitalize on the identificational properties of a known low type.

When chances of finding high types are low, firing costs exceed potential gains by so much that the employer never fires a worker even if both are of known low ability, an effect reinforced the less the future matters, i.e., the lower the discount factor.

2.4.3 Additional uncertainty

The assumption of fully revealing output of course is extreme. However, we expect our main intuitions to hold when output is only a noisy signal on ability, for example if there is a scope for effort or if other influences or shocks additionally drive the production result. If the employer can only imperfectly infer a worker's ability after low output, the expected ability of a worker decreases by less than in the full-information case, but this also reduces his informational value for identifying high ability types. Consider for example the situation that the employer keeps one worker after medium output and output in the next period then is low but not fully informative. If the posterior of the worker retained is still above average, it still is rational to further employ this worker but the information on him and his co-worker is less precise. Thus, while it might be that a worker is kept after low output, this happens rather for the potential chance of nevertheless being a high type than for his quality of being an identifier of high types.

2.5 Conclusion

We have discussed the problem of inferring ability when production is executed in teams. In a one-team setting the tools of the employer in any period consist of being able to decide whom to employ in future periods. We find that it is never optimal to keep a known low ability worker independent of the shape of the production technology: Keeping a fool is always dominated by another policy not keeping a fool, no matter whether production is sub- or supermodular. Whether to keep workers on whom there is improved but not full information depends on the parameters: The lower the chance of finding high ability workers and the lower the impact of future output, the more likely the employer is to keep workers of uncertain ability.

When the employer already has a team of known ability (the winning team) at hand and wants to expand the firm by hiring a second team, things partly change. As in the one-team case it is suboptimal to keep a known fool. But the winning team at hand adds the new tool of changing team composition in-house. When production is modular or submodular, it is always optimal for the entrepreneur to exercise this additional option: Information always is complete and no current revenue is lost. When production is supermodular, there is a true trade-off between current revenue and learning. In this case, for low shares of high ability types in the pool it is optimal only to split the winning team after observing medium output from the new team. For high discount levels and a high average ability level in the pool, immediately splitting the winning team is optimal, an effect reinforced the closer the production technology is to the modular case. For low discount factors and relatively high chances of finding a high ability worker, the winning team will always remain unchanged and the problem thus is similar to the one-team setting.

The analysis in this set-up of course is restricted: Uncertainty is fully resolved in two out of three states of the world, production only depends on ability and team size is limited to two members. Moreover, in the real world teams are often installed for other motives such as effort induction, match-specific productivity, learning and spill-over effects. Still, in abstracting from these other aspects we hope to have contributed to identifying the factors driving an employer's optimal decision under uncertainty about ability.

A2 Appendix

Proof of Proposition 1

To show that for any parameter combination $(p, \beta, Y_M) \in (0,1)^3$ each policy keeping a known fool is dominated by a policy not keeping a fool, we compare pairs of policies and then solve for the parameter combinations for which the policies yield the same payoff.

For values of $p < 1/2$ $(2,0,1)$ is dominated by $(2,1,1)$ as argued above, thus not optimal. Comparing $(2,0,1)$ to $(2,0,0)$ yields

$$\begin{aligned} u\left[(2,0,1)\right] &\geq u\left[(2,0,0)\right] \\ \Leftrightarrow \beta &\geq \frac{1}{2(1-p)} \ . \end{aligned}$$

Hence, for $p \geq 1/2$ $(2,0,1)$ is dominated by $(2,0,0)$ for arbitrary feasible values of β and Y_M.

Subtracting $u\left[(2,1,1]\right)$ from $u\left[(2,1,0)\right]$ and solving for β yields

$$\begin{aligned} u\left[(2,1,1)\right] &\geq u\left[(2,1,0)\right] \\ \Leftrightarrow \beta &\geq \underbrace{\frac{1}{1-p}}_{>1} + \underbrace{\frac{p}{(1-p)^2}(Y_M^{-1}-1)}_{>0} \ . \end{aligned}$$

Thus there is no feasible parameter combination for which $(2,1,1)$ does better than $(2,1,0)$.

To prove that $(2,2,1)$ is never optimal we show that whenever $(2,2,1)$ dominates $(2,1,1)$ it is itself dominated by $(2,2,0)$. Let us first determine when $u\left[(2,2,1)\right] \geq u\left[(2,1,1)\right]$ and solve for Y_M. The relevant condition is

$$Y_M \geq \frac{p(1-\beta+2\beta^2(1-p)p)}{1-\beta(2-(2-p)p-\beta(1-p)(1-p(1-2p)))} =: f_1 \ .$$

Note that both numerator and denominator are positive for all values of p and β and that the latter exceeds the former by $a := (1-\beta)(1-\beta(1-p))(1-p)$. Thus the right hand side lies in $(0,1)$. Hence, for any (p, β)-combination there is a feasible value of Y_M for which the utility from $(2,2,1)$ exceeds the one from $(2,1,1)$.

But $u\,[(2,2,1)] \geq u\,[(2,2,0)]$ if and only if

$$\begin{aligned} & \underbrace{(2p-1+\beta(3(1-p)^2-1)+\beta^2(1-p)^2(2p-1))}_{=:x} Y_M \\ \geq\ & \underbrace{2p-1+\beta(3(1-p)^2-1)+\beta^2(1-p)^2(2p-1)}_{=x} + a\ . \end{aligned}$$

It is straightforward to see that if the factor x in front of Y_M is positive, the inequality can only be fulfilled for values of Y_M greater than one, and that it cannot be fulfilled for $x = 0$. If the factor is negative, the condition translates into

$$Y_M \ \leq\ \frac{x+a}{x} =: f_2\ ,$$

i.e., an upper bound on medium output. Comparing this upper bound with the lower bound from above yields

$$f_1 - f_2 \ =\ -\frac{2(1-\beta)^2(1-\beta(1-p))^2(1-p)p}{(1-\beta(2-(2-p)p-\beta(1-p)(1-p(1-2p))))x} > 0\ ,$$

as the numerator and the first factor in the denominator (the denominator of f_1) are positive and x is negative by assumption. Hence, the lower bound an Y_M is higher than the upper bound. In conclusion, there is no value of Y_M for which $(2,2,1)$ generates higher utility than both $(2,1,1)$ and $(2,2,0)$. □

Proof of Lemma 6

To prove the lemma, we again show that the three relevant policies $CWT-KL$, $((2,1,1),s)$, and $((2,2,1),s)$ are dominated for any parameter constellation. The outline of the proof is as above. Starting with $((2,2,1),s)$, we have

$$\begin{aligned} & \tilde{u}\,[((2,2,1),s)] > \tilde{u}\,[((2,2,0),s)] \\ \Leftrightarrow\ & (p((1-\beta(1-p))^2+\beta^2(1-p)^2)-(1-p)(1-\beta(1-p)))Y_M \\ >\ & p((1-\beta(1-p))^2+\beta^2(1-p)^2) > 0\ . \end{aligned}$$

It is straightforward to see that if the factor in front of Y_M is greater than zero, Y_M has to be greater than one, while if the factor is smaller than zero, Y_M has to be negative as well for the inequality to hold.

$CWT{-}KL$ can be shown to be dominated by either $CWT{-}RL$ or ((2,2,0),s): $\tilde{u}_{CWT-KL} > \tilde{u}_{CWT-RL}$ translates into

$$Y_M \quad < \quad \frac{1}{2}\left(1 + \frac{p}{3p - 2(1-\beta(1-p)^2)}\right) \; . \tag{A2.1}$$

For $p \geq 1/2$ the right-hand side becomes negative (the last fraction then is smaller than $\frac{1}{\beta-1} < -1$). But if (A2.1) holds, then

$$\tilde{u}_{CWT-KL} - \tilde{u}\left[((2,2,0),s)\right] < -\frac{(1-p)^2 p(1-2\beta p)}{(1-\beta(1-p)^2)(2-2\beta(1-p)^2+3p)}$$

and the right-hand side is negative for $p \leq 1/2$. Thus for any parameter combination $CWT - KL$ is dominated.

For $((2,1,1),s)$ first assume that $p \leq 1/2$ and compare the utility to the one from $((2,2,1),s)$:

$$\begin{aligned}
& \tilde{u}\left[((2,1,1),s)\right] > \tilde{u}\left[((2,2,1),s)\right] \\
\Leftrightarrow \quad & (1-2\beta(1-p))(1-\beta(1-p)^2-2p)Y_M \\
> \quad & (1-2\beta(1-p))(1-\beta(1-p)^2-2p) + \beta(1-\beta(1-p))(1-p)^2 \; .
\end{aligned}$$

The right-hand side is positive for $p \in (0,1/2)$ (it attains its minimum of $(1-2p)(2p^2+p) > 0$ at $\beta = \frac{1-2r}{1-r}$). Thus as above, a positive factor on Y_M induces $Y_M > 1$ and a negative factor induces $Y_M < 0$. The contradiction for $p = 1/2$ is straightforward.

Finally, for $p > 1/2$, $((2,1,1),s)$ is dominated by $((2,1,0),s)$:

$$\begin{aligned}
& \tilde{u}\left[((2,1,1),s)\right] > \tilde{u}\left[((2,1,0),s)\right] \\
\Leftrightarrow \quad & ((1-\beta(1-p))(1-p)(1+2\beta(1-p)) - p)Y_M \\
< \quad & ((1-\beta(1-p))(1-p)(1+2\beta(1-p)) - p) - \beta(1-p)^2(1-\beta(1-p)) \\
= \quad & 1 - 2p - \beta^2(1-p)^3 \; .
\end{aligned}$$

The last term obviously is negative for $p > 1/2$. Hence, in analogy to the above, if the factor in front of Y_M is positive, Y_M has to be negative, while if the factor is negative, Y_M has to be greater than one, as the right-hand side is "more negative" (i.e., greater in absolute value) than the factor. This completes the proof. □

Chapter 3

Optimal Tolerance for Failure[**]

3.1 Introduction

When searching for new employees firms usually invest a non-negligible amount of time and resources in not only finding the most able employees but also the ones who are most motivated. Indeed, firms regularly claim a certain degree of ambition to be a relevant criterion for employment at the management level. The question "Where do you see yourself in five years?" belongs to the standard repertoire of job interviews and mirrors this concern. One way to think of ambition is that it reflects an employee's responsiveness to monetary incentives.

In this chapter we consider a situation where this responsiveness to incentives is endogenous and depends on the wealth that a manager has accumulated while working for a firm. The wealthier an agent, the lower his marginal utility of income. Hence, the prospect of earning a large bonus in case of success is less appealing to rich managers than to poor ones. Whether a manager has been able to accumulate wealth or not depends on his achievements. In case the manager has been successful in the past, he has earned higher bonuses and is harder to motivate in the future than an unsuccessful manager. At the same time, previous success is likely to carry some information on the ability of a man-

[**]This chapter is based on joint work with Caspar Siegert.

ager with respect to the task at hand. Hence, the principal faces a non-trivial trade-off between keeping only the most able employees and tolerating failure and renewing the employment contracts of unsuccessful but "hungry" managers.

The fact that endogenous changes in wealth influence the responsiveness to incentives is of high importance at hierarchy levels where incentive pay constitutes a large fraction of a manager's total compensation. This is in particular the case for senior executives and directors of large publicly held companies, whose wealth changed by almost US$ 670,000 for each 1% change in their company's stock price in the period between 1992 and 2002 (Brick et al., 2012). However, changes in the wealth of their employees are also a concern for younger companies that compensate their employees with stock options. In fact, in its I.P.O.-prospectus in 2012 the online network Facebook listed as a risk factor related to its business that *"we have a number of current employees [who] [...] after the completion of our initial public offering will be entitled to receive substantial amounts of our capital stock. As a result, it may be difficult for us to continue to retain and motivate these employees, and this wealth could affect their decisions about whether or not they continue to work for us".*[1]

We consider a two-period principal-agent model in which the probability that a project is successful in a given period depends on both the agent's ability and his unobservable effort. In the first period a principal hires an agent of unknown ability and offers him a wage that is contingent on the project's success. Conditional on success, the principal can then decide either to rehire the agent in the second period or to hire a new agent from a pool of ex ante identical employees. If the project is successful in the first period, the principal is going to adjust his belief on the agent's ability upwards. But a success will also trigger a bonus payment, which increases the agent's wealth and makes it more expensive to motivate him in the next period. While a higher wealth may reduce the agent's risk

[1] "Registration Statement under The Securities Act of 1933," by Facebook Inc., Published by The Securities and Exchange Commission, February 1, 2012, www.sec.gov/Archives/edgar/data/1326801/000119312512034517/d287954ds1.htm#toc287954_2, accessed May 31, 2012

aversion and make him more inclined to accept a bonus contract with uncertain future income, it also reduces the agent's marginal utility of income and makes it harder to compensate him for his effort. In this chapter we consider a situation where the second effect dominates and show that this is indeed the case under weak assumptions on the agent's utility function. Conversely, an unfavorable outcome in the first period reduces the principal's belief about the agent's ability. But it also reduces the agent's wealth since he will be financially punished in case of failure. It is thus not clear if the principal should rehire successful managers and if he should replace unsuccessful ones. Indeed, it may be optimal to tolerate failure, that is, to rehire unsuccessful managers since they have a high marginal utility of income and are hence more susceptible to monetary incentives — i.e., they are "hungry" for success.

