Essays on Financial Economics and the Cost of Incentives

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

2012

vorgelegt von

Caspar Siegert

Referent:	Prof. Dr. Klaus M. Schmidt
Korreferent:	Prof. Jacques Crémer, PhD
Promotionsabschlussberatung:	7. November 2012

Datum der mündlichen Prüfung: 23. Oktober 2012

Namen der Berichterstatter: Klaus M. Schmidt, Jacques Crémer, Florian Englmaier

Acknowledgements

First of all, I would like to thank Klaus Schmidt for being an excellent advisor and for both, challenging and encouraging me during the work on my thesis. He introduced me to the field of contract theory and shaped my research interests. I would also like to thank Jacques Crémer for inviting me to the Toulouse School of Economics and for serving as my local advisor during that time. He was an outstanding advisor and made my time in Toulouse very productive and memorable. I am grateful to Florian Englmaier who throughout the years gave me valuable advice concerning my research and agreed to join my thesis committee. Moreover, I would like to thank Niko Matouschek for inviting me to visit the Kellogg School of Management and for the opportunity to discuss various research projects with him.

I am highly indebted to my co-authors Piers Trepper and Robert Ulbricht. Working with them was a great pleasure and source of motivation. The cooperation resulted in two articles, one of which is part of this dissertation. Also, I would like to thank Anton Vasilev for invaluable help and support in collecting the data for one of these projects.

During my time in Munich I had the opportunity to work in a stimulating environment. In particular, I would like to thank Lukas Buchheim, Georg Gebhardt, Martin Kocher, Thomas Kolaska, Henrike Lindemann, Sandra Ludwig, Dominik Ruderer and Anno Stolper for their help and many insightful discussions. I am especially grateful to Benno Bühler for numerous inspiring discussions and helpful comments. Also, I would like to thank Werner Barthel, Michael Seitz, Sebastian Strasser and Stefan Vetter for the various ways in which they helped me during my dissertation. Financial support through the DFG via GRK 801 and SFB/TR 15 is gratefully acknowledged.

Finally, I am most grateful to my parents, my sister and my friends for their constant support and encouragement during the time of my dissertation.

Contents

List of Figures

Preface

This doctoral dissertation deals with different issues in the realm of financial economics and corporate finance. In particular, it considers a bank's optimal level of opacity and issues connected to incentive pay for managers. The first chapter offers a potential explanation on why banks regularly choose to issue highly opaque financial claims. Chapter 2 considers bonus payments in the financial industry and their effects on the propensity of managers to engage in undesirable actions. Chapter 3 analyses the optimal tenure of managers and investigates how this depends on a manager's past performance and endogenous changes in his wealth.

Since the beginning of the subprime mortgage crisis in 2007, the incentives of various actors on financial markets have attracted a lot of scrutiny. There are two different levels at which the incentives of key players in the financial industry can be examined. At the first level, we may ask ourselves if important institutions such as investment banks have the right incentives that are needed to guarantee the stability and efficiency of the financial sector. But even if the shareholders of banks and investment funds are facing proper incentives, it is unclear if the same holds true for their employees. Hence, at a second level, we need to consider the provision of incentives within organisations. In particular, do organisations write contracts that make sure that their employees act in their best interest? The first chapter is concerned with the first type of question, while Chapters 2 and 3 focus on incentives within organisations.

The fact that the recent global economic crisis originated in the financial industry helps to explain why incentives in this particular sector have attracted a large amount of public attention. In April 2010, the International Monetary Fund estimated that in the banking sector alone, the financial crisis had led to write downs of US$2.3 trillion. As of January 2011, this amounted to approximately 16 times the market capitalisation of Citigroup or to 2.4% of the value of all outstanding bonds worldwide according to calculations by the trade group TheCityUK. Obviously, we should ask ourselves how it was possible for an industry to accumulate this amount of losses in less than three years. But even if we abstract from recent events, it seems particularly important to understand the mechanisms that shape individual decisions in the financial sector. According to the US Department of Commerce in 2011 the financial sector contributed 8.3% to the gross domestic product of the United States. While this figure is somewhat lower for most of the large European economies, it is even larger for the United Kingdom. So the sheer size of this specific sector makes it worthwhile to try to better understand the underlying incentives. Moreover, recent events have highlighted that developments in the financial sector are prone to affect other sectors of the economy, too. A particular feature of banks and other financial intermediaries is that they act as input suppliers to almost any other industry. Companies from all sectors rely on the services of banks and investment banks in order to obtain credit, issue bonds or acquire other firms. Hence, flaws in the regulatory framework for banks do not only matter because we are looking at an industry that is in itself very important, but also because they can directly affect other industries by hampering their access to efficient financial services.

The common theme of this dissertation is that it tries to find rational explanations for the behaviour of firms in the financial industry. Arguably, firms do not always behave optimally in all aspects of their business. At the same time, the sector we consider is very competitive and shaped by a number of highly sophisticated institutions. Consequently, it would be surprising to see firms in this industry collectively adopt policies that are not in their own best interest. Yet, such claim have regularly been made with respect to the widespread opacity of financial assets or high-powered incentives for senior employees in the banking sector. In this dissertation we propose explanations as to why such seemingly irrational choices may make sense from an ex-ante perspective.

Preface

On the 9th of August 2007 BNP Paribas announced in a press release that three of its investment funds were no longer able to determine the value of assets that were backed by subprime mortgages originated in the United States due to their high level of complexity. One year later, banks around the world were struggling to absorb large losses incurred partly because of the trade of such mortgage-backed securities. While the potential repercussions of opacity in financial markets have since become very clear, it is much less obvious why banks found it optimal to design largely intransparent financial claims in the first place.

The first chapter analyses the incentives of a bank to structure its investments in ways that make it difficult for outsiders to obtain precise information on the quality of an investment project. Banks that have a need for liquidity can choose to obtain funds from other banks via ex-ante contracts that allow a lender to acquire non-verifiable information on the project. The quality of this information depends on the borrower's choice of transparency. Alternatively, a borrower can obtain funding on a spot market for interbank funds, where lenders are completely uninformed.

An optimal ex-ante contract grants credit to a borrower if and only if the lender announces that he has obtained favourable information on the project. This ensures that the lender has an incentive to reveal his information on the project's quality truthfully and can not gain by always claiming to face a borrower that is of low quality. Borrowers that have received a negative evaluation turn to the spot market in order to obtain funding. However, if informed lenders are highly effective at evaluating projects, uninformed lenders rationally expect the quality of debt traded on the spot market to be poor and the spot market may break down due to asymmetric information. Ex-ante, borrowers hence may want to choose financial structures that make it hard for informed lenders to obtain a lot of meaningful information. This increases the average quality of claims traded on the spot market and ensures that the market does not freeze.

Even though transparency may lead to a breakdown of the market for interbank loans, it also allows for a more efficient allocation of funds to the best projects. This implies that if the integrity of the spot market were guaranteed for, it would always be optimal to have as much transparency as possible. It hence may be optimal for the government to commit to interventions that keep the market for interbank loans liquid. This encourages borrowers to disclose more information and can increase welfare.

Chapter 1 contributes to the literature on optimal opacity by showing that financial intransparency can be optimal even if it is impossible to keep the person that originates a claim uninformed. While Tri Vi Dang, Gary Gorton and Bengt Holmström (2009) show that opacity can be optimal in a variety of scenarios, this effect is driven by the fact that it is possible to keep everybody in an economy uninformed. Consequently, opacity fully eliminates asymmetric information in their model. In our set-up, opacity affects the market composition of the spot market and may be desirable even if it increases the informational advantage of the borrower vis-à-vis lenders. The chapter is also related to the literature on government interventions in markets with asymmetric information by Thomas Philippon and Vasiliki Skreta (2012) and Jean Tirole (2012). However, we focus on the ex-ante effects of such interventions and show that they may induce borrowers to disclose more information. This is socially desirable as long as the government vouches for the integrity of the spot market.

In the second chapter we turn to the question of how incentives within organisations may have affected behaviour in the financial and other industries. Again, we consider an issue that has received a lot of attention in the wake of the subprime mortgage crisis, namely the compensation of CEOs and other senior executives. At the height of the financial crisis, British Prime Minister Gordon Brown described the size and structure of bonuses in the financial sector as "irresponsible" and the French minister of Economic Affairs Christine Lagarde judged incentive schemes for bankers as "perverse".[1] According to these arguments, high-powered incentives induced managers to take excessive risks and ultimately contributed to the crisis of the banking sector.

[1] Michael Peel and George Parker, "Brown attacks 'irresponsible' City bonuses," Financial Times, September 21, 2008, accessed May 31, 2012, http://www.ft.com/intl/cms/s/0/d35908c6-8810-11dd-b114-0000779fd18c.html

Preface

We show that large bonus payments may in fact have been a way to reduce gambling and reckless behaviour on the side of employees. In order to do so, we consider a moral hazard model where the agent can influence the probability that the firm reports high profits via two different actions: He can exert effort and he can engage in undesirable behaviour. While undesirable behaviour increases the observable revenues of a firm it has other, potentially very harmful, consequences such that a principal would always prefer his agent not to carry out these kinds of actions. Usually, the damages caused by misbehaviour are impossible to verify in court and it is hence impossible to condition the agent's compensation on these damages. However, with a small probability the principal obtains additional, hard evidence on misconduct and is able to punish the manager.

Whenever the principal wants the agent to exert effort, he must accept that the manager will also engage in some positive level of misbehaviour since both actions influence the performance figure that the agent's compensation is based on. Indeed, for small bonuses the level of effort and undesirable conduct are both monotonically increasing in the size of the bonus. Yet, this relationship breaks down once bonuses have reached a sufficiently high level. While the amount of effort that an agent exerts is still increasing in the size of the bonus, an increase in the bonus will reduce the amount of energy an agent invests into misconduct. Besides increasing the probability that the firm reports high profits, misbehaviour also increases the probability that the principal detects such conduct and does not pay any bonus at all. This threat is particularly effective if bonuses are substantial. The agent is therefore more likely to comply the larger the bonus he expects to earn.

The main prediction of the model is that the "gaming" of incentive schemes is not monotonically increasing in the power of incentives. Rather, it is for intermediate incentives that agents are most likely to misbehave. The central policy implication of the model is, hence, that legal caps on bonuses may be counterproductive: By limiting the admissible size of bonus payments, a policy maker may increase the level of misbehaviour.