Continuing employment relations only in case an agent has been successful is optimal whenever a success in period two is either extremely important or hardly matters at all. If success is very important, a principal will decide to offer a contract that induces the agent to exert maximal effort irrespective of the agent's employment history and more able managers are hence more likely to be successful. Moreover, the cost of remuneration is small relative to the profits in case the project turns out successful. Hence, the principal employs the agent who is most able in expectation. After a positive outcome in period one this is the current employee, while after failure this requires hiring a new employee. If the value of success is very low, the principal offers a contract that hardly implements any effort at all. Hence, the cost of incentives and the level of effort that a principal implements are very similar for agents with different track records and again it is optimal to hire the most able manager. Conversely, for intermediate values of success it may be optimal to tolerate failure and to rehire unsuccessful managers. In this case the cost of inducing effort will be an important determinant of firm profits and it may hence be optimal to hire managers that are "hungry", even if this comes at the cost of a lower expected ability. Moreover, the principal is more likely to reemploy an agent after low output if there is low ex ante uncertainty with respect to ability. In this case the principal infers little information from the fact that the agent has failed in period one. Hence, the benefit of employing an agent with low wealth outweighs the cost of having an employee with low expected ability. Similarly, if uncertainty is low a principal will be less optimistic with respect to the ability of a successful manager. Hence, he will be more likely to dismiss such an agent and to hire a new and “hungry” manager instead.

There is a wide strand of literature considering dynamic moral-hazard problems. The seminal papers on career concerns (Gibbons and Murphy, 1992) or the optimality of linear incentives (Holmström and Milgrom, 1987) abstract from wealth effects by assuming that agents have constant absolute risk aversion and that a change in wealth reduces not only the marginal utility of income, but also the marginal disutility of effort. Since these effects off-set each other, the cost of implementing effort is independent of wealth. One of the few papers considering wealth effects in a dynamic agency problem is Spear and Wang (2005). They consider situations in which an agent either becomes too wealthy to be susceptible to monetary incentives or “too poor to be punished” due to limited liability constraints. However, their model abstracts from differences in the ability of agents. In a similar vein, Biais et al. (2010) show that it may be necessary to reduce the scale of a project when the manager comes close to his limited liability constraint.

Thiele and Wambach (1999) discuss general conditions under which the cost of incentives is increasing in the agent’s wealth. This will be the case when the decrease in the marginal utility of income is large relative to the change of the agent’s risk aversion. The opposite can be true if agents have decreasing absolute risk aversion. In this case a richer agent will be less concerned about the income risk associated with performance pay and may be prepared to accept a lower remuneration than a poor agent. Thiele and Wambach do not, however, consider the question of optimal tenure or the interplay between wealth and ability.

An alternative explanation for why it may be optimal to treat unsuccessful managers favorably is presented by Manso (2011) who argues that a principal may need to reward short-term failure in order to encourage experimentation with technologies of uncertain productivity. In a related paper, Tian and Wang (2011) show empirically that startups financed by more failure-tolerant venture capital firms are more innovative. Landier (2006) stresses a different effect of leniency vis-à-vis failure: If banks offer to fund a new project in case an entrepreneur went bankrupt, good entrepreneurs are more likely to abandon bad projects and it is indeed optimal for banks to fund new projects. While there is an alternative equilibrium in which no entrepreneur obtains new funds after filing for bankruptcy, tolerance for failure may be socially beneficial. Finally, Grossman and Hart (1983) emphasize the general idea that an agent's compensation need not necessarily be monotonically increasing in firm profits. If very low earnings are likely to be caused by desirable actions like experimentation, it may instead be optimal to reward bad outcomes.

Our result that the optimal tenure of a manager may not be increasing in his success is consistent with a large body of empirical literature that finds low effects of firm performance on CEO turnover. On average, a manager in the 10th performance percentile is only two to six percentage points more likely to be forced out of his job than a manager from the 90th percentile.[2] Moreover, the responsiveness of CEO turnover to changes in performance does not seem to be systematically higher for firms with good corporate governance. Warner et al. (1988) observe that the probability of forced turnover does not depend linearly on firm success and that only the least successful managers are likely to be dismissed. In a recent paper Jenter and Lewellen (2010) show that the low responsiveness of forced CEO turnover to firm performance is likely to be driven by measurement error in the classification of resignations as "forced".[3] While they find significantly larger effects of firm performance on turnover than the previous literature, attrition is again concentrated in the lowest performance percentiles. This suggests that, for some reason, shareholders seem to be lenient vis-à-vis mildly unsuccessful CEOs.

The remainder of the chapter is structured as follows: In Section 3.2 we propose our main

[2]See, e.g., Warner et al. (1988); Denis et al. (1997); Hadlock and Lumer (1997); Murphy (1999); Huson et al. (2001) and Kaplan and Minton (2006)

[3]Note that in our theoretical model the distinction between forced and voluntary turnover does not bear any meaning: Since agents are always kept to their reservation utility, they are indifferent between accepting a new contract or retiring and any termination of an employment relation can be interpreted as both, forced and voluntary.

model and characterize the optimal contract in a static, one-period setting. In Section 3.3 we consider the dynamic problem. While we are able to derive some general insights with respect to the optimality of different employment policies, in Section 3.4 we turn to an example where the effort choice is binary. This allows us to characterize the optimal policies more closely. The main insights from the general case carry over to the example and we can additionally derive necessary and sufficient conditions for optimality of the different employment regimes. Section 3.5 concludes. Technical proofs are relegated to the appendix.

3.2 The Model

We consider a problem faced by a risk-neutral principal who owns a project and has to employ a risk-averse agent in order to manage the project. The life-cycle of the project can be divided into two periods. In each period, the project can yield profits that are either high or low: $\pi_t \in \{\underline{\pi}_t, \overline{\pi}_t\}$ where $\Delta_t = \overline{\pi}_t - \underline{\pi}_t > 0$ denotes the "value of success" for $t \in \{1, 2\}$. The probability of high profits in a given period is determined by the agent's effort as well as his ability.

There is a continuum of agents with unknown ability. Each agent has an additively separable lifetime utility of $U = u(W_T) - \sum_t C(e_t)$, where $u(.)$ is a standard increasing and concave utility function and $C(.)$ is an increasing and convex effort cost function. Effort is unobservable and can be chosen from an interval $e_t \in [0, \overline{e})$. We impose the usual Inada conditions: $C(0) = 0$, $C'(0) = 0$ and $\lim_{e \to \overline{e}} C'(e) = \infty$. Also, we assume that the agent has access to perfect markets for risk-free borrowing and lending, so his consumption utility $u(W_T)$ only depends on W_T, which is the sum of his initial wealth W and the wage payments that he earns in each period. Let us denote by $k \in \{H, 0, L\}$ an agent who has been successful (H) or not (L) in the first period. Similarly, 0 denotes an agent who has not been hired in period one. The agent's first-period compensation is given by w_k where without loss of generality we can assume that $w_0 = 0$. Similarly, w_{kl} denotes the payment in period two that can be contingent on the period-one outcome k

and the period-two outcome $l \in \{H, 0, L\}$. The principal can reemploy an agent that has been working for him before, but he may also decide to hire a new agent for the second period. We assume that the agent's employment history does not affect his outside option, i.e., $w_{k0} = 0$ for all k.

The probability that the project is successful in period t is given by $P(\pi_t = \overline{\pi}_t) = e_t + \tilde{\theta}$ and depends not only on the agent's effort, but also on his ability level $\tilde{\theta}$. We assume that $\tilde{\theta}$ describes how well the agent performs at the task at hand and that $\tilde{\theta}$ is fixed but unknown to the principal as well as to the agent. We will denote the distribution of talent across agents by $F(\tilde{\theta})$ and the expected quality of an agent by $\theta = E_0(\tilde{\theta})$. The support of $\tilde{\theta}$ consists of a strict subset of $[0, 1]$ where the upper bound satisfies $\overline{\theta} + \overline{e} \leq 1$ and the lower bound is denoted by $\underline{\theta}$.

The key trade-off a principal faces is as follows: On the one hand, an agent that has been performing well in the first period is likely to be of high ability, which makes it more attractive to reemploy him. On the other hand, he becomes more wealthy since he has earned high wages in period one. Under mild assumptions on the agent's risk preferences, this makes him more costly to motivate and may be the reason why hiring a new agent may turn out to be optimal. Let us define $h(v) = u^{-1}(v)$ as the wealth an agent needs in order to attain a level of consumption utility v.

Assumption 1 *We assume that*

$$h'''(v) \geq 0 \qquad \forall v \ .$$

This assumption will be imposed throughout the chapter. It is a sufficient condition that ensures that the effect higher wealth has on the principal's profits is negative: A larger level of wealth decreases the marginal utility of income and makes it harder to compensate the agent for his cost of effort. At the same time a richer agent may become less risk-averse and may be more inclined to accept a contract that offers him an uncertain future income. Assumption 1 ensures that the first effect dominates and that the principal would always prefer to hire a less wealthy agent. In our setting, it is equivalent to Assumption 1 (vii) of

Thiele and Wambach (1999). The assumption is satisfied for most of the commonly used utility functions, in particular utility functions exhibiting constant absolute risk aversion or constant relative risk aversion with a risk aversion parameter greater than or equal to one half.[4] More generally, it is satisfied whenever the agent's coefficient of absolute risk aversion is not decreasing too strongly in his wealth. For a more detailed discussion under which conditions the principal prefers a poorer agent to a richer one see also Chade and Vera de Serio (2011).

3.2.1 The single-period problem

In order to gain some insight into the principal's problem, let us start by considering the static problem a principal faces in case he never reemploys the agent in the next period. Faced with a given wage schedule, the agent chooses effort so as to maximize his expected utility $U(e_1|w_H, w_L) = (\theta + e_1)u(W + w_H) + (1 - \theta - e_1)u(W + w_L) - C(e_1)$. Hence, the level of effort the agent exerts is implicitly defined by

$$u(W + w_H) - u(W + w_L) - C'(e_1) = 0$$

which is independent of θ. The optimal contract that implements a given level of effort e_1 minimizes the expected wage payments subject to the participation constraint (PC) and incentive compatibility constraint (IC):

$$\max_{v_H, v_L} \Pi(W, \theta, e_1) = (\theta + e_1)(\overline{\pi}_1 - h(v_H) + W) + (1 - \theta - e_1)(\underline{\pi}_1 - h(v_L) + W)$$

$$\text{s.t.} \qquad (\theta + e_1)v_H + (1 - \theta - e_1)v_L - C(e_1) = u(W) \tag{PC}$$

$$v_H - v_L = C'(e_1) \tag{IC}$$

where $v_k = u(W + w_k)$. Since the agent has an initial wealth of W, the principal only has to pay him a wage of $h(v_k) - W$ in order to make sure that the agent has a consumption utility of v_k in state k. It is easy to see that $w_H \geq w_L$ where the inequality must be strict

[4]Applying the inverse function theorem yields $h''' = 1/(u')^3(2(A(W))^2 + A'(W))$ where $A(W) = -u''(W)/u'(W)$ is the measure of absolute risk aversion. In case of constant relative risk aversion with risk aversion parameter r the condition simplifies to $r(2r - 1)/(W^2(u')^3) \geq 0$.

whenever the principal implements some $e_1 > 0$. Since the principal only chooses two wages, the wage scheme is fully pinned down by the two constraints and we get

$$\begin{aligned} v_H(e_1) &= u(W) + C(e_1) + (1-\theta-e_1)C'(e_1) \ , \\ v_L(e_1) &= u(W) + C(e_1) - (\theta+e_1)C'(e_1) \ . \end{aligned}$$

Convexity of $C(e_t)$ implies that $v_L \leq u(W)$ and the agent earns negative wages in case he is unsuccessful. In order to make sure that the agent still has an incentive to accept the contract, he is payed a positive wage that compensates him for these losses as well as for the cost of effort in case of a positive outcome. Given the optimal wage scheme for each level of effort, we can characterize the optimal effort level e^* as follows:

$$(\overline{\pi}_1 - h(v_H)) - (\underline{\pi}_1 - h(v_L)) - E(h'(v_k)v_k'(e^*)) = 0 \tag{3.1}$$

and it is easy to show that $e^* > 0$. An increase in effort makes it more likely that the project is successful. The resulting benefit depends on the difference in profits net of wage payments between the two states. At the same time, increasing the effort level requires the principal to increase the wage payments that the agent can expect to earn for a given probability of success. This is captured by the term $E(h'(v_k)v_k'(e^*)) \geq 0$.[5]

The wealth effect

The fact that a richer agent has a lower marginal utility of income makes it harder to motivate him. While the cost of exerting effort is independent of an agent's wealth, the prospect of earning a high wage in case the project turns out successful is more attractive for poor agents. Consequently, a principal would always rather employ a poor agent than a rich one. Moreover, given that the appeal of poor agents is driven by the cost of incentives, it should not come as a surprise that the principal also decides to implement less effort the more wealthy an agent is.