The second chapter is related to the literature on multi-tasking in principal-agent relationships that was initiated by Bengt Holmström and Paul Milgrom (1991). In this literature it is argued that multiple actions should reduce the incentives offered in one dimension since the agent would otherwise neglect his other duties in order to focus on the well-incentivized task. We show that the opposite can hold true if the agent is protected by limited liability: Increasing incentives for one action increases the punishment that a principal can impose in the second dimension. Hence, the addition of a second action may increase optimal incentives in the first dimension. A related model of misbehaviour in the presence of limited liability constraints in considered by Roman Inderst and Marco Ottaviani (2009). Yet, in their model an agent chooses to misbehave once he has learned that he will not obtain a bonus by legitimate means. So unless he misbehaves he does not obtain a bonus for sure and the disciplining effect of bonuses that we consider in our model never plays a role. Finally, the predictions of our model are in line with empirical evidence by Rüdiger Fahlenbrach and René M. Stulz (2011) who find that in the banking sector, larger cash bonuses did not lead to more risk-taking in the years leading up to the financial crisis.

Chapter 3 of this dissertation is joint work with Piers Trepper and considers the question of why unsuccessful managers are not markedly more likely to lose their job than successful ones. There is a large body of empirical literature that finds the correlation between CEO success and forced turnover of CEOs to be very small. According to a broad number of studies of large, publicly held companies in the US, CEOs from the 10th percentile of firm performance are only 2-6% more likely to lose their job than managers from the 90th percentile. This low correlation seems surprising since we would typically assume that good performance is a signal of high managerial skill. It should, hence, be optimal to retain successful managers and to terminate the contracts of less successful CEOs.

We show that in the presence of performance pay, it can instead be optimal to let go of successful managers and to retain unsuccessful ones. In order to ensure that a manager has an incentive to exert effort, he has to be rewarded with a bonus in case he is successful. However, this implies that successful managers become wealthier than unsuccessful ones and it may be harder to motivate the former in subsequent periods. While unsuccessful managers are of poor expected ability, they are "hungry" and easy to motivate via monetary incentives. Therefore it can be optimal to tolerate failure and to renew the contracts of unsuccessful managers.

A change in the agent's wealth affects the cost of supplying incentives in two ways. On the one hand, a richer agent may become less risk-averse and may be more likely to accept a bonus contract with uncertain future income. On the other hand, he has a lower marginal utility of income and it is harder to compensate him for his cost of effort. We consider a situation where the second effect dominates and show that this is the case under mild assumptions on the agent's utility function.

Whenever it is extremely important that a project turns out to be successful, it is optimal to reemploy a manager only if he has a positive track record. If a success is crucial, the principal will induce any manager to exert maximal effort. Moreover, the differences in the cost of compensation are small relative to the benefit of employing an agent with a higher expected ability that is more likely to deliver positive results. If a successful conclusion of the project is not important at all, it is again optimal to reemploy only successful managers. In this case the principal offers contracts that implement very limited effort and the cost of compensation is low irrespective of the agent's employment history. Again, the principal only cares about talent and it is optimal to always hire the agent with the highest expected ability. However, for intermediate values of success, the cost of compensation is an important determinant of firm profits and it can be optimal to keep unsuccessful managers while not renewing the contracts of successful CEOs.

The third chapter builds on results by Henrik Thiele and Achim Wambach (1999) who show general conditions under which a principal prefers a less wealthy agent to a richer one. However, they neither consider endogenous changes in wealth, nor the trade-off between wealth and ability. Starting with Robert Gibbons and Kevin J. Murphy (1992), there is an extensive literature on dynamic agency problems in which the principal revises his belief about the agent's ability. Yet, these models typically assume that changes in the agent's wealth affect not only the marginal utility of income, but also the marginal disutility of work. If the agent has constant absolute risk aversion these two effects cancel each other out and the cost of incentives does not depend on wealth. Chapter 3 combines these two strands of literature by considering a situation where both, the agent's expected ability and the cost of incentives change as a function of the agent's employment history.

All three chapters of this dissertation are self-contained and include their own introductions and appendices such that they can be read independently. The respective appendices contain all proofs that are not included in the text.

Chapter 1

Optimal Opacity and Market Support

The recent financial crisis has often been attributed to wide-spread opacity in the design of financial claims and on the balance sheets of banks. From a theoretical point of view the amount of intransparency was indeed staggering: A more transparent financial system should not only allow for funds to be directed to the most efficient projects, but it should also reduce the informational asymmetry between informed sellers and potentially uninformed buyers of financial claims. So it seems surprising that banks voluntarily issued complex asset backed securities and other financial products that are considered to be highly opaque. In this chapter we show that it can indeed be optimal for banks to choose to be intransparent, even if intransparency increases informational asymmetries.

We consider an interbank market on which one bank is endowed with a profitable investment opportunity and needs to obtain funds from banks that have excess liquidity. It can either do so by turning to a relationship lender who is able to obtain information on the quality of the project. Or it can choose to borrow from perfectly uninformed transactional lenders. Less-than-perfectly transparent financial claims make it harder for relationship lenders to evaluate the quality of a project. This reduces the stigma attached to a borrower who does not obtain funding from a relationship lender and has to turn to a spot market for interbank credit. Transactional lenders on the spot market rationally expect that their informational disadvantage vis-à-vis relationship lenders is small if opacity is high. This in turn reduces the interest rates demanded by transactional lenders and may prevent a breakdown of the spot market for interbank funds. While borrowers that receive credit from their relationship lender have to pay a higher interest rate if opacity is high, intransparency may still be optimal in order to ensure the liquidity of the spot market. This chapter shows that even though opacity unambiguously increases the informational asymmetry between contracting parties, it can nevertheless play a role in guaranteeing market liquidity and may be both, privately and socially desirable.

We show that even though borrowers can always avoid a spot market breakdown ex-ante by designing sufficiently opaque financial claims, a government may find it optimal to commit to intervene by offering schemes in the fashion of the original plans for the Troubled Asset Relief Program ("TARP 1") in case the market does threaten to freeze. As part of this scheme, the government offers to buy claims on the spot market at a loss. By doing so it subsidises trade on the spot market which ensures that transactional lenders find it optimal to trade claims and that the market does not break down. From an ex-ante perspective such a policy may be desirable since it allows borrowers to adopt more transparent financial structures without provoking a market freeze. While opacity may be necessary in order to prevent a market breakdown, it is costly since it reduces the financial sector's ability to allocate more funds towards the more efficient projects. A government intervention allows the financial sector to enjoy the benefits of transparency without jeopardising the liquidity of the spot market. Our model suggests that concerns about the ex-ante moral hazard created by schemes like TARP may be exaggerated. Indeed, we will see that the anticipation of a government intervention can even have positive welfare effects.

Starting with George A. Akerlof (1970) a large body of literature has discussed the costs of asymmetric information and has typically argued that from an ex-ante perspective, a seller would be best off by disclosing as much information as possible in order to guarantee symmetric information. The present model shows that this needs not be the case if a seller of financial claims is concerned about his access to alternative sources of funding. While this insight is surprising from a theoretical point of view, it is in line with a large number of stylised facts that suggest that actors on financial markets choose to obfuscate information on a regular basis. In particular, the rise in the prevalence of complex Asset Backed Securities (ABS) and other classes of highly structured financial products can be seen as a sign of such voluntary opacity. While financing investments via Asset Backed Securities is only one of many ways in which banks may increase financial intransparency, it is a very important one: According to the Securities Industry and Financial Markets Association, the value of outstanding ABS in the US reached almost US$3 trillion at the beginning of the subprime mortgage crisis in 2007. The fact that banks heavily relied on securitizing even small mortgages meant that a large number of claims had to be bundled in order to create a tradable ABS. Yet, it is often argued that many of these assets were bundled in ways that were more intransparent than necessary. In its proposal to revise the disclosure regulations for Asset Backed Securities the US Securities and Exchange Commission argues that "*...in many cases, investors did not have the information necessary to understand and properly analyse structured products...*" . This suggests that these financial products could indeed have been made more transparent and that the opacity of these assets was a choice rather than a technological necessity.

We consider a model in which one bank is endowed with an investment opportunity that requires additional outside funding and may be more or less profitable. The bank finds it optimal to invest its own capital into the project regardless of the project's type. On top of that, it can obtain funds from a relationship lender or from transactional lenders in order to increase the scale of the investment. It is only when the borrower is turned down by a relationship lender that he turns to the spot market for transactional funding. However, at this point in time the borrower has obtained information on the quality of his project and the market is subject to asymmetric information.

A relationship lender is able to obtain information about the project in order to assess, albeit imperfectly, the quality of the investment opportunity. In order to make it incentive compatible for the lender to reveal the outcome of the project evaluation truthfully, it is optimal to agree on a contract that grants credit only to those projects that received a positive evaluation and to terminate relationships with borrowers who have projects that are likely to be of bad quality. Even though the borrower is free to take out a loan on a competitive spot market made up of uninformed transactional lenders, the fact that he has to resort to an outside market may give rise to an inefficiency: Whenever the average quality of debtors on the spot market falls below some threshold, this will give rise to a “market for lemons” and a borrower with a good project that has received a wrongful evaluation will not find it worthwhile to obtain credit on the spot market. This is not welfare-maximising since it sharply reduces the scale of the most profitable projects.

A market breakdown becomes more likely the more information a relationship lender obtains: Better screening reduces the number of good projects that end up in need of spot market funding. Hence, lenders on the spot market have to demand a higher interest rate in order to break even and borrowers with highly profitable projects who end up on the spot market may no longer find it optimal to ask for credit. In order to prevent such a market breakdown, a borrowing bank can create purposefully opaque assets that make it harder for relationship lenders to assess the quality of an investment project. Typically, the fact that financial claims are less-than-perfectly transparent is observable even to otherwise uninformed parties. Hence, designing intransparent assets convinces transactional lenders that they are not at a large informational disadvantage vis-à-vis a relationship lender. This in turn reduces the interest rate demanded on the spot market for interbank trade and ensures the liquidity of this market. However, intransparency comes at a cost since a borrower that received a positive evaluation has to pay a higher interest rate if opacity is high. Moreover, intransparency is costly even from an ex-ante perspective. In case the spot market does not break down, opacity increases the interest rate payed by borrowers with a positive screening and it reduces the rate paid by borrowers with a negative evaluation. This reduces the scale at which projects of a high expected quality can be set up and increases the scale at which the less profitable projects are financed. Since borrowers pay the price for this reduction in efficiency, they only choose positive levels of opacity if this is necessary to prevent a market breakdown.

The fact that loosing funding from a well-informed lender carries a stigma and can lead to a market breakdown may explain why borrowers are not interested in close ties to lenders. In this respect our model resembles the hold-up problem discussed in the relationship banking literature (see, e.g., Raghuram G. Rajan, 1992; Steven A. Sharpe, 1990; Ernst-Ludwig von Thadden, 1995). In these models lenders may use their information monopoly to extract rents from entrepreneurs, a problem that can be addressed by multiple banking relationships (Steven Ongena and David C. Smith, 2000).[1] The effect we consider in this model is considerably different. In particular, our model implies that a government can encourage transparency by preventing a breakdown of the spot market via schemes like TARP. The chapter closely relates to the literature on market liquidity: Dang, Gorton and Holmström (2009) present a model in which information acquisition is harmful since it generates asymmetric information. However, their model relies on the assumption that it is possible to keep *everybody* in an economy uninformed by choosing appropriate contracts. In contrast to this, in our setting the borrower will always learn his type, so opacity increases the informational asymmetries between contracting parties. Even though the mechanism is quite different, our results are complementary to Dang, Gorton and Holmström (2009) in that they also emphasise the role of opacity in ensuring market liquidity. In a related paper Marco Pagano and Paolo Volpin (2008) show that opacity can be desirable since it increases the time sophisticated investors need in order to learn the quality of an asset, which can reduce informational asymmetries. Our model of government intervention in markets with asymmetric information builds on a recent body of literature initiated by Philippon and Skreta (2012) and Tirole (2012), who analyse optimal mechanisms by which a government can re-start a market hampered by asymmetric information at minimal fiscal cost. Finally, Frédéric Malherbe (2010) shows that expectations of a market breakdown may be self-fulfilling: If a borrower expects the resale market for his assets to break down, he finds it optimal to hoard liquidity in order to pay for potential reinvestments in the future. But if any borrower is able to finance reinvestments without selling his assets, all trade on the resale market will be opportunistic and a market break-

[1] A related rationale for multiple banking relationships is put forward by Enrica Detragiache, Paolo Garella and Luigi Guiso (2000). In their model single banks are fragile and may fail to supply credit.

down is indeed an equilibrium. Conversely, sellers who have an urgent need for liquidity are likely to sell assets of high quality since they do not trade for opportunistic reasons. In our setting the opposite holds true: A need for liquidity implies that the project has failed to secure other sources of funding and is likely to be of *bad* quality.

Note that the optimal contracts signed with relationship lenders resemble real-world credit lines. An optimal contract specifies conditions at which an eligible borrower can take out a loan in the future, but a lender reserves the right to withhold funding altogether. This is in accordance with the empirical observation by Amir Sufi (2009) that credit lines are usually tied to strict covenants and are frequently withdrawn if a debtor's economic outlook worsens. However, this interpretation of credit lines differs greatly from the traditional views of credit lines as insurance against high liquidity needs (Tim S. Campbell, 1978; Bengt Holmström and Jean Tirole, 1998), as a way to give investment incentives (Arnoud W.A. Boot, Anjan V. Thakor and Gregory F. Udell, 1987) or to reduce transaction costs (J. Spencer Martin and Anthony M. Santomero, 1997). The rest of the chapter is organised as follows: Section 1.1 formulates the baseline model and Section 1.2 derives the optimal financial contracts. Section 1.3 analyses the benefits of opacity in this framework. Section 1.4 explores the policy implications of the model and consider how government interventions on the spot market can increase efficiency. Section 1.5 concludes.

1.1 The Model

We consider a setting with two types of banks: One bank, the borrower B, is endowed with some limited amount of capital A and an investment project that can be expanded by an arbitrary scale $S \in \mathbb{R}^+$. We can think of this investment opportunity as a large-scale project that requires the bank to obtain additional funds on an interbank market. While we typically think of banks as lenders, at any given point in time a bank may have more investment opportunities than it has liquidity available and can approach other banks in order to obtain additional funds. Each unit of investment has a cost of \$1 and will generate future, non-pledgable income of ρ_{np}. This income can not be promised to debt holders and includes private benefits or information that may prove useful in future transactions. Additionally, with probability θ_i the investment will generate financial returns ρ_p that can be pledged to outside investors. The probability that a project is commercially successful depends on the type of the project: $i \in \{H, L\}$ where $\theta_H > \theta_L$. We denote the probability of a project being of type $i = H$ by α. All other banks, called "lenders", lack an investment opportunity and are prepared to lend their capital. For simplicity, we assume that lenders do not face any capital constraints.

The timing of the game can be described by three periods: At $t = 0$ B does not yet posses any information on the type of his project and can secure funding for the upcoming investment opportunity. At $t = 1$ he learns the type of his investment project, can obtain additional funding on a spot market and make investments. At $t = 2$ uncertainty is resolved, projects are terminated and creditors are repaid.

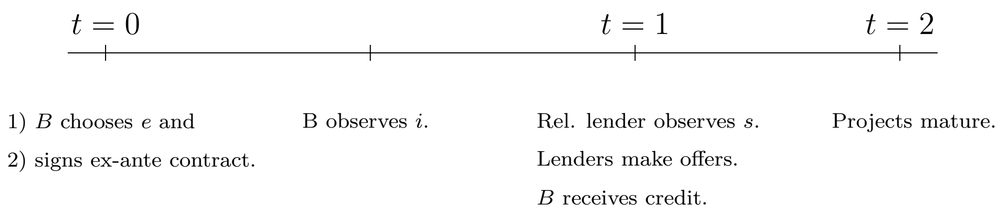

Figure 1.1: Timing of the Game.

We assume that there are two types of financial contracts that the borrower can sign: First, he can agree to a contract that makes payoff-relevant information available to a lender. An informed lender can evaluate the quality of the investment opportunity and receives a non-verifiable, imperfect signal $s \in \{H, L\}$ on the project's type. This allows him to offer B different terms of financing depending on the outcome of the evaluation process. Contracts are signed behind the veil of ignorance (i.e. at $t = 0$), so asymmetric information is not an issue. Since this type of contract is typically based on debtor-specific information, we will refer to the respective lender as a "relationship lender" (see Arnoud W.A. Boot, 2000). More generally, we will call this kind of contractual arrangement an "ex-ante contract" to reflect the fact that it is signed before the borrower learns anything about the type of his project. Additionally, the borrower can obtain funding from perfectly uninformed transactional lenders. In Section 1.2.2 we will see that without loss of generality, we can assume that the corresponding contracts are signed in period 1 when the borrower has already learned the type of his project.

Let us consider period 0 in more detail: At this point B is endowed with an investment project that may be expanded in the future, but he is not yet able to assess the type of the project. He can sign an ex-ante contract that specifies an amount of credit $q^{\hat{s}}$ that the borrower gets in period 1 and a corresponding repayment $d^{\hat{s}}$ that he has to make in period 2. Both of these quantities can be made contingent on a report $\hat{s} \in \{H, L\}$ that the lender makes about s, the outcome of the screening process. We assume that for any report $\hat{s}$ the lender offers a unique $q^{\hat{s}}$ and $d^{\hat{s}}$. In principle, a lender could offer B a choice between multiple financing terms for each $\hat{s}$ in order to induce borrowers of different types i to select into different options in period 1, once they have learned their own type. However, in Appendix A2 we show that it is indeed an equilibrium for all borrowers that receive a given report $\hat{s}$ to accept the same terms. Hence, while a contract may depend on the information that a lender receives, the borrower's information will never influence the terms of trade with a relationship lender. It may therefore be necessary to ensure that the lender has an incentive to report his signal s truthfully.

Whenever the borrower is unable to meet his obligations in period 2, he defaults on his debt but is still able to enjoy the non-pledgable income that accrued from the project. Throughout this chapter, we assume that contracts are unobservable to other lenders.

Before signing any contract, the borrower can influence the precision of the signal s that a lender will receive by choosing the level of financial transparency e of his project. A borrower may decide to structure the investment project in a less transparent fashion, e.g. by bundling various distinct assets into a portfolio of Asset Backed Securities, by securitizing fractions of future claims on outside markets, or by moving revenues into off-balance sheet vehicles. All of these decisions will make it harder for any outside party to assess the type of a given investment project.

If the borrower makes more information available, this enables the lender to evaluate the the quality of the project with higher precision. For simplicity we assume that for any e we have $Prob(s = H \cap i = L) = Prob(s = L \cap i = H)$, i.e. the lender is equally likely to err in either direction when assessing the type of a project. The probability of either type of misjudgment is given by $\alpha(1-\alpha)(1-e)$. For $e = 0$ the borrower does not disclose any information and the signal is completely uninformative, while for $e = 1$ the probability of misjudgment becomes zero. We assume that the borrower can choose e at zero cost from an interval $(0, \overline{E}]$ where $\overline{E} < 1$. Hence, even if B goes out of his way to offer maximal transparency, the relationship lender will be unable to asses the quality of the project perfectly. This implies that it is never possible to fully eliminate asymmetric information.

We subscribe to the notion that it is impossible to hide information from the borrower himself: While the relationship lender will only receive a noisy signal s, the borrower will eventually learn his type perfectly. The borrower knows how his structured claims have been set up and what will end up on (or off) his balance sheets. Hence, he is able to deduce what any new information implies for the prospects of his investment opportunity irrespective of the level of transparency. A lender on the other hand does not possess all of this information and will find it difficult to asses the potential of a project when opacity is high.

In period 1 the borrower can obtain additional funding on a spot market for credit should he not be able to satisfy his liquidity needs via ex-ante contracts. On this market a large number of uninformed lenders simultaneously offer to lend money at an interest rate of r_M that allows lenders to break even.[2] Since a borrower can not pledge any future earnings to more than one lender, he can raise at most q_M units of capital on this market where $\left(A + q^{\hat{s}} + q_M\right)\rho_p = d^{\hat{s}} + q_M r_M$.[3] A fundamental difference between relationship lending

[2]Note that Andrea Attar, Thomas Mariotti and François Salanié (2011) show competition in linear prices to be indeed an equilibrium under asymmetric information if borrowers can not commit to deal with only one lender. While there are other equilibria, the equilibrium allocation is always unique.

[3]Since contracts are unobservable, the assumption that a borrower can be kept from pledging future income to more than one borrower may not seem very natural. Yet, we can interpret this assumption as follows: When dealing with any given lender, the borrower sets up a new legal entity and transfers some

and transactional financing is the ad-hoc nature of the latter: Ex-ante contracts are signed when the borrower does not yet have a clear grasp of the quality of the project, while the borrower already knows the type of his project perfectly once he tries to obtain funding on the spot market. At this point in time the project is fully developed and the borrower has a precise idea of the risks and opportunities associated with it. This implies that a borrower may decide to drop out of the spot market in case he considers the interest rate to be unduly high relative to the profits he could make if he doesn't take out any loan and sets up the project using only his own capital A. At the end of period 1 the borrowing bank makes its investments and expands the project by the desired scale S.