[5]To see that this term is positive note that $h' > 0$ and $h'' > 0$ by concavity of $u(.)$.

Proposition 1 *The principal's profit is decreasing in the agent's wealth:* $d\Pi^*/dW < 0$*. Additionally, the optimal effort level that a principal implements is decreasing in the agent's wealth:* $de^*/dW < 0$*.*

Proof: Consider the impact a change in the agent's initial wealth has on the principal's optimal profits. Applying the envelope theorem, for any effort level e^* this effect is given by

$$-\left[(\theta + e^*)h'(v_H) + (1 - \theta - e^*)h'(v_L)\right] u'(W) + 1 \ .$$

From Assumption 1 we know that $h'(v)$ is convex, so $h'\left((\theta + e^*)v_H + (1 - \theta - e^*)v_L\right) u'(W) > 1$ is sufficient for the expression to be negative. Using the inverse function rule we can check that this is indeed the case. Hence, the principal's profit is decreasing in the agent's wealth.

Moreover, taking the derivative of the marginal return to effort (3.1) with respect to W yields

$$u'(W)\left[-(h'(v_H) - h'(v_L)) - E\left(h''(v_k)v_k'(e)\right)\right]$$

which is strictly negative for all $e > 0$. Hence, the returns to effort are smaller the larger an agent's wealth and the optimal level of effort implemented by a principal must be decreasing in the agent's wealth. □

Two comments are in order: If the agent's coefficient of absolute risk aversion is strongly decreasing in his wealth, Assumption 1 is violated and Proposition 1 may no longer hold. While a richer agent still has a strictly smaller marginal utility of income, he is also considerably less risk averse. This implies that a rich agent is less concerned about the possibility of earning negative wages in case he is unsuccessful and he is more likely to accept a given contract. A positive wealth effect would trivially lead to optimality of a policy of only reemploying successful agents, as the trade-off between wealth and ability vanishes.

The fact that our results hold for constant absolute risk aversion may be surprising at first. The literature on the dynamic provision of incentives (see, e.g., Holmström and

Milgrom, 1987, and Gibbons and Murphy, 1992) typically exploits the fact that optimal incentives are independent of wealth for CARA utility. Yet, in those settings consumption utility and the cost of effort are not additively separable. Instead, an increase in wealth reduces the marginal utility of income but also results in effort being less painful.

The ability effect

While at the beginning of period one the principal does not have any information on the quality of a specific agent, at the end of period one he can draw conclusions on the ability of an agent from the profits the project has generated. In order to derive optimal employment decisions it is hence necessary to consider how the principal's one-period profits depend on an agent's presumed ability.

Proposition 2 *For a given level of wealth, the principal always prefers a more able agent to a less able one:* $d\Pi^*/d\theta > 0$.

Proof: The first order condition for effort (3.1) tells us that even if we account for wage payments, the principal still makes strictly larger profits in case of a positive outcome. For any given contract, all agents that accept the contract exert the same level of effort and the probability of a positive outcome in period two is hence increasing in the belief θ, making more able agents more attractive. Finally, agents with a high ability anticipate that they are more likely to be successful than less able agents. Since $v_{kH} \geq v_{kL}$ and all agents exert the same level of effort, this implies that a more able agent will accept any contract that would be accepted by a less able one. □

Once the principal has observed period-one output, his posterior belief about the ability of an agent who has earned profits π_1 is given by

$$\begin{aligned}
\theta_H = E_1\left[\tilde{\theta}|\pi_1 = \overline{\pi}_1\right] &= \theta + \frac{\sigma^2}{\theta + e_1} , \\
\theta_L = E_1\left[\tilde{\theta}|\pi_1 = \underline{\pi}_1\right] &= \theta - \frac{\sigma^2}{1 - \theta - e_1} ,
\end{aligned}$$

where σ^2 is the variance of the prior distribution $F(\tilde{\theta})$.[6] Whenever the agent has been earning high profits, this is good news about his ability: Since all agents exert the same level of effort in equilibrium, a highly able agent is more likely to be successful than a less able agent. By the same logic, an unfavorable outcome in period one reduces an agent's expected ability. The amount of updating depends on the variance of the prior distribution $F(\tilde{\theta})$: The more uncertain the agent's ability was ex ante, the more a principal optimally infers from realized profits. If the principal did, however, have a precise idea about the agent's ability beforehand, there is little new information he obtains by observing period-one outcomes.

As an aside, it should be noted that $E_0[(\theta_k - \theta)^2]$ reaches a local maximum at $e_1 = 0$ and $e_1 = \overline{e}$. In these cases the principal adjusts his prior most strongly in expectation and the amount of learning is maximized. If the principal implements no effort at all, a high outcome is very informative about the agent's ability since it cannot be due to the agent's hard work. Conversely, in case $e_1 = \overline{e}$ a low outcome is very informative: Due to the high level of equilibrium effort, any agent is likely to earn high profits. If an employee is nevertheless unsuccessful, this implies that he is probably not very suitable for the task at hand.

3.3 The Dynamic Problem

Let us now turn to the dynamic setting. In particular, we will focus on the question if it is optimal to reemploy an agent for a second period and how the decision to do so depends on period-one outcomes. If the principal anticipates that he will reemploy an agent, this also affects the optimal period-one contract. However, in Section 3.3.1 we will see that, in order to derive our key results, we can ignore the change in period-one contracts since

[6]Following Bayes' rule the conditional density of $\tilde{\theta}$ after high profits is

$$f(\tilde{\theta}|\pi = \overline{\pi}) = f(\tilde{\theta})\frac{P(\overline{\pi}|\tilde{\theta})}{P(\overline{\pi})} = f(\tilde{\theta})\frac{\tilde{\theta} + e_1}{\int_{\underline{\theta}}^{\overline{\theta}}(\tilde{\theta} + e_1)df} = f(\tilde{\theta})\frac{\tilde{\theta} + e_1}{E[\tilde{\theta}] + e_1} \ .$$

Taking conditional expectations and applying Steiner's theorem yields θ_H.

certain reemployment policies will turn out to be dominated for any period-one contract that implements positive levels of effort. Throughout the chapter we will assume that Δ_1 is sufficiently large for the principal to indeed implement strictly positive effort in period one. This abstracts from the uninteresting case in which agents do not earn any wages in the first period and there are no wealth effects.

In Section 3.3.2 we consider the adjustment in period-one contracts. Our analysis is simplified by the fact that the period-one contract is fully pinned down by the effort level that a principal decides to implement in period one. Hence, while anticipating certain reemployment decisions may affect the effort that a principal decides to implement in period one, the structure of the contract remains unchanged. Nevertheless, we will see that it is unclear in which direction a principal decides to adjust effort as a function of different reemployment policies.

We assume that the principal always has full bargaining power and that he offers a series of short-term contracts. The assumption that the principal has full bargaining power even if he wants to rehire a specific agent in period two simplifies notation without affecting our results.[7] Similarly, in Appendix A3.2 we show that the the results by Fudenberg et al. (1990) can be applied to our setting and that ruling out long-term contracts is without loss of generality. However, restricting attention to short-term contracts allows us to abstract from issues of deferred compensation. In period one an agent does not care about period-two wages since he will be kept to his reservation utility no matter whether he is offered a new contract or not. Consequently, the agent's expectations with regard to period two payments do not influence his actions in period one. Instead, the agent maximizes $E\left(u(W+w_k)\right)-C(e_1)$.

[7] If the principal did not have full bargaining power in period two, he could always use the period-one contract in order to extract the rents that an agent will get in the future. Since period-two contracts will always maximize total surplus, the distribution of bargaining power does not affect actions in period two and the principal's overall profits remain unchanged.

3.3.1 Optimal reemployment

We can describe the possible reemployment choices that a principal faces after the end of period one by four policies: He can decide to never continue employment (NC) or always renew the contract irrespective of period-one outcomes (AC). Alternatively, he may decide to continue with an agent if and only if he was successful, i.e., continue after high output only (HC) or if and only if he was unsuccessful in the first period, i.e., after low output only (LC). The decision whether to reemploy an agent in period two or not after a certain period-one outcome is influenced by two factors. Ceteris paribus, the principal would prefer to employ a less wealthy agent. However, an agent who has earned a negative bonus in the first period is also more likely to be of low ability. It is therefore not obvious if the principal would want to reemploy him or not. Similarly, a successful agent is of high expected ability, but he has also earned positive wages in period one and is therefore harder to motivate. Nevertheless, we are able to show that there are always parameter constellations such that it is either optimal to keep an agent only after high profits (HC) or to keep him in case he was unsuccessful (LC).

Proposition 3 *Assume that $F(\tilde{\theta})$ is non-degenerate. In this case, as $\Delta_2 \to 0$ or $\Delta_2 \to \infty$ it becomes optimal to reemploy an agent if and only if he was successful in the first period. For any strictly positive Δ_2, as $\sigma^2 \to 0$ it becomes optimal to reemploy an agent if and only if he was unsuccessful in the first period.*

Proof: See the appendix. □

Whenever the value of success Δ_2 is small, the principal offers negligible incentives. This implies that differences in period-two effort and differences in the cost of remuneration between agents of different wealth levels are very small. At the same time, there are non-negligible differences in the expected ability of agents with different employment histories. It is therefore optimal to make reemployment decisions solely on the basis of talent and to rehire an agent only in case he has been successful in the past. If an agent has been unsuccessful, it is optimal to hire a new manager who has strictly higher expected ability. Similar reasoning applies if Δ_2 is very large. In this case the principal implements effort levels that are arbitrarily close to $\overline{e}$ irrespective of the agent's wealth. This implies that a more able (but richer) agent is successful with a strictly larger probability than a less able (and poorer) agent. If $\Delta_2 \to \infty$ it follows directly that it is optimal to employ the most able agent that is available since doing so maximizes the probability of a success.

Even though policy HC is always optimal for very large and very small levels of Δ_2, the same does not necessarily hold true for intermediate values of success. In this case, different wealth levels can result in significant differences in the cost of compensation. At the same time, the advantage of employing an agent with higher ability may not be large enough in order to off-set the negative effect of higher wealth. In Section 3.4 we analyze this potential non-monotonicity in the appeal of HC more closely by looking at a situation where effort is binary. However, Proposition 3 already tells us that a policy of reemploying only successful managers will not always be optimal. If there is little ex ante uncertainty with respect to the talent of potential managers a principal obtains very little additional information on the agent's ability by observing period-one output. Hence, while different employment histories are still associated with considerable differences in wealth, agents hardly differ with respect to their expected ability. This makes it optimal to hire the poorest manager a principal can get hold of and to only reemploy an agent in case he was unsuccessful. An illustration of the regions for which each policy is preferred is provided in Figure 3.1.