In the last period, uncertainty is resolved and creditors are repaid whenever possible. For simplicity, we normalise the discount factor in the economy to one. Moreover, for the rest of the chapter, we impose the following parameter restrictions:

$$\rho_{np} + \theta_i \rho_p > 1 \qquad \forall i \in \{H, L\} \tag{1.1}$$

$$\theta_i \rho_p < 1 \qquad \forall i \in \{H, L\} \tag{1.2}$$

These restrictions imply that while investing is always socially efficient (1.1), it is never possible to increase the scale of a project using credit alone (1.2). Hence, the size of any project is restricted by the amount of capital A that a debtor can invest into it.

The key source of inefficiency in the model is the fact that relationship lenders only receive soft information on the quality of a project. If the lender could commit to reveal his signal truthfully, a borrower would make available as much information as possible. He would sign a contract that supplies him with a large, subsidised loan in case his project is likely to be of high quality and that expropriates him in case his investment opportunity is likely to be bad. Moreover, he would never need to turn to the spot market. The fact that the relationship lender has to find it individually optimal to reveal

of his own capital to this entity. The lender in turn grants credit to this entity and obtains senior claims on all future cash-flow.

his signal truthfully introduces two restrictions: First, a lender can not offer a contract that subsidises borrowers with a favourable signal at the expense of borrowers that are likely to be of low quality. Otherwise, he would always have an incentive to misreport his signal as being $s = L$.

Even worse, in order to give a lender incentives to report his signal truthfully, an optimal ex-ante contract does not offer any credit to borrowers in case of a negative signal and these borrowers have to turn to the spot market instead. In this respect the model is reminiscent of the literature on “up-or-out” clauses in labour contracts (e.g. Charles Kahn and Gur Huberman, 1988).[4] In order to make the revelation of a lender's non-verifiable signal incentive compatible, it is optimal to terminate contracts of borrowers that are likely to be of bad quality. Hence, whenever the lender is to reveal his signal, it is impossible to circumvent the spot market that is subject to asymmetric information. While it is possible to design ex-ante contracts that supply the borrower with funding irrespective of $\hat{s}$, such contracts can not make use of any of the lender's information and turn out never to be optimal.

1.2 Equilibrium Contracts

1.2.1 The Spot Market

We start our analysis by looking at the period-1 spot market for a borrower who wants to take out a loan from uninformed lenders. In order to do so we assume that only borrowers who have received a negative evaluation consider borrowing on the period-1 market and that these borrowers trade only on the spot market. In the next section we will see that this assumption is without any loss of generality.

The spot market is characterised by asymmetric information since B has already learned his type while the banks on the spot market, which we refer to as transactional lenders, are completely uninformed. Given that only borrowers with $s = L$ end up on this market,

[4] Very similar clauses in financial contracts are also considered by von Thadden (1995).

B's type can be fully characterised by his true quality $i \in \{L, H\}$. While i is unobservable to the prospective lenders, the borrower may make the decision whether to accept credit at a given interest rate contingent on his knowledge about his investment opportunity's quality.

The linear nature of the problem implies that a borrower who does not receive funding from his relationship lender will either not take out any loan at all or he will borrow as much as possible on the spot market. If he decides not to pledge any future cash-flows to another bank, the size of his investment will be equal to his endowment A and he receives an expected utility of

$$A\left(\rho_{np} + \theta_i \rho_p\right). \tag{1.3}$$

This utility under self-financing is larger for borrowers that have a high-quality project.

If, on the other hand, the borrower decides to take out a loan, this allows him to invest at a larger scale while promising the pledgable income to lenders. Since lenders make zero profits, they offer to buy claims from the borrower at an interest rate of $r_M = 1/\hat{\theta}_M$ where $\hat{\theta}_M$ is the market's rational expectation of θ given that B has decided to trade on the spot market. The total amount of credit that a borrower can obtain at this interest rate is given by q_M which is implicitly defined by the feasibility constraint $(A + q_M)\,\rho_p = q_M r_M$. This implies that the borrower can invest at scale $S = A + q_M = \frac{A}{1-\rho_p\hat{\theta}_M}$ and receives an expected utility of

$$\left(\frac{A}{1-\rho_p\hat{\theta}_M}\right)\rho_{np}. \tag{1.4}$$

Assume that both types of borrowers choose to take out a loan at $t = 1$. In this case the share of claims traded on the spot market that are of high quality is given by

$$\alpha_M = \frac{Prob(s = L \cap i = H)}{Prob(s = L)} = \frac{\alpha(1-\alpha)(1-e)}{\alpha(1-\alpha)(1-e) + [(1-\alpha) - \alpha(1-\alpha)(1-e)]} = \alpha(1-e)$$

and the market's rational expectation concerning a project's quality is given by $\hat{\theta}_M = \alpha_M\theta_H + (1-\alpha_M)\theta_L$. Indeed, whenever $\hat{\theta}_M$ is sufficiently large, (1.4) is larger than (1.3) irrespective of i and everybody participates in the $t = 1$ market.

However, this does not hold true once the expected quality of borrowers who are in need of funding at $t = 1$ drops below some threshold $\tilde{\theta}_M$. In this case lenders can no longer offer an interest rate that allows them to break even and that still induces high types to borrow. Instead, any borrower with a good project chooses self-financing and we end up with a market for lemons where only low-quality debtors take out loans and pay an interest rate of $r_M = 1/\theta_L$.

The expected quality of borrowers who consider taking up credit on the spot market is determined by two factors: The overall share of high quality projects α and the quality of information available to relationship lenders. A higher share of good projects increases the expected quality of claims traded on the spot market, while the quality of information obtained by relationship lenders negatively affects the quality of claims traded in $t = 1$. If relationship lenders are more effective at evaluating projects, fewer good projects will be misjudged and will end up on the spot market. Yet, the overall volume of credit traded on the period-1 market is unaffected by the level of transparency. Consequently, the expected quality of debtors on the spot market is decreasing in the level of transparency and a market breakdown becomes more likely the more information a borrower discloses to his relationship lender. Indeed, high types will always drop out of the market when screening is sufficiently precise and the following assumption is satisfied:

Assumption 1. *The asymmetric information problem is "severe" relative to the surplus that can be generated by investing:*

$$\rho_{np} + \theta_H \rho_p < \frac{\theta_H}{\theta_L}.$$

Assumption 1 implies that high types will not find it attractive to take out a loan if they are taken to be of bad quality for sure (and have to pay the corresponding interest rate that would allow banks to break even). Conversely, when the asymmetric information problem is rather mild, the efficiency gains of setting up the project at a larger scale are sufficiently large to justify paying high interest rates and adverse selection will never turn out to be a problem on the period-1 market. Also, we will make the following assumption:

Assumption 2. *The share of good projects α is sufficiently high:*

$$\rho_{np} + \theta_H \rho_p > \frac{\theta_H}{\alpha\theta_H + (1-\alpha)\theta_L}.$$

This condition guarantees that if all lenders are perfectly uninformed about the project's quality, all claims will be traded in equilibrium. We will impose Assumptions 1 and 2 throughout the rest of the chapter.

Under these assumptions we can show that the spot market breaks down if and only if the level of transparency e exceeds some threshold:

Lemma 1. *There exists some level of transparency $\tilde{e} \in (0,1)$ such that high types who end up on the spot market take out a loan if and only if $e \leq \tilde{e}$ where*

$$\tilde{e} = 1 - \frac{1}{\alpha(\theta_H - \theta_L)}\left[\frac{\theta_H}{\rho_p\theta_H + \rho_{np}} - \theta_L\right].$$

Proof. See the appendix.

If relationship lenders are very effective at screening projects, transactional lenders can be almost certain that any borrower who ends up on the spot market is of low quality. Consequently, they demand a high interest rate and the spot market breaks down. However, limited information disclosure can ensure that the spot market works smoothly. Opacity increases the portion of high types seeking to take out a loan on the spot market. This guarantees favourable interest rates and makes sure that in equilibrium high types do indeed trade on this market. Using Lemma 1 we can define $\tilde{\theta}_M = \hat{\theta}_M(\tilde{e})$ as the lowest expected quality of claims on the spot market that does not lead to a market breakdown.

We will see that in case the spot market does not break down, the borrower would always prefer to be as transparent as possible. Nevertheless, it may pay for a borrower to choose intransparent financial structures ex-ante in order to prevent a market freeze.

1.2.2 Optimal Ex-ante Contracts

Having lined out how the spot market for credit in period 1 works, we can now turn to the contracts that relationship lenders offer in period 0. For the time being, we will assume that e is exogenously given in order to fix ideas. To avoid confusion, we will call borrowers on whom a lender has received a positive signal ($s = H$) "approved" borrowers and those on whom a lender has received a negative signal ($s = L$) "disapproved" borrowers.

Since contracts are signed behind the veil of ignorance, asymmetric information is not an issue. In fact, we would expect borrowers to commit to trade with the relationship bank irrespective of the report that the lender makes about the outcome of the screening process and to circumvent the period-1 market altogether in order to avoid the cost of asymmetric information. However, we will see that the set of contracts a relationship lender can offer is restricted by the fact that the lender gains only soft information on B's type. This will explain why it is usually optimal not to grant any credit to disapproved borrowers.[5]

Recall that an ex-ante contract specifies an amount of credit $q^{\hat{s}}$ that a borrower receives in period 1 and a repayment $d^{\hat{s}}$ that he has to make in period 2, given that the relationship lender reports $\hat{s} \in \{L, H\}$ as the outcome of the project evaluation. Without loss of generality, we assume that the lender reports his signal truthfully in equilibrium, i.e. $s = \hat{s}$. Since the market for ex-ante contracts is competitive we must have $\alpha(d^H\hat{\theta}^H - q^H) + (1-\alpha)(d^L\hat{\theta}^L - q^L) = 0$ where $\hat{\theta}^s$ denotes the expected quality of a borrower conditional on receiving signal s and where $\hat{\theta}^H > \hat{\theta}^L$ for all $e > 0$. At the same time, lenders must have an incentive to reveal their signal truthfully, which requires that

$$d^H\hat{\theta}^H - q^H \geq d^L\hat{\theta}^H - q^L \tag{1.5}$$

$$d^L\hat{\theta}^L - q^L \geq d^H\hat{\theta}^L - q^H. \tag{1.6}$$

[5]For some levels of transparency e it may be optimal to offer credit to disapproved borrowers. However, we will see in Section 1.3 that this does not hold true once we allow B to choose the level of opacity optimally.

Additionally, the contract must be feasible, i.e. the borrower can not default on his debt with certainty:

$$d^{\hat{s}} \leq \rho_p(A + q^{\hat{s}}) \tag{1.7}$$

We can check that constraint (1.5) must always be binding in equilibrium. Otherwise, the lender could offer a contract that makes slightly higher profits when he disapproves of a borrower and slightly lower profits in case he approves of a borrower. In expectation, this contract would transfer wealth towards borrowers with more profitable investment opportunities. Hence, ex-ante a borrower would strictly prefer this new contract. Since it is efficient to transfer funds to those projects that are more profitable in expectation, the lender will treat approved borrowers as favourable as he can without having an incentive to misreport a good signal as a bad one. Finally, we can ignore the second truth-telling constraint and check that it is indeed not binding in any optimum. This allows us to characterise the optimal ex-ante contract as follows:

Proposition 1. *An optimal ex-ante contract will take either of the two following forms:*

- *It offers credit if and only if the lender has received a positive signal. Credit is priced fairly and approved borrowers promise all pledgable income to their relationship lender:* $q^L = d^L = 0$, $q^H = d^H \hat{\theta}_H$ *and* $\left(A + q^H\right)\rho_p = d^H$.
- *It offers to fund borrowers irrespective of the signal* s*. All borrowers receive the same interest rate and promise all pledgable income to their relationship lender:* $q^{\hat{s}} = d^{\hat{s}}\left(\alpha\theta_H + (1-\alpha)\theta_L\right)$ *and* $\left(A + q^{\hat{s}}\right)\rho_p = d^{\hat{s}}$ *for all* $\hat{s} \in \{H, L\}$.

Proof. See the appendix.

Lenders on the spot market rationally expect that they will only trade with borrowers that the relationship lender disapproves of. So they charge an interest rate that exceeds the expected cost of lending to approved borrowers. Hence, it is optimal for borrowers to obtain credit exclusively from their relationship lender in case of a positive evaluation. At the same time, a relationship lender must have an incentive to report his signal s truthfully. One way to achieve this is by offering a contract that makes zero profits on approved borrowers and that grants no credit at all in case of a negative evaluation. In this case the lender has no incentive to misreport a good signal as a bad one and borrowers with a negative evaluation can turn to transactional lenders instead. Alternatively, the lender can offer funding to borrowers with an unfavourable rating, too. But if he does so, he must grant them credit at advantageous conditions that do not reflect the true risk of lending to them. Approved borrowers, on the other hand, have to cover the corresponding losses by paying an interest rate d^H/q^H that lies above the actuarially fair one. This ensures that a relationship lender can not increase his profits by misreporting a favourable signal as a bad one. The more advantageous the conditions for disapproved borrowers are, the smaller are the interest payments that a lender can extract from approved borrowers by misreporting his signal. Additionally, by misreporting the lender would lose out on the profits that he makes on approved borrowers in case he truthfully reveals his information. Since a disapproved borrower can obtain additional funding on the spot market at a constant interest rate r_M, a borrower's expected utility is linear in the size of the credit that he obtains from his relationship bank. This implies that an optimal contract will either not grant any credit to disapproved borrowers or it will allow disapproved borrowers to borrow as much as possible without violating the feasibility constraint (1.7).[6] In this case the truth-telling constraint requires the subsidy to disapproved borrowers to be particularly high and financing conditions do not depend on $\hat{s}$ at all.

If the equilibrium interest rate r_M that is demanded on the spot market is low, an optimal ex-ante contract does not offer credit to disapproved borrowers and these borrowers turn to the spot market instead. Any borrower that does receive funding is likely to have a high quality project. These borrowers get favourable conditions that allow them to invest at a large scale and they do not trade on the spot market. Disapproved borrowers on the other hand need to turn to the spot market, pay higher interest rates in equilibrium and invest at a smaller scale.

Alternatively, consider a situation in which the interest rate demanded on the spot market

[6]We abstract from the knife-edge case in which the borrower is indifferent between a continuum of different contracts stipulating different levels of q^L. This simplification does not affect any of our results.

is high. In this case it is optimal to sign a contract that offers as much credit as possible to both types of borrowers. Borrowers receive the same level of funding irrespective of the signal the relationship lender has obtained, i.e. we have a contract that has pooling between borrowers with different signals. While a borrower is able to circumvent the spot market that might otherwise be subject to a market freeze, the relationship lender can not make use of his information in order to channel funds towards the more efficient projects. So all borrowers receive credit at the same conditions and will invest at the same scale. Note that for $e = 0$ the expected utility of a borrower is the same for both types of contracts. If the signal a relationship lender obtains is completely uninformative, it does not matter if relationship banks offer credit to both types of borrowers or not. Transactional lenders know that the spot market constitutes a random sample of borrowers and they offer exactly the same conditions as a relationship lender.

For either type of contract, any borrower who does trade with the relationship bank receives as much credit as possible and whenever borrowers trade with a lender irrespective of $\hat{s}$, they are treated equally. This allows us to fully describe contracts by an interest rate $r_R = d^{\hat{s}}/q^{\hat{s}}$ that a borrower has to pay and a set of reports $\Sigma \in \{(H, L), (H)\}$ for which he has the right to borrow at this rate. The quantity of credit that a borrower receives is pinned down by the feasibility constraint. So the optimal ex-ante contracts resemble credit lines as they can be observed in reality: Credit lines specify an interest rate at which a borrower is able to withdraw funds in the future. While in some cases this right is unaffected by new information that a lender learns about a potential borrower, Sufi (2009) shows that credit lines are often subject to strict covenants and are regularly withdrawn by lenders. Indeed, many lenders reserve the right to withdraw credit lines whenever they see fit. This corresponds to a situation where the lender's report on a non-verifiable signal s determines whether a prospective borrower receives credit at pre-specified conditions or not.

Let us now turn to the welfare effects of transparency. We will denote the borrower's expected utility in case he signs a contract that only grants credit in case of a positive rating by $W(e)$. In case the borrower signs a contract that offers credit to all debtors, the

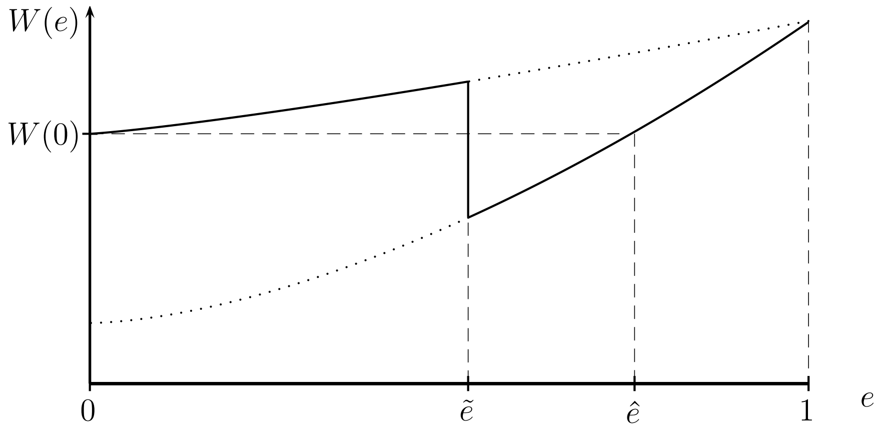

Figure 1.2: Welfare as a Function of Transparency.

borrower receives an expected utility of $W(0)$ irrespective of e. Since the contract does not condition on the lender's information, the borrower receives the same level of utility as if the lender received no information at all. Hence, we can restrict our discussion to the case where the lender offers credit to a borrower only if he has obtained a favourable signal. In general, transparency is beneficial since it allows banks to direct funds to the more efficient projects. Slightly increasing the scale of the more efficient projects while reducing the scale of the less efficient ones would always be socially beneficial. Improved transparency allows the banking sector to do just that: Since projects of approved borrowers become increasingly likely to be of type $i = H$, an approved borrower can obtain more funds when pledging his future income to a relationship lender and the scale of his project can be increased. Disapproved borrowers on the other hand need to pay higher interest rates on the spot market and, hence, need to reduce the size of their investments. An additional advantage of information disclosure is that it increases the probability that good projects (i.e. projects with $i = H$) are the ones financed at a large scale. Since all pledgable income ends up with the lenders and lenders make zero profits, this efficiency gain of transparency materialises in a larger average scale of the investment projects. However, Figure 1.2 illustrates that social welfare drops sharply once a certain level of transparency has been exceeded. Whenever $e > \tilde{e}$, the market for uninformed credits breaks down and borrowers with good projects that end up on the period-1 market will not take out a loan.

The stigma attached to a bank having to borrow on the spot market becomes too large for the spot market to keep working.

In case $e \leq \tilde{e}$ it is easy to see that borrowers always want to sign an ex-ante contract that offers credit only to approved borrowers. Such a contract makes use of the lender's information and the financial sector allocates funds more efficiently than if all borrowers were to receive the same financing terms. At the same time information is sufficiently coarse as to not put the integrity of the spot market in jeopardy.

Lemma 2. *Assume that $e \leq \tilde{e}$. The borrower signs an ex-ante contracts that offers credit if and only if the lender has received a positive signal. This allows him to invest at scale $S = \frac{A}{1-\rho_p \hat{\theta}^H}$ in case of a positive screening.*

Proof. See the appendix.

Now assume that $e > \tilde{e}$: In this case uninformed lenders rationally expect the quality of any borrower who is willing to borrow on the spot market to be θ_L, i.e. the spot market for credit freezes. If relationship banks are to reveal their information, disapproved borrowers have to turn to the spot market. This comes at a cost since high types that received a negative evaluation will now abstain from taking up debt on the spot market. However, as $e \to 1$ the contract characterised in Lemma 2 is still optimal. Even though a market breakdown is inevitable, its costs are negligible. Screening is almost perfect and very few good investment projects end up in need of additional funds on the period-1 market. So while the efficiency gains from making use of the lender's information are seizable, the difference between the borrower's utility in case of a market breakdown and the utility in a hypothetical scenario where the spot market remains intact (depicted by a dotted line in Figure 1.2) is arbitrarily small.