Since the second part of Proposition 3 holds true for any strictly positive Δ_2, we can fix some $\hat{\sigma}^2$ such that policy LC is optimal. Yet, the first part of the proposition implies that at $\hat{\sigma}^2$ there are still arbitrarily large and arbitrarily small values for Δ_2 for which HC is optimal. In between those extreme cases, however, it is optimal to "reward" failure by only reemploying agents in case they were unsuccessful in period one.

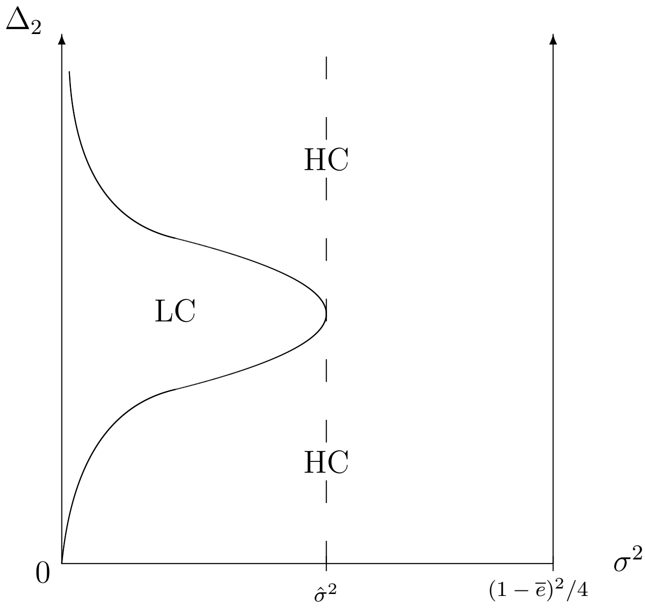

Figure 3.1: Comparison between HC and LC.

Corollary 1 *There exists some $\hat{\sigma}^2 > 0$ such that for all $\sigma^2 \in (0, \hat{\sigma}^2)$ it is optimal to reemploy a successful agent in case Δ_2 is either very large or very small, while for intermediate values of Δ_2 it is optimal to keep unsuccessful managers.*

Accordingly, there is always an interval of values of σ^2 for which the appeal of HC is not monotonic in the value of success: While it is optimal to hire only the most able employee available for low and high values of success, hiring the least wealthy agent is optimal for intermediate values of success.[8]

So far we have only identified conditions such that the extreme polices under which a principal always hires the most able (HC) or the least wealthy (LC) agent are optimal. But it may also pay for a principal to use a more nuanced approach: A principal may try to avoid very rich agents but may still decide in favor of a more wealthy agent in case he is faced with the choice between an unsuccessful agent and hiring a new manager. In this case he never continues employment (NC). Similarly, he may reemploy his agent in a second

[8]Note that depending on the parameter constellation $\hat{\sigma}^2$ might coincide with the upper bound on the variance $(1-\overline{e})^2/4$ that stems from the finite support of $\tilde{\theta}$.

period irrespective of past success (AC). Indeed, situations in which a CEO's tenure is largely independent of his performance seem to be empirically much more relevant than settings in which previous success reduces the probability that a manager's contract is extended.

Unfortunately, the problem is complicated by a plethora of countervailing forces: The value of reemploying a particular agent depends on the effort level a principal has implemented in period one. A period-one contract that implements a high level of effort reduces the expected ability of both, successful and unsuccessful agents.[9] At the same time, a higher level of period-one effort results in large differences in wealth between successful and unsuccessful agents, which makes it less attractive to keep a successful agent. Finally, a principal who anticipates that he will want to reemploy an agent in case he is successful will take this into account when choosing the level of effort he implements in period one. Increasing incentives does not only carry the direct cost of having to pay a larger bonus, but it also implies that the principal may end up with a richer agent in period two. Overall, the appeal of policies AC and NC strongly depends on the agent's higher order risk preferences as well as on the precise shape of the effort cost function. So we will defer the discussion of whether there are parameter constellations for which AC or NC are optimal to Section 3.4.

However, we can obtain some general insights on how the profit earned under the different policies depends on the variance of the prior distribution $F(\tilde{\theta})$. A policy of retaining only successful agents becomes more attractive the more uncertain an agent's ability: higher uncertainty makes good news even better and bad news worse, and a policy of only retaining successful agents capitalizes on positive ability updates. For the same reason a policy of solely reemploying unsuccessful agents becomes less attractive the more uncertainty there is. In contrast to this, profits earned under policy NC do not depend on the variance in the ability of different agents at all: The principal does not learn anything about his period-two agent from period-one output. Hence, the expected quality of agents

[9] Yet, since agents are more likely to be successful in equilibrium, the expected ability of an agent does of course remain constant.

is constant over time and independent of σ^2. The results for AC, however, are ambiguous: Ex ante, the expected quality of an agent still stays constant over time. Yet, the cost of compensating an agent in period two depends on the probability the agent attributes to earning high or low wages in equilibrium. These probabilities in turn are affected by what the agent himself learns about his ability from period one output which explains why profits under AC can be increasing or decreasing in σ^2.

3.3.2 Optimal first-period effort

If a principal anticipates that he will reemploy an agent after the first period with positive probability, this is clearly going to affect the contract that the principal offers in period one: The amount of effort that a principal implements in period one affects the agent's wealth and his perceived ability in the next period. However, we can see that the structure of the first-period contract remains unchanged. Wage payments are pinned down by incentive compatibility and the participation constraint. As in the static setting, the participation constraint will always be binding: Otherwise it would be possible to reduce payments in period one and thus the agent's future wealth in both states of the world, which is strictly beneficial. Hence, the introduction of a second period only affects the effort that a principal optimally implements in the first period.

Proposition 4 *All policies under which a principal rehires an agent with positive probability may lead to optimal first-period effort levels below or above the optimal one-period effort level. When policy NC is employed, effort in the first period equals the optimal one-period effort level.*

Proof: Let us denote the principal's expected period t profit by Π_t and his overall surplus by $\Pi = \Pi_1 + \Pi_2$. In order to gain an insight into the different determinants of period-one effort we will start off by looking at a case where the principal reemploys the agent irrespective of his first-period success (AC). In this case, the return to setting slightly higher incentives in the first period is given by

$$\begin{aligned}\frac{d\Pi}{de_1} \quad = \quad & \overbrace{\frac{d\Pi_1}{de_1}}^{\text{period-one effect}} \quad \overbrace{+\Pi_2^*\left(W_H, \theta_H\right) - \Pi_2^*\left(W_L, \theta_L\right)}^{\text{direct period-two effect}} \\ & \underbrace{+(\theta+e_1)\left(\underbrace{\frac{d\Pi_2^*}{d\theta_H}}_{>0}\underbrace{\frac{d\theta_H}{de_1}}_{<0} + \underbrace{\frac{d\Pi_2^*}{dW_H}}_{<0}\underbrace{\frac{dW_H}{de_1}}_{>0} + \right)}_{\text{indirect period-two effect after high output}} \underbrace{+(1-\theta-e_1)\left(\underbrace{\frac{d\Pi_2^*}{d\theta_L}}_{>0}\underbrace{\frac{d\theta_L}{de_1}}_{<0} + \underbrace{\frac{d\Pi_2^*}{dW_L}}_{<0}\underbrace{\frac{dW_L}{de_1}}_{<0}\right)}_{\text{indirect period-two effect after low output}} .\end{aligned}$$

As in the static setting, an increase in e_1 has a direct effect on period-one profits by making a success of the project more likely while increasing the agent's expected compensation. Additionally, an increase in e_1 affects the profits a principal can expect to make in the next period via three distinct channels: First of all, it becomes more likely that the principal is faced with a successful agent in period two since a high period-one outcome is more likely. Since the principal does generally not make the same amount of profits with each type of agent, this is going to affect his expected profits. We will refer to this as the "direct" period-two effect. Secondly, there are indirect effects on period-two profits: A change in period-one incentives affects the profits a principal can expect to earn with either type of agent. If incentives are large, a positive outcome becomes less informative about the agent's ability and a negative outcome becomes more informative. Successes will partly be attributed to higher effort, while failure despite increased effort is an even worse signal on ability. Hence, the expected ability of both types of agents decreases in e_1.[10] Moreover, an increase in period-one incentives increases the agent's wealth in case of success and it reduces his wealth after low outcomes. While the first effect reduces period-two profits, the second effect has a positive impact on the principal's expected surplus. Whether optimal effort increases or decreases in comparison to one-period optimal effort is ambiguous, as the sign of the direct period-two effect may vary depending on (initial) ability and wealth. Additional ambiguity is introduced by the two indirect effects that may take either sign on aggregate.

Similar reasoning yields ambiguous effects under policies HC or LC. Under those regimes

[10] Again, note that albeit both posteriors decrease with increased effort the agent's expected ability remains constant as he is more likely to be successful.

the direct period-two effect must be positive for the former policy and negative for the latter as otherwise the principal could increase expected profits by not rehiring any agent. A change in e_1 does not have any indirect effect in case the principal does not rehire his old manager. For LC this is the case if the project was successful and for HC this is the case after a bad period-one outcome. Overall, the indirect effect will be negative for successful agents under HC and ambiguous for unsuccessful ones under LC, such that in each case the sum of the direct and indirect period-two effects may take either sign. Again, the optimal first-period effort might increase or decrease relative to the static problem in both cases.

Finally, in case the principal finds it optimal never to extend the employment contract of an agent, e_1 is trivially the same as in the static problem: If the principal never rehires the agent, changes in the agent's period-two wealth or presumed ability do not affect the principal's earnings. Also, second-period profits are independent of the outcome in period one, so there is no direct effect of first-period effort on second-period surplus. □

Whether the optimal first-period effort level increases or decreases relative to the one-period problem will depend on the actual parameters and the shapes of the utility and the cost function. In the next section we will abstract from such issues by only allowing for binary effort levels and by assuming that the principal always wants to implement effort in the first period. This allows us to characterize the optimal employment policies more closely while preserving the key trade-offs of the more general model.

3.4 The Case of Binary Effort

We have seen that in general, a principal may not be best off by hiring the most able employee he can get hold of if we account for endogenous differences in wealth. If differences in ability are small, a principal cares more about hiring an agent that is easy to motivate than one who has a track record of success. However, fully characterizing the optimal policies is non-trivial: Even for a given prior distribution $F(\tilde{\theta})$ it can be optimal to rehire only successful agents in case effort is either of very high or of very low importance, while for intermediate values of Δ_2 a principal would only rehire unsuccessful managers.

In this section we look at a specific example where the agent's effort choice is binary in order to develop a better understanding of the rehiring strategies that may turn out to be optimal. Whenever the agent exerts effort e, he suffers a non-monetary cost of C. Alternatively, a manager can choose not to exert any effort and does not suffer any disutility from doing so. The agent has a large initial wealth W and a consumption utility of $u(W_t)$. In the interest of simplicity, we assume that $u(W_t) = \sqrt{W_t}$ for all $W_t \geq 0$ and $u(W_t) = -\infty$ otherwise. The appeal of square root utility lies in the simple functional form that allows us to derive explicit conditions for the optimality of the different policies. However, the key trade-offs are the same for all utility functions that satisfy Assumption 1.[11] In particular, once we assume that consumption utility $u(W_t)$ and effort costs are additively separable, CARA utility does no longer simplify the problem faced by the principal: Changes in wealth always affect the cost of incentives and this effect is positive for any function that satisfies Assumption 1.

In order to concentrate on the interesting case, we assume that the principal always finds it worthwhile to implement effort in the first period, i.e., we assume that Δ_1 is sufficiently high. This abstracts from the trivial case in which the agent does not earn any wages in period one and the principal is only interested in screening for ability: He will fire any unsuccessful agent and reemploy any successful agent.[12] Whenever the principal decides to induce effort in period two he has to offer a wage scheme that satisfies

$$v_{kH} - v_{kL} = C/e \ ,$$

$$(\theta_k + e)v_{kH} + (1 - \theta_k - e)v_{kL} - C = v_k \ ,$$

where $v_k \equiv \sqrt{W + w_k}$ and $v_{kl} \equiv \sqrt{W + w_k + w_{kl}}$. Since the wage schedule is uniquely pinned down by the incentive and participation constraints we can directly solve for the

[11] We discuss the key implications of square root utility at the end of Section 3.4.1.