Yet, if e is small and the market freezes, the opposite may hold true. In this case the negative consequences of a market freeze are quite damaging since there is a large probability that good projects end up without funding in period 1 and these borrowers will choose to self-finance, which comes at a high efficiency cost. It may therefore be optimal to sign a contract that does not depend on $\hat{s}$ and that results in expected utility of $W(0)$ in order to circumvent the spot market. The appeal of this kind of contractual arrangement is reinforced by the fact that information is imprecise anyway. If the lender's signal does not contain a lot of information, there would be little benefit in conditioning the contract on the lender's information even if the spot market did not break down. Hence, signing a contract that does not condition on $\hat{s}$ is attractive in comparison.

Proposition 2. *There exists a threshold* $\hat{e}$ *such that whenever* $e \notin (\tilde{e}, \hat{e})$ *the borrower signs an ex-ante contract as described by Lemma 2. If* $e \in (\tilde{e}, \hat{e})$ *the borrower signs a contract that does not depend on the report* $\hat{s}$ *and he invests at scale* $S = \frac{A}{1-\rho_p(\alpha\theta_H+(1-\alpha)\theta_L)}$.

Proof. See the appendix.

We have seen that it is optimal to sign an ex-ante contract that makes us of the lender's information if e is either small or very large. So the only question that remains is whether the interval $(\tilde{e}, \hat{e})$ for which $W(e) < W(0)$ and for which it is optimal to sign a contract that does not depend on $\hat{s}$ is non-empty. This will be the case if θ_L is sufficiently small relative to θ_H.[7] In this case low levels of transparency suffice to lead to a breakdown of the spot market (i.e. $\tilde{e}$ is small). We know from the discussion above that for these low levels of transparency a market breakdown is particularly costly. Hence, it is optimal to circumvent the spot market and to sign contracts that do not condition on the lender's information instead.

To sum up, we have seen that for a given level of transparency e, the borrower faces a choice between foregoing all of the benefits of information disclosure by signing an ex-ante contract that does not depend on $\hat{s}$. Or he signs a contract that forces him to seek credit on a spot market which may be subject to a market freeze. In the next section we allow borrowers to choose their level of financial transparency. This enables borrowers to release some information while at the same time making the information sufficiently opaque for the spot market to keep working. Note that our assumption that contracts with uninformed lenders can only be entered in period 1 is without loss of generality. Instead of trading with uninformed lenders in period 0, a borrower might as well sign

[7] More precisely, the interval $(\tilde{e}, \hat{e})$ will be non-empty whenever $\theta_L \to (\theta_H/(\rho_{np}+\theta_H\rho_p)-\alpha\theta_H)/(1-\alpha)$.

an ex-ante contract that does not condition on $\hat{s}$. In either case he receives the same financing terms with certainty. Moreover, it is impossible to write a forward contract in which the borrower commits to trade with an uninformed lender in case of a negative report since contracts with the relationship lender are unobservable.

1.3 Optimal Opacity

In the last section we have seen that the imperfect nature of the information that relationship lenders receive may have very costly consequences. With a positive probability borrowers that face highly profitable investment projects are not approved and do not receive credit from their relationship lender. Moreover, they may decide to abstain from taking up new debt on the spot market due to asymmetric information. Yet, imperfect information disclosure may also be a solution to this problem. If the information released to lenders is sufficiently imprecise, this will reduce the stigma attached to a borrower who has to resort to the market for uninformed credit and may prevent a market breakdown. Hence, there is a fundamental convexity in the returns to information disclosure. If a borrower can enable his relationship lender to distinguish between projects of different quality with high precision, he should always try to do so. But if the nature of the investment opportunity is such that a certain amount of noise will always persist in the evaluation process, it may be optimal to make sure that even less information on the quality of the project is made available.

Assume B was able to promise a relationship lender ex-ante to make more information available while keeping the decision to increase transparency shrouded from the public eye. In this case the relationship lender would offer more attractive conditions for the case of a favourable rating. But if the borrower ends up with a negative evaluation, he enters the spot market for credit where the terms of credit are exogenously given. The very appeal of increasing opacity is that such tactics are readily observable. While financial obfuscation makes it impossible for a relationship lender to precisely assess the quality of a given claim, the fact that a bank has been engaging in high levels of financial engineering and has created opaque claims is commonly observable. A high level of opacity reassures transactional lenders that they are not at a large informational disadvantage vis-à-vis the relationship lender. Hence, transactional lenders will infer less information from the fact that a borrower is in need of funds in period 1. Note that even if the borrower is able to reduce opacity once an ex-ante contract has been signed, he does not have any incentive to do so.[8]

Proposition 3. *There exists a threshold* $\tilde{E} \in (\tilde{e}, 1)$ *such that*

- *if* $\overline{E} > \tilde{E}$ *a borrower discloses as much information as possible and chooses* $e = \overline{E}$.
- *if* $\overline{E} \leq \tilde{E}$ *a borrower chooses potentially less-than-perfectly transparent financial structures:* $e = \min\left\{\tilde{e}, \overline{E}\right\}$.

Proof. See the appendix.

If it is possible to disclose sufficiently precise information (i.e. if $\overline{E}$ is large), it is always optimal to do so. This allows the banking sector to allocate funds in the most efficient way possible, while the cost of a market breakdown is small since few good projects end up without funding in period 1. But if screening is bound to be sketchy anyway, a borrower optimally restricts the amount of information a relationship lender can obtain and chooses the maximal amount of transparency that does not lead to a freeze of the spot market for debt. This results in an expected utility of $W(\tilde{e})$ and is strictly more attractive than choosing a slightly larger e. Doing so would only marginally reduce the interest rate demanded by relationship lenders but lead to a breakdown of the spot market. It is easy to see that we must have $\tilde{E} > \hat{e}$: In Section 1.2.2 a borrower faced the choice between an ex-ante contract that makes use of the lender's information but leads to a breakdown of the spot market and a contract that does not condition on any information at all.

[8]Assume that an ex-post increase in e is unobservable and in equilibrium the spot market works smoothly. Since an increase in e does not alter the interest rates r_R and r_M or the probability of qualifying for r_R and since the borrower's utility does not depend on his type directly, the change in e does not affect the borrower's expected utility. If the spot market does break down, an increase in transparency makes the borrower strictly worse off: In this case qualifying for r_R is more important for low types than for high types and the borrower has an interest in opacity.

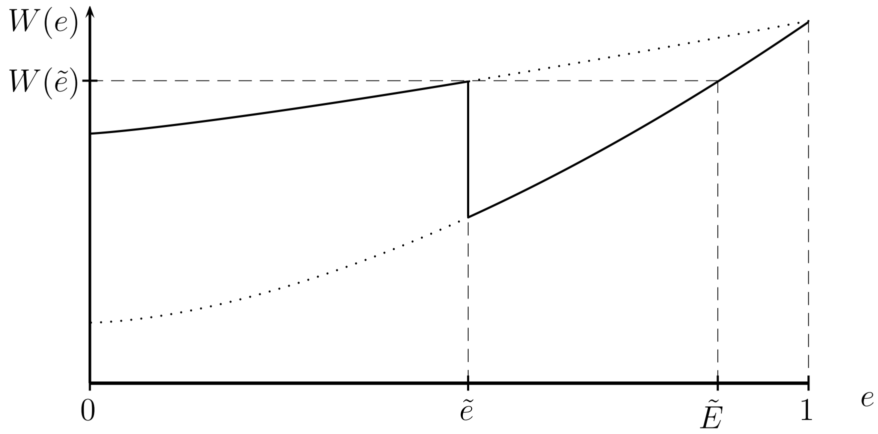

Figure 1.3: The Optimal Level of Transparency.

If the borrower is able to fine-tune the amount of transparency such that he is able to benefit from a functioning spot market while still enjoying some of the efficiency gains of information disclosure, a market breakdown becomes more unattractive in comparison: If a borrower decides to prevent a market breakdown he can now enjoy a utility of $W(\tilde{e})$ instead of $W(0)$.

Note that our model does not try to answer the question of how opacity should best be achieved. Structuring assets in a way that makes it hard to asses their value may be one option. Alternatively, the borrower may reduce the lender's incentive to invest into information acquisition by designing information-insensitive securities as in Dang, Gorton and Holmström (2009), by signing contracts with banks that have poorly incentivized employees, etc. Instead of analysing in detail how opacity can be achieved, we are more interested in the optimal degree of (in)transparency. A fundamental difference to the model of opacity proposed by Dang, Gorton and Holmström (2009) is that in our setting opacity does not ameliorate asymmetric information between the contracting parties: While the borrower always learns his type perfectly, relationship lenders receive less information the more opaque the borrower's financial structure. Moreover, in equilibrium uninformed lenders know that they lend only to those borrowers that received an unfavourable evaluation. So the less precise the information of a relationship lender is, the less does an uninformed lender learn from the fact that he is trading with a particular borrower. Nevertheless, a positive amount of opacity may be optimal because it has positive effects on the market composition of the spot market.

1.4 Government Intervention and Market Support

The danger of a market breakdown can impose considerable costs on the economy: Either asymmetric information does indeed lead to a freeze of the spot market and prevents the best projects from being set up at a large scale. Or borrowers reduce transparency in order to prevent a market breakdown, which comes with efficiency costs of its own. One way for a government to increase efficiency is to subsidise trade on the spot market in order to prevent a freeze. Indeed, the original plans for the Troubled Asset Relief Program ("TARP 1") discussed by the US government in 2008 can be seen as such a policy. By systematically buying up bad assets at a loss, the government can increase the average quality of claims remaining on the market and can help to jump-start interbank trade.

In our analysis we will compare two cases: In the first scenario, we assume that the government can commit to time-inconsistent policies by announcing the maximal level of transparency e for which it is prepared to offer market support before B chooses the level of opacity. In the alternative case the government can not commit to a bailout policy and decides whether to buy claims once borrowers have chosen their level of transparency and the market threatens to break down. The main insight from our analysis is that while banks face considerable leeway to force the government into action ex-post, a government does not always want to rule out any bailouts ex-ante, even if it can commit to do so. Part of the reason why it may be optimal not to rule out bailouts ex-ante is that they can have positive incentive effects and encourage borrowers to disclose more information. This increases the allocative efficiency of the financial sector and is socially desirable as long as the government ensures that the spot market remains liquid.

In order to allow for interventions in which the government only buys some of the claims that would otherwise end up on the spot market, in this section we consider the case where there is a continuum of ex-ante identical borrowers, each of which may have a need for additional funds at $t = 1$. Moreover, there is an infinite supply of funds, so borrowers do not compete for funding. Our model of market support follows along the lines of a stripped-down version of Tirole (2012). More specifically, we assume that a government may enter the period-1 market in order to act like an ordinary (uninformed) lender by publicly offering to purchase claims at an interest rate of r_G. The timing is as follows: At the beginning of $t = 1$ a borrower receives offers from his relationship lender and the government. After this market has closed and before investments have to be made, any borrower who is still in need of funds has the option of selling claims on the competitive spot market made up of transactional lenders. As before, a borrower can not pledge any future income to more than one lender, so the maximal amount of credit he can obtain from the government q_G is implicitly defined by

$$\left(A + q^{\hat{s}} + q_M + q_G\right)\rho_p - d^{\hat{s}}r - q_M r_M = q_G r_G.$$

The fundamental difference between the government and commercial lenders is that the government is concerned with total welfare and does not have to make zero profits since it can use tax revenues to finance welfare-enhancing deals on the interbank market. We assume that taxation creates a deadweight-loss of $\lambda - 1 > 0$, so the government will only restart the market for financial claims if the cost of fiscal funds λ is sufficiently low. In order to concentrate on the case where there is a trade-off between subsidising borrowers and preventing a spot market freeze we make the following assumption:

Assumption 3. *The cost of public funds is not too low:*

$$\lambda > \underline{\lambda} = \frac{\rho_{np}}{1 - \rho_p \theta_H}.$$

Assumption 3 implies that even if the government could be perfectly sure to face high types, it would not want to transfer wealth to these borrowers unconditionally. This implies that a policy like the final version of the Troubled Asset Relief Program ("TARP

2”) that injected money into failing banks is never optimal. Transferring funds to banks is very costly in the best of circumstances. But since banks that are in need of new funds are below average in terms of their financial prospects, injecting money into these banks is even less attractive.

As a consequence of Assumption 3 the government only offers subsidised credits if doing so is necessary in order to prevent a market breakdown. This implies that if the government intervenes, it offers an interest rate that just allows the spot market to keep working, i.e. $r_G = 1/\tilde{\theta}_M$. Moreover, whenever both, commercial lenders on the spot market and the government trade claims with a positive probability, we must have $r_M = r_G$ since borrowers have to be indifferent between taking money from the government or private investors. For simplicity we assume that any given borrower trades either with the government or with a private lender. Still, there is a large number of equilibria which differ with respect to the probability with which a given borrower sells claims to the government. Since all of these equilibria result in the same cost of an intervention and the same utility for borrowers, we assume that borrowers trade with the government as rarely as possible.[9] In this equilibrium good claims never end up in the hands of the government while the government funds bad projects with a positive probability.

Assume that absent any intervention the spot market for financial claims would break down and that the government is willing to intervene. The spot market only remains liquid if the expected quality of debtors who are in need of funding is at least equal to

$$\tilde{\theta}_M = \theta_H/(\rho_p\theta_H + \rho_{np}).$$

This implies that the fraction of claims on the spot market that are of high quality must be at least

[9] In particular, there exists an equilibrium in which the spot market is inactive and all borrowers that are not funded by their relationship lender obtain funds from the government. While in this equilibrium we may have $r_G > r_M$, the equilibrium would again have the same consequences for welfare as the one we consider.

$$\tilde{\alpha}_M = \alpha_M(\tilde{e}) = \frac{\theta_H - (\rho_{np} + \rho_p\theta_H)\,\theta_L}{(\rho_{np} + \rho_p\theta_H)\,(\theta_H - \theta_L)}.$$

In equilibrium, sellers with low-quality claims trade with the government with a probability that guarantees that the average quality of borrowers that are still in need of funds is $\tilde{\theta}_M$. Hence, commercial lenders find it optimal to offer the same interest rate as the government and the spot market for interbank funds doesn't break down. If all borrowers choose the same level of transparency this implies that the government finances a fraction $\left(1 - \frac{\alpha_M(e)}{\tilde{\alpha}_M}\right)$ of the projects that do not receive funding from relationship lenders and makes a loss of $\left(\tilde{\theta}_M - \theta_L\right)\frac{A\rho_p}{1-\rho_p\tilde{\theta}_M}$ on each borrower that it trades with.

Commitment Solution

Consider the case where the government announces a bailout policy before borrowers have chosen e and is able to commit to it. The cost of an intervention is increasing in the level of transparency e since the average quality of claims on the spot market is decreasing in e and the government has to purchase a larger number of "bad" claims in order to guarantee that the average quality of the remaining borrowers is $\tilde{\theta}_M$. Hence, a government may either want to rule out interventions altogether, or it may want to announce a maximal level of transparency e for which it is prepared to offer support if the spot market threatens to break down.

Proposition 4. *Assume the government can commit to a bailout policy and that $\overline{E} > \tilde{e}$. The government will ensure that the spot market for credit remains intact for any level of transparency $e \leq \overline{E}$ if the cost of public funds λ is weakly below $\overline{\lambda}^c(\overline{E})$. Otherwise it does not offer any intervention at all. Whenever the government does offer market support, borrowers choose the maximal level of transparency $e = \overline{E}$.*

Proof. See the appendix.

Even though the government could make interventions conditional on the borrowers' choices of transparency, it will never choose to do so. The cost of an intervention is linearly increasing in the level of transparency that borrowers choose. At the same time, the expected utility of borrowers that receive funding from their relationship lender is convex in the transparency level since they benefit from more advantageous terms of credit. Moreover, given that the government intervenes, the utility of all borrowers that do not obtain funds from their relationship lender does not depend on e since they can always borrow at an interest rate of $1/\tilde{\theta}_M$. Hence, the government will either announce to ensure the liquidity of the spot market irrespective of the borrowers' choices of opacity, or it will announce not to intervene at all.

Since there are no strings attached, any borrower who ends up in need of funds in period 1 will qualify for state aid. So we do not have to consider any participation decision on the side of borrowers and only need to check if the government does indeed find it optimal to offer market support. Recall that there exists some threshold $\tilde{E}$ such that absent any government intervention, a borrower chooses limited transparency in order to guarantee the integrity of the spot market if and only if $\overline{E} \leq \tilde{E}$. In this case, the government is prepared to intervene whenever the benefits of increases transparency exceed the fiscal cost of market support and

$$\left[\frac{1}{1-\rho_p\hat{\theta}^H(\overline{E})} - \frac{1}{1-\rho_p\hat{\theta}^H(\tilde{e})}\right]\frac{\alpha}{1-\alpha}A\rho_{np} \geq \lambda\left(1-\frac{\alpha_M(\overline{E})}{\tilde{\alpha}_M}\right)\left[\frac{\rho_p\left(\tilde{\theta}_M-\theta_L\right)}{1-\rho_p\tilde{\theta}_M}\right]A. \quad (1.8)$$

The threshold $\overline{\lambda}^c(\overline{E})$ denotes the highest cost of public funds for which this condition is satisfied.

Now, consider the case where $\overline{E} > \tilde{E}$: If the government decides not to intervene, borrowers choose maximal transparency and the spot market freezes. In this case an intervention does not affect the level of transparency that borrowers choose, but it prevents a market breakdown. This is socially beneficial and justifies the cost of an intervention whenever

λ is sufficiently small, i.e. if

$$\left[\frac{1}{1-\rho_p\tilde{\theta}_M} - \frac{1}{1-\rho_p\theta_L}\right](1-\alpha_M(\overline{E}))A\rho_{np} \geq \lambda\left(1-\frac{\alpha_M(\overline{E})}{\tilde{\alpha}}\right)\left[\frac{\rho_p\left(\tilde{\theta}_M-\theta_L\right)}{1-\rho_p\tilde{\theta}_M}\right]A. \quad (1.9)$$

Again, the threshold $\overline{\lambda}^c(\overline{E})$ denotes the level of λ for which the condition is satisfied with equality. Note that in case $\overline{E} > \tilde{E}$ the benefit of an intervention is restricted to borrowers with bad projects who end up on the spot market. The terms offered by relationship lenders are the same with or without intervention since borrowers choose the same level of transparency $\overline{E}$. Moreover, high types who end up on the spot market receive a utility that is equal to the one they could obtain by borrowing at a rate of $1/\tilde{\theta}_M$ in either case: For this interest rate they are indifferent between taking out a loan or choosing self-financing. But the participation of high types in the spot market has a positive externality on low types who are now able to borrow at a considerably lower interest rate.

Henceforth, we will call an intervention "unconditional" if the government is prepared to offer market support regardless of a borrower's level of transparency and "conditional" otherwise. More generally, we will refer to an intervention as "larger" if borrowers have chosen more transparent financial structures, since in this case the government needs to buy a larger number of claims in order to keep the spot market working.

No Commitment Solution

Let us now consider the problem faced by a government that lacks commitment power. In this case the government can not convince borrowers that it will follow any policy that is not ex-post optimal. Hence, without loss of generality we can assume that the government only decides to intervene once borrowers have chosen a level of transparency e that is incompatible with a functioning spot market. The key difference to the case where the government has commitment power is that now, borrowers can strategically choose their level of transparency in order to affect the government's decision whether to intervene or not. In the case of commitment power, this decision was taken by the government ex-ante and could no longer be influenced by the borrowers.

Proposition 5. *Assume the government can't commit to a bailout policy and that* $\overline{E} > \tilde{e}$*. The government will ensure that the spot market for credit remains intact if and only if* $e \leq \hat{E}^{nc}$ *where* $\hat{E}^{nc}$ *is implicitly defined by*

$$\alpha_M(\hat{E}^{nc}) = \frac{\lambda\,(1-\rho_p\theta_L)-\rho_{np}}{\lambda\,(1-\rho_p\theta_L)-\tilde{\alpha}_M\rho_{np}}\tilde{\alpha}_M.$$

If $\overline{E} \leq \tilde{E}$ *borrowers always benefit form an intervention and choose a level of transparency* $e = \min\{\hat{E}^{nc}, \overline{E}\}$.
If $\overline{E} > \tilde{E}$ *borrowers benefit from an intervention and choose* $e = \min\{\hat{E}^{nc}, \overline{E}\}$ *if* $\lambda \leq \overline{\lambda}^{nc}(\overline{E})$ *and they choose* $e = \overline{E}$ *otherwise.*

Proof. See the appendix.

We can check that $\hat{E}^{nc} > \tilde{e}$, i.e. the government always intervenes if transparency is not too high: If the level of transparency is sufficiently close to $\tilde{e}$, the government only has to buy an arbitrarily small number of claims at a loss in order to prevent a full-scale market freeze. Hence, it is always optimal to do so. Moreover, the higher the fiscal cost of an intervention λ, the smaller the maximal level of transparency for which the government is prepared to intervene, i.e. $\hat{E}^{nc}$ is decreasing in λ.