[12] For the sake of simplicity we also abstract from hybrid cases where the principal induces effort for a subset of policies only.

utility an agent has to receive in each state of the world which yields

$$\begin{aligned} v_{kH} &= v_k + \frac{C}{e}(1-\theta_k) \ , \\ v_{kL} &= v_k - \frac{C}{e}\theta_k \ . \end{aligned}$$

In case of a positive outcome the agent is paid a positive bonus and his consumption utility increases. But if profits are low, he pays a fine and ends up with a lower level of utility than before. Using the fact that $w_{kl} = h(v_{kl}) - h(v_k) = v_{kl}^2 - v_k^2$ we can express the expected cost of implementing effort $(\theta_k + e)h(v_{kH}) + (1 - \theta_k - e)h(v_{kL}) - h(v_k)$ as

$$\left(\frac{C}{e}\right)^2 [(\theta_k + e)(1-\theta_k) - e\theta_k] + 2Cv_k \ . \tag{3.2}$$

Alternatively, the principal can hire an agent at zero cost if he does not implement any effort.

The cost of inducing effort depends positively on the agent's wealth v_k. As before, it is more costly to align the incentives of a manager with the ones of the principal if the manager is wealthy rather than "hungry". Additionally, the cost of incentives depends on θ_k, where the sign of this effect is ambiguous. Higher ability increases the probability that the principal has to pay a bonus. But at the same time it allows the principal to reduce the level of wages since the agent anticipates that he is less likely to be punished. Hence, the expected wage payments can increase or decrease in θ_k. Note that this insurance effect of ability is completely independent of the agent's wealth and it does not hinge on the agent's specific risk preferences. Overall, it is not clear if the cost of incentives is increasing in previous success and we will need an additional assumption to ensure that this is the case. The benefit of inducing effort, however, is given by $e\Delta_2$ and is independent of the agent's employment history.

We will denote by $\hat{\Delta}_k$ the lowest value of Δ_2 for which the principal finds it optimal to implement effort, given that he faces an agent with history k. This threshold is implicitly defined by the point at which the benefit of effort $e\Delta_2$ equals its cost as stated in (3.2). Whenever the effect of θ_k on the insurance properties of a contract is not too large, we expect the threshold to increase in the agent's past performance. In this case the principal will only offer incentives to successful agents if the returns to effort are substantial, while for unsuccessful agents he is willing to induce effort even if Δ_2 is low. Henceforth, we will make the following assumption:

Assumption 2 *Conditional on high period-one outcomes, the agent's expected ability is sufficiently small:*

$$\theta_H \leq \frac{1}{2} \ .$$

There are two reasons why the agent's expected ability may be low after a positive outcome in period one. Either the prior concerning the agent's ability θ is low, such that even a positive period one outcome does not leave the principal too optimistic concerning the agent's ability. Or the variance of the prior distribution σ^2 is small. In this case a positive outcome in period one contains little additional information on the agent's ability and Assumption 2 is satisfied for any $\theta < 1/2$. Assumption 2 ensures that the effect of θ_k on the insurance properties of a contract is sufficiently small and we obtain the ordering $\hat{\Delta}_L < \hat{\Delta}_0 < \hat{\Delta}_H$:

Lemma 1 *Under Assumption 2 the lowest values of success for which the principal finds it optimal to implement effort in period two when facing an agent with history $k \in \{L, 0, H\}$ have the "natural ordering"*

$$\hat{\Delta}_L < \hat{\Delta}_0 < \hat{\Delta}_H \ . \tag{3.3}$$

Proof: See the appendix. □

Note that Assumption 2 is only broadly sufficient for the natural ordering to hold. We can check that this ordering also obtains for arbitrary values of θ_H as long as σ^2 is sufficiently small. Moreover, as long as $\hat{\Delta}_L < \hat{\Delta}_0 < \hat{\Delta}_H$ all of our results remain unchanged even if Assumption 2 is violated. Conversely, we can easily derive the implications of a violation of the natural ordering property: If $\hat{\Delta}_L \geq \hat{\Delta}_0$, an unsuccessful agent is not only of lower expected ability, but he is also more costly to motivate and will never be reemployed. By the same line of argument there is no reason to replace a successful agent in case $\hat{\Delta}_0 \geq \hat{\Delta}_H$.

3.4.1 Optimal continuation

In order to see which employment policy is optimal for a given set of parameters, we can look at the principal's decision problem after each of the possible period-one outcomes separately. In the one case a principal faces the choice between reemploying an unsuccessful manager or hiring a new one, while in the other case he chooses between a successful agent and a new one. Since the probability of either period one outcome does not depend on the reemployment policies, looking at the two cases separately is sufficient in order to derive the optimal employment policies.

Continuing after low output

Consider the problem a principal faces after a bad outcome in the first period. In this case he can choose between reemploying the unsuccessful manager or hiring a new agent of unknown ability. The following lemma establishes conditions on the size of profit differentials such that the former is indeed optimal.

Lemma 2 *It is optimal to rehire an unsuccessful agent if and only if* $e > \theta - \theta_L$ *and*

$$\Delta_2 \in [\frac{e}{e-(\theta-\theta_L)}\hat{\Delta}_L, \frac{e}{\theta-\theta_L}(\hat{\Delta}_0 - \hat{\Delta}_L)] \ ,$$

where the interval might be empty. Whenever the interval is non-empty, it contains $\hat{\Delta}_0$.

Proof: The expected profit from continuing with the old manager after low output is given by

$$\underline{\pi}_2 + \theta_L \Delta_2 + e(\Delta_2 - \hat{\Delta}_L)\mathbb{1}_{\{\Delta_2 > \hat{\Delta}_L\}} \ , \tag{3.4}$$

where $\mathbb{1}_{\{.\}}$ denotes the indicator function.

The project makes profits of at least $\underline{\pi}_2$ with certainty. Even though the ability of an unsuccessful agent is strictly lower than θ in expectation, there is still some positive probability θ_L that the firm earns high profits due to the agent's ability. Finally, in case the value of success is sufficiently high, the principal decides to offer a contract that induces the agent to exert effort. The benefit of effort is given by $e\Delta_2$ while the cost of remuneration can be expressed as $e\hat{\Delta}_L$ since the principal is indifferent between inducing effort and not inducing effort for $\Delta_2 = \hat{\Delta}_L$. Similarly, the expected profit when hiring a new agent is given by

$$\underline{\pi}_2 + \theta\Delta_2 + e(\Delta_2 - \hat{\Delta}_0)\mathbb{1}_{\{\Delta_2 > \hat{\Delta}_0\}} \ . \tag{3.5}$$

For new agents, the expected ability is strictly larger than for previously unsuccessful agents. But at the same time, the net profit from inducing effort is smaller than for old managers. While the direct benefit of effort is the same as for unsuccessful managers, the wealth effect makes it more expensive to motivate a new agent. Consequently, it pays to reemploy an unsuccessful agent whenever (3.4) is greater than (3.5), i.e., if

$$-(\theta - \theta_L)\Delta_2 + e(\Delta_2 - \hat{\Delta}_L)\mathbb{1}_{\{\Delta_2 \in [\hat{\Delta}_L, \hat{\Delta}_0]\}} + e(\hat{\Delta}_0 - \hat{\Delta}_L)\mathbb{1}_{\{\Delta_2 > \hat{\Delta}_0\}} > 0 \ . \tag{3.6}$$

Employing an unsuccessful agent always has the disadvantage of lower expected ability as expressed by the first term. Yet, for intermediate values $\Delta_2 \in [\hat{\Delta}_L, \hat{\Delta}_0]$ it pays to induce effort for unsuccessful agents but not for new agents. If e is sufficiently large, this effect can off-set the negative ability effect. A necessary condition for this to be the case is

$$e > \theta - \theta_L \ . \tag{3.7}$$

If this condition is satisfied, the probability of making high profits is larger when hiring an unsuccessful manager. Hence, the appeal of rehiring an unsuccessful agent is increasing in the value of success. So whenever it is optimal to keep an unsuccessful agent instead of hiring a new agent on the market for some Δ_2, it is also optimal to do so for all larger values of success.

However, once $\Delta_2 > \hat{\Delta}_0$ this logic does no longer apply. In this case it pays to induce effort for both types of agents and due to their higher expected ability, new managers are now more likely to earn high profits. Hence, it becomes more attractive to hire a new agent the larger the value of success. While the cost of incentives is lower by $e(\hat{\Delta}_0 - \hat{\Delta}_L)$ for unsuccessful agents, this quantity does not depend on the value of success and loses in relative importance as Δ_2 becomes larger.

Thus, there exists some interval of values for Δ_2 for which the principal prefers to reemploy unsuccessful agents to hiring new ones. The boundaries of this follow immediately by solving for the roots of (3.6) within $[\hat{\Delta}_L, \hat{\Delta}_0]$ and $[\hat{\Delta}_0, \infty)$ if they exist. The second part of the lemma follows from the discussion above. □

Continuing after high output

In a similar fashion we can compare the benefit of hiring a new agent to rehiring an old agent in case the project turned out successful in period one. In this case we have

Lemma 3 *It is optimal to replace a successful agent if and only if $e > \theta_H - \theta$ and*

$$\Delta_2 \in [\frac{e}{e-(\theta_H-\theta)}\hat{\Delta}_0, \frac{e}{\theta_H-\theta}(\hat{\Delta}_H - \hat{\Delta}_0)] \ ,$$

where the interval might be empty. Whenever the interval is non-empty, it contains $\hat{\Delta}_H$.

Proof: The proof follows the same line of argument as the one of Lemma 2: The benefit from rehiring an agent is

$$\underline{\pi}_2 + \theta_H\Delta_2 + e(\Delta_2 - \hat{\Delta}_H)\mathbb{1}_{\{\Delta_2 > \hat{\Delta}_H\}} \ . \tag{3.8}$$

As before, the principal earns low profits for sure and may earn high profits due to the agent's ability. Since previous success is a positive signal on the agent's ability, an old agent has an expected ability that is above θ. Additionally, the principal may choose to induce effort which comes at a cost of $e\hat{\Delta}_H$ but increases the probability of a positive outcome by an additional constant of e. Note that for previously successful agents the cost of inducing effort is higher than for unsuccessful and new agents: Successful managers have already earned a positive bonus in period one and are, therefore, harder to motivate.

Since the value of hiring a new agent is the same as after a low period one outcome, a principal will rehire an old agent if and only if the difference between (3.5) and (3.8) is negative:

$$-(\theta_H - \theta)\Delta_2 + e(\Delta_2 - \hat{\Delta}_0)\mathbb{1}_{\{\Delta_2 \in [\hat{\Delta}_0, \hat{\Delta}_H]\}} + e(\hat{\Delta}_H - \hat{\Delta}_0)\mathbb{1}_{\{\Delta_2 > \hat{\Delta}_H\}} < 0 \ . \tag{3.9}$$

By the same logic as before, a principal will decide to hire the less able (and less wealthy) manager in case the effort effect outweighs the ability effect. A necessary condition for this to be the case is

$$e > \theta_H - \theta \ . \tag{3.10}$$

For intermediate values of success a principal will only induce effort in case he hires a new manager. Whenever (3.10) holds the probability of making high profits is hence larger when hiring a new manager and the appeal of a new manager vis-à-vis a successful manager is increasing in the value of success. However, once the value of success exceeds $\hat{\Delta}_H$ the principal implements effort regardless of which type of agent he hires. Hence, any further increase in the value of success makes it more attractive to retain the successful manager: Due to his higher expected ability, a previously successful agent will earn high profits with a larger probability than a new agent. By the same reasoning as above, hiring the less wealthy (and less able) agent is optimal for an intermediate interval of Δ_2. □

Optimal policies

So far we have looked at the decisions a principal takes after either period-one outcome in separation. The results are depicted in Figure 3.2, where the areas outside of the respective triangles describe situations in which the principal chooses to hire the most able employee that is available. We can now translate these regions into optimal employment policies, such as shown in Figure 3.3. If the two triangles in Figure 3.2 overlap, the principal finds it optimal to continue with his current employee only after low period-one outcomes (LC). Similarly, the remaining areas in the two triangles correspond to situations where the principal either continues employment irrespective of profits (AC) or never (NC). In all other cases, the principal chooses to renew an agent's contract only in case he was successful. Characterizing the respective areas more closely allows us to establish the following relationship between the optimal policy and the value of success:

Proposition 5 *For each policy other than HC the set of values of Δ_2 for which a policy is optimal can be described by one (possibly empty) interval. The ordering of these intervals is always such that a policy of never continuing employment contracts (NC) is optimal for larger values of success Δ_2 than policies LC and AC and policy LC is optimal for larger values of success than policy AC.*

Proof: The argument of the previous section yields that the lowest value of success for which it is optimal to hire a new agent after high output is strictly above the lowest value of success for which a principal wants to rehire an unsuccessful agent: The first threshold lies between $\hat{\Delta}_0$ and $\hat{\Delta}_H$ while the latter must lie below $\hat{\Delta}_0$. Moreover, we can calculate the difference of the two upper bounds provided in Lemma 3 and 2 to be

$$\frac{C^2}{\theta + e}\left(2\frac{e}{\theta_H - \theta} - \frac{\theta - \theta_L}{e}\right).$$

From the necessary conditions (3.7) and (3.10) this is strictly positive. Thus the interval for which it is optimal to continue with an agent after low output starts earlier and ends earlier (i.e., at lower values of success) than the interval for which it is optimal to hire a new agent after high output. For some intermediate values of success the intervals might overlap such that it is optimal to reemploy unsuccessful managers but to replace successful ones. In this case the principal employs policy LC and expected profits are given by the probability-weighted sum of (3.4) and (3.5). For slightly lower values of Δ_2 it is optimal to employ policy AC and profit equals the weighted sum of (3.4) and (3.8). For slightly larger values of success NC is optimal and in case of extremely high or low values of success HC is optimal. Profits in these cases are derived similarly. Finally, in case the two intervals do not overlap, NC is still optimal for strictly larger values of Δ_2 than AC and LC is never optimal. $\square$

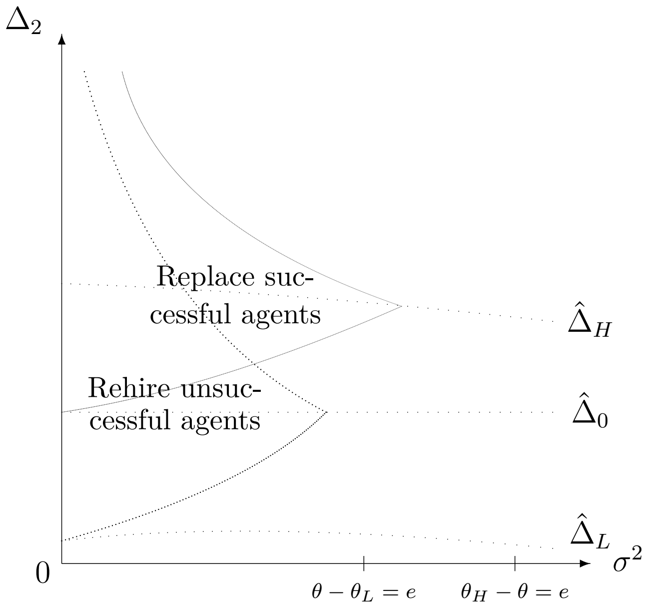

Figure 3.2: Optimal continuation decisions in the σ^2-Δ_2-space.

So far we have concentrated on how the optimal reemployment decision depends on the value of success Δ_2. Let us now investigate the impact of uncertainty with respect to the agent's ability. We can see in Figure 3.2 that the larger the uncertainty with respect to the agent's ability, the more likely HC is to be optimal relative to all other policies: Larger uncertainty makes period one output more informative and thus favors the policy conditioning most severely on ability. Conversely, for $\sigma^2 \to 0$ all reemployment policies turn out to be optimal for some values of success as we will show in the next section.

Note that Figures 3.2 and 3.3 depict a situation where $\theta_H - \theta < \theta - \theta_L$, i.e., the update on ability after high output is smaller in absolute size than the update after low output. Under such a parameter constellation the ability advantage of a successful agent may be off-set by the effort advantage of a new agent for intermediate levels of volatility, while the disadvantage of an unsuccessful agent is still too large for the agent to make up for it via increased effort. Hence, the tip of the upper triangle will typically lie further to the right than the tip of the lower triangle.[13] The opposite situation will usually appear if

[13]We refrain from stating the precise technical condition for the relative position of the tips for it bears no further intuitive interpretation and is rather involved.

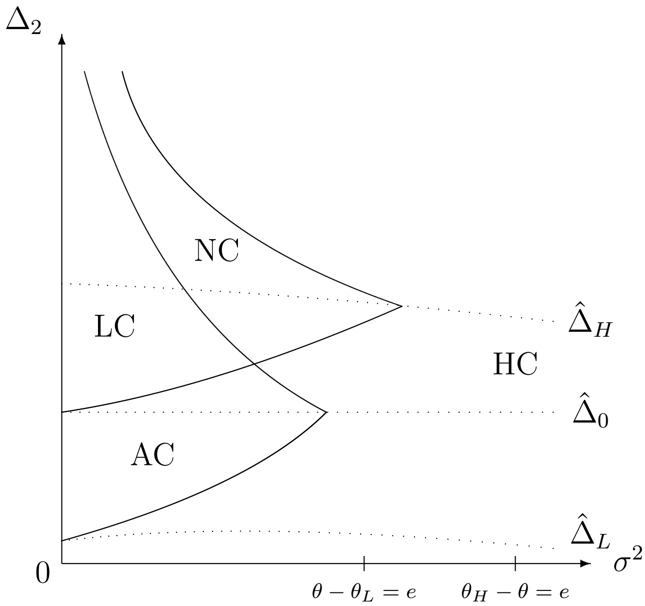

Figure 3.3: Optimal policies in the σ^2-Δ_2-space.

$\theta_H - \theta > \theta - \theta_L$: The expected ability of successful agents is so high that it always pays to reemploy these agents, while unsuccessful agents get a chance to make up for failure for intermediate levels of Δ_2. In this case the tip of the lower triangle will lie further to the right than the one of the upper triangle.

While the assumption of square root utility is not crucial for most of the analysis, it does influence our results in two ways: First, it allows us to derive sufficient conditions for the "natural ordering" to hold that we discussed at the beginning of this section. More importantly, it implies that the largest value of success for which it is optimal to replace a successful manager is larger than the largest value of success for which rehiring unsuccessful agents is optimal. This result is driven by the fact that for square root utility the cost of incentives is linear in v_k. If the cost of compensation is sufficiently concave in v_k this may no longer be the case and AC may in some cases be optimal for larger values than LC. Nevertheless, comparing the two upper bounds from Lemma 2 and 3 yields that our description of optimal policies will hold true for any utility function that satisfies Assumption 1 if the expected cost of giving incentives when always continuing with an old agent exceed the cost of motivating a new agent, i.e., whenever $(\theta+e)\hat{\Delta}_H+(1-\theta-e)\hat{\Delta}_L \geq \hat{\Delta}_0$.

3.4.2 The cost of effort

We have seen that whenever policies other than hiring the most able manager available are optimal, they follow a particular order. However, we have said very little on the optimality of the different policies itself. In this section we will see that the optimality of policies other than HC crucially depends on the cost of effort C. The larger the cost of effort, the more expensive it is to compensate an agent for the disutility arising from his work. Since previously successful agents value monetary compensation less strongly than new or "hungry" managers, this effect is largest for agents with a positive track record. So the larger C, the more likely policies are to be optimal that involve hiring agents of inferior ability.

Before proceeding, let us recap the necessary conditions for optimality of the different policies that we have discussed in the last section.

Lemma 4 *(i) AC is the optimal policy for some values of success* Δ_2 *if and only if it is optimal at* $\hat{\Delta}_0$*. A necessary condition for this is* $e > \theta - \theta_L$.

(ii) NC is the optimal policy for some values of success Δ_2 *if and only if LC or NC itself is optimal at* $\hat{\Delta}_H$*. A necessary condition for this is* $e > \theta_H - \theta$.

(iii) LC can only be optimal for any value of success if AC is optimal at $\hat{\Delta}_0$*. A necessary condition for optimality of LC is* $e > \max\{\theta - \theta_L, \theta_H - \theta\}$.

(iv) HC is optimal for all sufficiently low and high levels of Δ_2*. It is globally optimal if and only if it is optimal at* $\hat{\Delta}_0$ *and at* $\hat{\Delta}_H$*. A sufficient condition for this is* $e < \min\{\theta - \theta_L, \theta_H - \theta\}$.

Proof: See the appendix. □

Using those results, we can now derive necessary and sufficient conditions on the cost of effort for the different policies to be optimal.

Proposition 6 *For each policy $i \in \{AC, NC, LC\}$ there is a lower bound $\underline{C}^i$ on the cost of effort C such that policy i is optimal for some values of success if and only if $C \geq \underline{C}^i$. However, in some cases we may have $\underline{C}^i = \infty$.*

Conversely, for any given effort cost C there is a threshold $\tilde{\sigma}_i^2$ such that policy i is optimal for some values of success if $\sigma^2 < \tilde{\sigma}_i^2$.

Proof: See the appendix. □

Let us start by concentrating on the question if it can be optimal to always continue a relationship (AC). Indeed, a policy of accepting failure is commonplace in many industries and seems to be empirically highly relevant. Generally, a policy of never terminating employment relationships becomes more attractive if the cost of effort C is high: While the attraction of HC relies on the fact that it enables the principal to weed out the least able employees, adopting a policy of lenience with regard to past failure allows the principal to reduce the cost of supplying incentives. This effect becomes more pronounced the larger the private cost of effort that the agent has to be compensated for. We can hence derive a lower bound on the cost of effort that guarantees that tolerance for failure is indeed optimal for some profit differentials. From Lemma 4 we know that AC is optimal for some values of success if and only if it is optimal for $\hat{\Delta}_0$. Moreover, the only other policy that may be optimal at $\hat{\Delta}_0$ is to reemploy an agent if and only if he has been successful. So the lower bound on C is defined by the cost of effort that makes a principal indifferent between policies AC and HC at a value of success of $\hat{\Delta}_0$.

Similar reasoning can be extended to all other policies that prescribe hiring an agent of inferior ability. This allows us to define lower bounds on C for policies LC and NC to be optimal as well. The technical proof is relegated to the appendix.

We can easily derive that $\underline{C}^{LC}$ must be higher than $\underline{C}^{AC}$ and $\underline{C}^{NC}$. A policy of only rehiring unsuccessful managers can only be optimal for some values of success if policies

AC and NC are so for some other values of success, too. Yet, if costs are sufficiently high, it pays to always employ the least wealthy agent since the manager's wealth has a significant influence on the cost of compensation.

Note that for high levels of σ^2 the thresholds $\underline{C}^i$ may be infinitely high, in which case the corresponding policy never turns out to be optimal irrespective of C. This will always be the case if the necessary conditions mentioned in Lemma 4 are violated. In this case the ability effect is sufficiently strong such that it can never be off-set by implementing effort. Hence, the cost of effort does not affect reemployment decisions. Conversely, the second part of Proposition 6 tells us that for low levels of uncertainty with respect to the agent's ability all policies turn out to be optimal for some values of success. If the ability effect becomes arbitrarily small, any cost advantage of employing inferior agents suffices in order to make hiring a low-ability agent optimal. This result is similar to the observation in Section 3.3 that LC always becomes optimal as uncertainty about the agent's ability vanishes. However, there is one important difference between the case of continuous effort and the example of binary effort that we consider here. In case of binary effort, the principal never implements any effort at all if Δ_2 is low. For these values of success, the cost of incentives is irrelevant and the principal finds it optimal to employ policy HC even as $\sigma^2 \to 0$.

Finally, Proposition 6 implies that keeping only successful agents will be the optimal policy for any value of success if the cost of effort is low:

Corollary 2 *Policy HC will be optimal for any values of success if the cost of effort C is below* $\min\{\underline{C}^{AC}, \underline{C}^{NC}\}$.

If the agent has little disutility from exerting effort, the principal faces very similar costs for implementing effort regardless of the agent's employment history. Hence, ability becomes the relevant criterion for employment and the principal will always hire the agent with the highest expected ability. In case the project was successful in the first period this is his current employee, while in case of a low period-one outcome the principal hires a new manager.

3.5 Conclusion

We have considered optimal reemployment decisions in a principal-agent model where agents differ with respect to their ability and are prone to wealth effects. The basic trade-off is that previously unsuccessful agents are less likely to be of high ability but at the same time they have low wealth and are hence less costly to motivate in the future. This explains why it may be optimal for an employer to tolerate failure and to renew the contracts of unsuccessful managers. The idea that it may not be optimal to keep only the most successful managers is consistent with a large body of empirical literature that finds low correlations between a firm's success and forced CEO turnover.

The main insight from our analysis is that the incentives to fire unsuccessful managers are not monotonically increasing in the value that a success has to the principal. Instead, a principal will always find it optimal to retain only the most able employees if the additional benefit of a success is very high or very low, but he may be lenient vis-à-vis failure if the importance of effort lies between these extremes. In this case the cost of managerial compensation is key in determining firm profits and it pays to hire managers that are easy to motivate. Moreover, hiring unsuccessful managers is particularly appealing if uncertainty with respect to managerial skills is low and if the the agent has to invest a high amount of energy into unobservable tasks: If uncertainty is low, the previous track record of an agent contains little information on his ability and it is optimal to hire "hungry" managers. If the cost that an agent has to bear in case he exerts effort is high, differences in the agent's responsiveness to monetary incentives are crucial, which makes it again optimal to hire agents that are easy to motivate.

Our analysis shows that the optimal length of employment relationships depends on a number of factors that are more subtle than is typically assumed. Additionally, the optimal degree of tolerance for failure is likely to change over time if we allow for relationships that last for more than two periods. With increasing tenure, the principal obtains little additional information on the agent's ability in any given period and he hence may become more lenient towards an unsuccessful agent over time. While not covered in this dissertation, the optimal degree of leniency in such long-term relationships is an interesting area for future research.

A3 Appendix

A3.1 Mathematical appendix

Proof of Proposition 3

Proof: Assume that in period one the principal has offered some contract that implements a strictly positive level of effort e_1. Depending on the period-one outcome this contract results in utility levels of $v_H > v_L$. Now consider the first order condition for effort in period two

$$(\overline{\pi}_2 - h(v_{kH})) - (\underline{\pi}_2 - h(v_{kL})) - E(h'(v_{kl})v'_{kl}(e^*)) = 0 \ .$$

If $\Delta_2 \to 0$ the level of effort a principal implements in period two must go to zero regardless of the period-one outcome $k \in \{H, 0, L\}$. This implies that irrespective of the agent's wealth, the principal implements negligible incentives at costs that are close to zero. So for a given period-one contract, differences in wealth do not matter and the principal always employs the most able manager. Since this holds for any period-one contract, HC must be optimal as $\Delta_2 \to 0$.

Conversely, consider the situation where $\Delta_2 \to \infty$ and the principal offers an arbitrary period-one contract. By the first order condition we know that in this case a principal chooses to set maximal incentives irrespective of k, i.e., $e_2 \to \overline{e}$. This implies that the expected probability of a project being successful in period two is highest for managers that were previously successful and lowest for unsuccessful managers. Hence, using the envelope theorem we can show that the effect of an increase in Δ_2 on expected profits is largest for policy HC and this policy will be optimal for sufficiently large values of success. Again, since this holds for any period-one contract, HC must be optimal as $\Delta_2 \to \infty$.

Now, let us fix some Δ_2 and consider the effect of a change in σ^2. Again assume that the principal offers an arbitrary period-one contract. From Proposition 1 we know that for a given ability, the principal would always want to hire the least wealthy agent. So as $\sigma^2 \to 0$ a policy of only reemploying unsuccessful managers (LC) will become optimal.

Moreover, this holds for any $\Delta_2 > 0$ and any period-one contract. $\square$

Proof of Lemma 1

Proof: Applying the definition of $\hat{\Delta}_k$ as the level of Δ_2 at which the principal is just willing to implement effort when facing an agent with history k we get

$$e\hat{\Delta}_k = \left(\frac{C}{e}\right)^2 [(\theta_k + e)(1-\theta_k) - e\theta_k] + 2Cv_k \ ,$$

where the right hand side represents the cost of implementing effort. Using the fact that $v_H = v_0 + \frac{C}{e}(1-\theta)$ and $v_L = v_0 - \frac{C}{e}\theta$ this gives us

$$\begin{aligned}
\hat{\Delta}_H &> \hat{\Delta}_0 \Leftrightarrow e > \frac{\theta(1-\theta) - \theta_H(1-\theta_H)}{2(1-\theta_H)} \ , \\
\hat{\Delta}_0 &> \hat{\Delta}_L \Leftrightarrow e > \frac{\theta_L(1-\theta_L) - \theta(1-\theta)}{2\theta_L} \ .
\end{aligned}$$

If $\theta_H \leq \frac{1}{2}$ the numerators of both conditions are negative and the inequalities are trivially satisfied since $e > 0$. $\square$

Proof of Lemma 4

Proof: *(i)* Necessity: *AC* includes rehiring unsuccessful agents. By Lemma 2 this can only be optimal if it is optimal at $\hat{\Delta}_0$. Sufficiency: The only other policy including reemployment of unsuccessful agents is *LC*. But *LC* also incorporates replacing successful agents, which by Lemma 3 cannot be optimal at $\hat{\Delta}_0$, as the lower bound of the stated interval is larger than this threshold. The necessary condition for optimality at $\hat{\Delta}_0$ is equivalent to (3.7) and has been derived in the main text.

(ii) Necessity: *NC* includes rehiring unsuccessful agents. By Lemma 3 this can only be optimal if it is optimal at $\hat{\Delta}_{\mathrm{H}}$. Sufficiency: The only other policy including replacement of successful agents is *LC*. As derived in the main text the upper bound for rehiring unsuccessful agents lies strictly below the upper bound for replacing successful agents. Thus if *LC* is optimal at $\hat{\Delta}_{\mathrm{H}}$ *NC* must become so for higher

profit differentials. The necessary condition is equivalent to (3.10) and has been derived in the main text as well.

(iii) Follows from the proof of (*i*). The necessary condition is the stricter of the two conditions (3.7) and (3.10) for reemploying inferior agents.

(iv) The first part is a direct consequence of Lemma 2 and 3. The remainder follows from (*i*), (*ii*), and (*iii*). □

Proof of Proposition 6

Proof: We know from Lemma 4 that AC is optimal for some values of success if and only if it is optimal at $\hat{\Delta}_0$. Moreover, it is easy to verify that the only other policy that may be optimal at $\hat{\Delta}_0$ is HC. As the two policies only differ with respect to reemployment after low output this is equivalent to (3.6) being positive at $\Delta_2 = \hat{\Delta}_0$. We can restate this as

$$e\hat{\Delta}_L \leq (\theta_L + e - \theta)\hat{\Delta}_0 \ .$$

Substituting expression (3.2) gives us

$$e\left(\left(\frac{C}{e}\right)^2 [(\theta_L + e)(1-\theta_L) - e\theta_L - 2e\theta] + 2Cv_0\right) \leq \\ (\theta_L + e - \theta)\left(\left(\frac{C}{e}\right)^2 [(\theta + e)(1-\theta) - e\theta] + 2Cv_0\right)$$

or simply

$$2e^2 v_0(\theta - \theta_L) \leq C\left[2e^2\theta_L - (\theta - \theta_L)\left[\theta(1-\theta) - e(\theta - \theta_L)\right]\right] \ .$$

For $2e^2\theta_L \leq (\theta - \theta_L)\left[\theta(1-\theta) - e(\theta - \theta_L)\right]$ this condition is violated for any positive level of costs C. In this case there is no finite cost level such that AC is preferred to HC. Thus, we can define

$$\underline{C}^{AC} := \begin{cases} \frac{2(\theta-\theta_L)e^2 v_0}{2e^2\theta_L - (\theta-\theta_L)[\theta(1-\theta) - e(\theta-\theta_L)]} & \text{if the denominator is positive,} \\ \infty & \text{else,} \end{cases}$$

such that AC is optimal for a non-empty interval of values for Δ_2 if and only if $C \geq \underline{C}^{AC}$.

For sufficiently low levels of uncertainty there will always be a non-empty interval of values for Δ_2 for which AC is optimal. As $\sigma^2 \to 0$ we have $\theta_L \nearrow \theta$, so $\underline{C}^{AC}$ is finite and approaches zero. Thus, for any positive C there is a threshold $\tilde{\sigma}^2_{AC}$ such that $\underline{C}^{AC} \leq C$ whenever $\sigma^2 \leq \tilde{\sigma}^2_{AC}$.

Next, let us turn to the question of when it is optimal never to continue employment. By Lemma 4 we know that policy NC is optimal for some values of success if and only if LC or NC is optimal at $\hat{\Delta}_H$. Hence, NC is optimal for some values of success if and only if it is optimal to let go of a successful manager at $\hat{\Delta}_H$, which is the case whenever (3.9) is non-negative at $\hat{\Delta}_H$:

$$-(\theta_H - \theta)\hat{\Delta}_H + e(\hat{\Delta}_H - \hat{\Delta}_0) \quad \geq \quad 0 \ .$$

Plugging in the expressions for $\hat{\Delta}_H$ and $\hat{\Delta}_0$ and simplifying, this is equivalent to

$$2e^2(\theta_H - \theta)v_0 \leq C\left[2e^2(1-\theta_H) - (\theta_H - \theta)[\theta_H(1-\theta_H) + e((1-\theta_H) + (1-\theta))]\right] \ .$$

If the term on the right hand side is negative, NC can not be optimal for any cost level C. If it is positive, we get a lower bound on C. Again, we can define the lower bound by

$$\underline{C}^{NC} \quad := \quad \begin{cases} \frac{2e^2(\theta_H-\theta)v_0}{2e^2(1-\theta_H)-(\theta_H-\theta)[\theta_H(1-\theta_H)+e((1-\theta_H)+(1-\theta))]} & \text{if the denominator is positive,} \\ \infty & \text{else.} \end{cases}$$

For $\sigma^2 \to 0$ ability levels converge, i.e., $\theta_H \searrow \theta$. Again, this implies that $\underline{C}^{NC} \searrow 0$ and the inequality holds for arbitrary levels of C.

Third, a policy of only continuing employment with unsuccessful agents will be optimal if the upper bound for reemploying unsuccessful agents as defined in Lemma 2 lies above the lower bound for replacing successful agents as defined in Lemma 3, i.e., if

$$\frac{e}{\theta - \theta_L}(\hat{\Delta}_0 - \hat{\Delta}_L) \geq \frac{e}{e - (\theta_H - \theta)}\hat{\Delta}_0 \ .$$

Plugging in the expressions for $\hat{\Delta}_L$ and $\hat{\Delta}_0$ and simplifying yields

$$2e^2(\theta-\theta_L)v_0 \;\; \leq \;\; C[2e^2\theta_L-(\theta-\theta_L)[(1-\theta_H)\theta_H-(\theta_H-\theta)\theta_L]+e(\theta^2+\theta_L^2-2\theta_L\theta_H)] \; .$$

Again this cannot be satisfied if the factor on the right hand side is negative and we get the following lower bound for the cost C:

$$\underline{C}^{LC} \;\; := \;\; \begin{cases} \frac{2e^2(\theta-\theta_L)v_0}{2e^2\theta_L-(\theta-\theta_L)[(1-\theta_H)\theta_H-(\theta_H-\theta)\theta_L]+e(\theta^2+\theta_L^2-2\theta_L\theta_H)} & \text{if the denominator is positive,} \\ \infty & \text{else.} \end{cases}$$

As above, as $\theta-\theta_L \to 0$ and $\theta_H-\theta \to 0$ when $\sigma^2 \to 0$, the lower bound converges to zero. The thresholds $\tilde{\sigma}^2_{NC}$ and $\tilde{\sigma}^2_{LC}$ are derived analogously to $\tilde{\sigma}^2_{AC}$. This completes the proof. $\square$

A3.2 Long-term contracts

We intend to show in this appendix that the series of short-term contracts that we considered in this chapter is equivalent to a setting that allows for long-term contracts. In order to do so, we only need to show that any long-term contract can (and will) be replicated by a series of short term contracts at the same cost. Since long-term contracts must be weakly more attractive than short-term contracts this is sufficient to show equivalence.

Assume that a long-term contract implements effort levels $(\hat{e}_1, \hat{e}_2^H, \hat{e}_2^L)$ where $\hat{e}_2^k$ is the effort level that a principal implements after a period-one outcome of k. In order to simplify notation, we will treat cases where the principal does not rehire a manager after an outcome k as if he implemented zero effort: $\hat{e}_2^k = 0$. This is without loss of generality since doing so is costless for the principal. If he hires a new agent, the contract of the agent that is newly hired in period two is trivially a short-term contract and will hence be the same no matter if we allow for long-term contracts or not. A long-term contract can be fully characterized by the consumption utility that an agent receives for any given j and k. We will denote these levels of utility by v_{kj}^{LT}. Incentive compatibility implies that the contract must satisfy

$$(\theta_H + \hat{e}_2^H)v_{HH}^{LT} + (1 - \theta_H - \hat{e}_2^H)v_{HL}^{LT} - C(\hat{e}_2^H) - \\ \Big[(\theta_L + \hat{e}_2^L)v_{LH}^{LT} + (1 - \theta_L - \hat{e}_2^L)v_{LL}^{LT} - C(\hat{e}_2^L)\Big] = C'(\hat{e}_1)$$

$$v_{LH}^{LT} - v_{LL}^{LT} = C'(\hat{e}_2^L)$$
$$v_{HH}^{LT} - v_{HL}^{LT} = C'(\hat{e}_2^H)$$

and the participation constraint requires that

$$[\theta + \hat{e}_1]\Big[(\theta_H + \hat{e}_2^H)v_{HH}^{LT} + (1 - \theta_H - \hat{e}_2^H)v_{HL}^{LT} - C(\hat{e}_2^H)\Big] + \\ [1 - \theta - \hat{e}_1]\Big[(\theta_L + \hat{e}_2^L)v_{LH}^{LT} + (1 - \theta_L - \hat{e}_2^L)v_{LL}^{LT} - C(\hat{e}_2^L)\Big] - C(\hat{e}_1) = v_0 \; .$$

The only reason why a long-term contract mighty be preferable to a short-term contract is that the period-two participation constraints only need to be satisfied in expectation.

However, since the agent's expected utility after either period-one outcome is fully pinned down by the incentive compatibility constraint for period-one effort, the principal can never exploit this additional degree of freedom and he can equivalently resort to short-term contracts: In a setting with short-term contracts we must have

$$v_H - v_L = C'(\hat{e}_1)$$
$$v_{LH} - v_{LL} = C'(\hat{e}_2^L)$$
$$v_{HH} - v_{HL} = C'(\hat{e}_2^H)$$

if the contract implements the same levels of effort. Moreover, the principal will push the agent down to his reservation utility in period two and offer contracts that have $(\theta_k + \hat{e}_2^k)v_{kH} + (1 - \theta_k - \hat{e}_2^k)v_{kL} - C(\hat{e}_2^k) = v_k$. Similarly, at the beginning of period one the principal will offer a contract that has $(\theta + \hat{e}_1)v_H + (1 - \theta - \hat{e}_1)v_L - C(\hat{e}_1) \geq v_0$. Since the principal prefers a less wealthy agent in period two, it is easy to check that this condition will always be binding. By substituting the participation constraint for period two into the period-one participation constraint and the period-one incentive compatibility constraint we get the same constraints as in the case of long-term contracts. Since the four utility levels v_{HH}, v_{HL}, v_{LH} and v_{LL} are fully pinned down by the constraints this implies that the agent will have the same levels of wealth at the end of period two irrespective of whether we allow for long-term contracts or not. It follows directly that the principal makes the same level of profit if he offers short-term contracts. Finally, it is easy to check that the principal does indeed make the same effort and employment decisions under short-term contracts, i.e., there are no time-inconsistencies. The principal can replicate *any* long-term contract at the same cost. Hence, if the long-term contract was optimal, the principal must find it optimal to implement the same level of effort and to make the same reemployment decisions in a setting with short-term contracts. Since period-one compensation is independent of period-two outcomes, it does not distort the principal's choice between different policies in period two once period-one payments have been made.

Bibliography

Alchian, A. A. and H. Demsetz (1972): "Production, Information Costs, and Economic Organization," *American Economic Review*, 62, 777–795.

Anderson, A. and L. Smith (2010): "Dynamic Matching and Evolving Reputations," *Review of Economic Studies*, 77, 3–29.

Anderson, S. P. and S. Coate (2005): "Market Provision of Broadcasting: A Welfare Analysis," *Review of Economic Studies*, 72, 947–972.

Armstrong, M. (2006): "Competition in Two-Sided Markets," *RAND Journal of Economics*, 37, 668–691.

Bar-Isaac, H. (2007): "Something to Prove: Reputation in Teams," *RAND Journal of Economics*, 38, 495–511.

Bayes, T. (1763): "An Essay towards Solving a Problem in the Doctrine of Chances." *Philosophical Transactions*, 53, 370–418.

Belleflamme, P. and M. Peitz (2010): *Industrial Organization: Markets and Strategies*, Cambridge: Cambridge University Press.

Biais, B., T. Mariotti, J.-C. Rochet, and S. Villeneuve (2010): "Large Risks, Limited Liability, and Dynamic Moral Hazard," *Econometrica*, 78, 73–118.

Bolton, P. and M. Dewatripont (1994): "The Firm as a Communication Network," *Quarterly Journal of Economics*, 109, 809–839.

Bolton, P. and C. Harris (2000): "Strategic Experimentation: The Undiscounted Case," in *Incentives, Organizations and Public Economics – Papers in Honour of Sir James Mirrlees*, ed. by P. Hammond and G. Myles, Oxford: Oxford University Press, 53–68.

Breton, M., P. St-Amour, and D. Vencatachellum (2003): "Dynamic Production Teams with Strategic Behavior," *Journal of Economic Dynamics and Control*, 27, 875–905.

Brick, I. E., O. Palmon, and J. K. Wald (2012): "Too Much Pay-Performance Sensitivity?" *Review of Economics and Statistics*, 94, 287–303.

Chade, H. and V. N. Vera de Serio (2011): "Wealth Effects and Agency Costs," Mimeo.

Che, Y.-K. and S.-W. Yoo (2001): "Optimal Incentives for Teams," *American Economic Review*, 91, 525–541.

Coase, R. H. (1937): "The Nature of the Firm," *Economica*, 4, 386–405.

Denis, D. J., D. K. Denis, and A. Sarin (1997): "Ownership Structure and Top Executive Turnover," *Journal of Financial Economics*, 45, 193–221.

Easley, D. and N. M. Kiefer (1988): "Controlling a Stochastic Process with Unknown Parameters," *Econometrica*, 56, 1045–1064.

Facebook (2012): "Registration Statement under The Securities Act of 1933," *www.sec.gov/Archives/edgar/data/1326801/000119312512034517/d287954ds1.htm*, Published by The Securities and Exchange Commission, February 1, 2012.

Fudenberg, D., B. R. Holmström, and P. R. Milgrom (1990): "Short-term Contracts and Long-term Agency Relationships," *Journal of Economic Theory*, 51, 1–31.

Gibbons, R. and K. J. Murphy (1992): "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," *Journal of Political Economy*, 100, 468–505.

Grossman, S. J. and O. D. Hart (1983): "An Analysis of the Principal-Agent Problem," *Econometrica*, 51, 7–45.

Hadlock, C. J. and G. B. Lumer (1997): "Compensation, Turnover, and Top Management Incentives: Historical Evidence," *Journal of Business*, 70, 153–187.

Harrington Jr., J. E. (1995): "Experimentation and Learning in a Differentiated-Products Duopoly," *Journal of Economic Theory*, 66, 275–288.

Holmström, B. R. (1982): "Moral Hazard in Teams," *Bell Journal of Economics*, 13, 324–340.

——— (1999): "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66, 169–182.

Holmström, B. R. and P. R. Milgrom (1987): "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55, 303–328.

Huson, M. R., R. Parrino, and L. T. Starks (2001): "Internal Monitoring Mechanisms and CEO Turnover: A Long-Term Perspective," *Journal of Finance*, 56, 2265–2297.

Jenter, D. and K. Lewellen (2010): "Performance-Induced CEO Turnover," Mimeo.

Jeon, S. (1996): "Moral Hazard and Reputational Concerns in Teams: Implications for Organizational Choice," *International Journal of Industrial Organization*, 14, 297–315.

Kaiser, U. and J. Wright (2006): "Price Structure in Two-sided Markets: Evidence from the Magazine Industry," *International Journal of Industrial Organization*, 24, 1–28.

Kaplan, S. N. and B. A. Minton (2006): "How has CEO Turnover Changed? Increasingly Performance Sensitive Boards and Increasingly Uneasy CEOs," *National Bureau of Economic Research Working Paper No. 12465*.

Keller, G. and S. Rady (1999): "Optimal Experimentation in a Changing Environment," *Review of Economic Studies*, 66, 475–507.

——— (2003): "Price Dispersion and Learning in a Dynamic Differentiated-Goods Duopoly," *The RAND Journal of Economics*, 34, 138–165.

Kim, E. H., A. Morse, and L. Zingales (2009): "Are Elite Universities losing their Competitive Edge?" *Journal of Financial Economics*, 93, 353 – 381.

Knight, F. H. (1921): *Risk, Uncertainty, and Profit*, Boston: Houghton Mifflin Co.

Landier, A. (2006): "Entrepreneurship and the Stigma of Failure," Mimeo.

Liptser, R. S. and A. N. Shiryayev (1977): *Statistics of Random Processes I: General Theory*, Applications of Mathematics, New York: Springer Verlag.

Manso, G. (2011): "Motivating Innovation," *Journal of Finance*, 66, 1823–1860.

Marschak, J. and R. Radner (1972): *Economic Theory of Teams*, Cowles Foundation for Research in Economics at Yale University, Monograph 22, New Haven: Yale University Press.

McAfee, R. P. and J. McMillan (1991): "Optimal Contracts for Teams," *International Economic Review*, 32, 561–577.

Meyer, M. A. (1994): "The Dynamics of Learning with Team Production: Implications for Task Assignment," *Quarterly Journal of Economics*, 109, 1157–1184.

Mirman, L. J., L. Samuelson, and E. E. Schlee (1994): "Strategic Information Manipulation in Duopolies," *Journal of Economic Theory*, 62, 363–384.

Mookherjee, D. (1984): "Optimal Incentive Schemes with Many Agents," *Review of Economic Studies*, 51, 433–446.

Murphy, K. J. (1999): "Executive Compensation," in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Amsterdam: Elsevier, vol. 3, 2485–2563.

Nocke, V., M. Peitz, and K. Stahl (2007): "Platform Ownership," *Journal of the European Economic Association*, 5, 1130–1160.

Prescott, E. C. (1972): "The Multi-Period Control Problem Under Uncertainty," *Econometrica*, 40, 1043–1058.

Rochet, J.-C. and J. Tirole (2003): "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 1, 990–1029.

——— (2006): "Two-Sided Markets: A Progress Report," *RAND Journal of Economics*, 37, 645–667.

Rothschild, M. (1974): "A Two-armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 9, 185–202.

Rysman, M. (2004): "Competition Between Networks: A Study of the Market for Yellow Pages," *Review of Economic Studies*, 71, 483–512.

——— (2009): "The Economics of Two-Sided Markets," *Journal of Economic Perspectives*, 23, 125–143.

Sah, R. K. and J. E. Stiglitz (1986): "The Architecture of Economic Systems: Hierarchies and Polyarchies," *American Economic Review*, 76, 265–298.

Spear, S. and C. Wang (2005): "When to Fire a CEO: Optimal Termination in Dynamic Contracts," *Journal of Economic Theory*, 120, 239–256.

Thiele, H. and A. Wambach (1999): "Wealth Effects in the Principal Agent Model," *Journal of Economic Theory*, 89, 247–260.

Tian, X. and T. Y. Wang (2011): "Tolerance for Failure and Corporate Innovation," *Review of Financial Studies*, forthcoming.

Warner, J. B., R. L. Watts, and K. H. Wruck (1988): "Stock Prices and Top Management Changes," *Journal of Financial Economics*, 20, 461–492.

Curriculum Vitae

seit Oktober 2007	Doktorandenstudium in Ökonomie Munich Graduate School of Economics, LMU München
Oktober 2005 - Februar 2007	Unternehmensberatung The Boston Consulting Group GmbH, München
Oktober 2000 - August 2005	Diplomstudium in mathematischer Finanzökonomie Universität Konstanz
August 1999 - Juni 2000	Zivildienst Evangelische Altenpflege, Wuppertal
Juni 1999	Abitur Carl-Duisberg-Gymnasium, Wuppertal
August 1980	Geboren in Wuppertal, Deutschland