The most important difference to the case where the government has commitment power is that a government may end up offering interventions only conditionally. That is, it will only intervene if borrowers have not chosen too high a level of transparency. The reason why such limited bailouts can occur in case the government lacks commitment power is that they are a direct result of extortionary practices. By choosing intermediate levels of transparency, borrowers reduce the cost of an intervention. But more importantly, they increase the cost of a market breakdown in case the government fails to intervene. If transparency is low, a large number of high types fail to receive funding from their relationship lender. If the government doesn't intervene, these borrowers don't take out any loans which comes at a high cost in terms of efficiency. Hence, by making the alternative to an intervention less appealing, borrowers can force the government into action ex-post. However, conditional interventions are never socially optimal and will not obtain in equilibrium if the policy maker can commit to a bailout regime.

If borrowers anticipate that the government may intervene, this will affect the level of transparency they choose in period 0. Let us first consider the case where $\overline{E} \leq \tilde{E}$, i.e. absent any government intervention a borrower would choose to limit transparency in order to guarantee the integrity of the spot market. In this case, a government intervention allows borrowers to disclose additional information without provoking a breakdown of the spot market. A larger level of transparency reduces the interest rate demanded by relationship banks while the interest rate that borrowers need to pay on the spot market is given by $r_M = 1/\tilde{\theta}_M$ and does not depend on e. Hence, borrowers will always choose the maximal level of transparency for which the government is willing to intervene.

Instead, assume that $\overline{E} > \tilde{E}$ so absent any intervention, a borrower would choose to disclose as much information as possible. This allows the borrower to enjoy the full efficiency gains of information disclosure but results in a freeze of the spot market. In this case a borrower will only choose to benefit from a government intervention if he doesn't have to restrict the level of transparency too much in order to do so. This is the case if λ is small and the government is prepared to intervene even for high levels of transparency. If on the other hand λ is large, $\hat{E}^{nc}$ is close to $\tilde{e}$ and the government can only be induced to offer very limited interventions. If the borrower preferred to disclose as much information as possible in the absence of a government intervention, he will still do so now and will not make use of a government scheme.

After having looked at the cases where the government does or does not have commitment power separately, we can now compare the two cases and summarise our results.

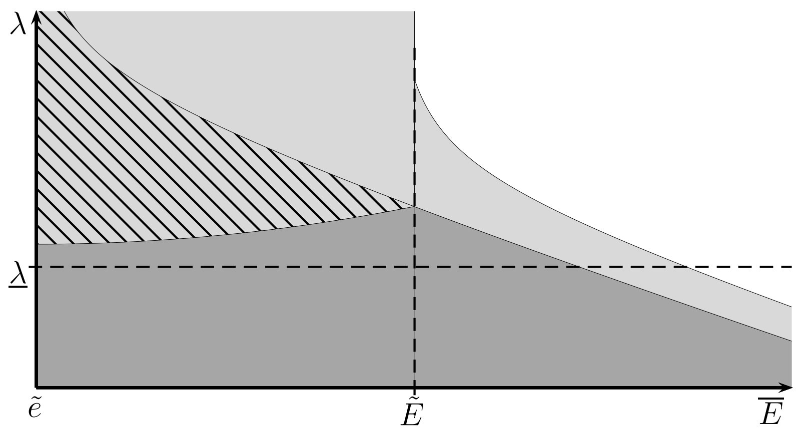

Figure 1.4: Government Interventions with/without Commitment.

Proposition 6. *If the government is unable to commit to a time-inconsistent policy, it offers weakly larger interventions in equilibrium.*

Nevertheless, if α *is large there exists a non-empty set* $[\underline{\lambda}, \overline{\lambda}^c(\overline{E})]$ *such that the government offers market support if* $\lambda \in [\underline{\lambda}, \overline{\lambda}^c(\overline{E})]$ *and* $\overline{E} \in (\tilde{e}, \tilde{E}]$ *even if it can commit to a bailout policy. In this case an intervention induces borrowers to implement maximal transparency.*

Proof. See the appendix.

A summary of the optimal policies is given in Figure 1.4. The dark area depicts situations in which the government will offer unconditional interventions regardless of whether it has commitment power or not. The dashed area describes situations in which a government would like to commit to a laissez-faire policy but will end up offering unconditional interventions if it is unable to do so. The light area corresponds to parameter constellations in which a government that lacks commitment power will end up offering conditional interventions (i.e. $\hat{E}^{nc} < \overline{E}$), while a government with commitment power doesn't offer any intervention at all. Finally, the white area depicts situations in which the government does never intervene in equilibrium.

In order to see that the government will offer weakly larger interventions in equilibrium if it has no commitment power, let us first consider the case where $\overline{E} \leq \tilde{E}$. Absent any intervention, a borrower would restrict information disclosure. Hence, if a government can commit to a bailout policy, it will compare the benefits of an intervention to the alternative of limited information disclosure. However, if the government is unable to commit to a policy, a borrower can choose to adopt maximal transparency before the government decides on a bailout policy. Now the government compares the benefits of an intervention to the alternative of a market breakdown. For $\overline{E} \leq \tilde{E}$ a market breakdown is more costly than limiting information disclosure and a government that lacks commitment power is hence more likely to offer unconditional interventions than a government that can commit to time-inconsistent policies. Moreover, in case the government lacks commitment power, there are values for λ for which the government is no longer prepared to offer unconditional bailouts, but it can still be induced to offer small interventions, i.e. $\hat{E}^{nc} < \overline{E}$. By choosing a level of transparency that is sufficiently close to $\tilde{e}$ a borrower can always force the government to offer some (potentially small) intervention regardless of λ.

Instead, assume that $\overline{E} > \tilde{E}$. In this case a government will compare the benefits of an intervention to the alternative of a market breakdown even if it can commit to a bailout policy. So a government with commitment power will offer unconditional market support if and only if a government that lacks commitment power would do so. However, a government that lacks commitment power may still end up offering conditional interventions in cases where it can not be induced to intervene irrespective of e.

It is often argued that while schemes such as TARP may be ex-post optimal, a government might want to rule out interventions ex-ante in order not to create incentives for banks to rely on such bailouts. Proposition 6 verifies this intuition by showing that borrowers can create situations in which the government feels compelled to intervene ex-post, even though it would prefer to commit to a laissez-faire policy ex-ante. But there are still situations in which market support is optimal even from an ex-ante perspective. This may be due to two different reasons: Either because government intervention prevents an inevitable market breakdown, or because it induces borrowers to disclose more information, making the financial sector less opaque. The latter effect is particularly interesting, since it implies that offering a bailout may be optimal because of the incentives that it creates ex-ante. If borrowers expect to receive government support in case they end up on the spot market, they are prepared to release more information and to make the financial system more transparent. This insight qualifies standard concerns about the moral hazard created by government interventions like TARP.

1.5 Conclusion

We have seen that banks may choose to design purposefully opaque financial instruments in order to limit the amount of information that a relationship lender can obtain. If informed lenders are unable to collect much information, their actions have less informational content and will be interpreted less strongly by transactional lenders: Transactional lenders may now be prepared to lend money to borrowers who have lost funding from their relationship lender at conditions that everybody finds agreeable. There is an optimal amount of information that borrowers choose to disclose and that allows them to enjoy as much of the efficiency gains of transparency as possible while at the same time keeping the cost of asymmetric information in check. Moreover, we have seen that there is a fundamental non-concavity in the returns to information disclosure: If borrowers can disclose very precise information, they will always choose to do so. But if information disclosure is bound to be imprecise anyways, borrowers may find it optimal to make even less information available as to prevent a breakdown of the spot market for interbank credit.

Even though opacity exacerbates asymmetric information between contracting parties, it is beneficial since it affects the composition of the spot market for interbank credit and may thereby prevent a market freeze. Our model extends previous results on the optimality of opaque financial structures by showing that such structures can contribute to market liquidity even if the issuer of a claim can not be kept uninformed. This insight contributes to our understanding of why many firms decide to issue opaque financial claims, choose off-balance-sheet vehicles, intransparent accounting practices etc. even though they reduce the financial sector's ability to allocate funds towards the most efficient projects.

Finally, we have seen that in a framework where banks optimally choose to be intransparent, government interventions such as the Troubled Asset Relief Program can have desirable ex-ante effects. If borrowers expect the government to prevent a market freeze in the future, they are prepared to disclose more information and to increase financial transparency. This increases the allocative efficiency of the financial sector and may warrant the fiscal cost of a bailout program.

A1 Mathematical Appendix

Proof of Lemma 1. Without loss of generality, we can assume that lenders on the spot market offer an interest rate in the interval $[1/\theta_H, 1/\theta_L]$ in equilibrium since otherwise they would make losses (profits) irrespective of the market composition. Hence, borrowers with bad projects will always trade on the spot market: By condition (1.1) we know that $\rho_{np} + \theta_L\rho_p < \frac{\rho_{np}}{1-\rho_p\theta_L}$ and if it is beneficial to take out a loan at an interest rate of $1/\theta_L$, it is optimal to do so for all lower interest rates.

If both types of borrowers decide to borrow on the spot market lenders offer an interest rate of $r_M = 1/\hat{\theta}_M(e) = 1/(\theta_L + \alpha(1-e)(\theta_H - \theta_L))$ and make zero profits. Assumption 2 ensures that for $e = 0$ high types want to take out loans on the spot market, while Assumption 1 implies that they prefer autarky if $e = 1$. Since taking up debt becomes monotonously less attractive in the interest rate there is a unique threshold $\tilde{e} \in (0,1)$ such that high types will indeed take out a loan on the spot market whenever $e \leq \tilde{e}$. This threshold is implicitly defined by $A(\rho_{np} + \rho_p\theta_H) = \frac{\rho_{np}A}{1-\rho_p\hat{\theta}_M(\tilde{e})}$. It is easy to verify that for all $e < \tilde{e}$ the equilibrium is unique.

Instead, consider the case where $e > \tilde{e}$. The lowest interest rate that lenders can offer without making losses is given by $1/(\theta_L + \alpha(1-e)(\theta_H - \theta_L))$. However, for this rate high types will choose autarky and only low types take out loans. Since lenders have to make zero profits this implies that in equilibrium they have to offer an interest rate of $r_M = 1/\theta_L$. □

Proof of Proposition 1. First, let us show that for any optimal contract, constraint (1.5) must be binding. Consider otherwise. In this case the lender could reduce q^L by one unit and d^L by ρ_p units while increasing q^H by $\frac{1-\alpha}{\alpha}$ units and d^H by $\frac{1-\alpha}{\alpha}\frac{\hat{\theta}^L}{\hat{\theta}^H}\rho_p$ units. This change of the contract leaves the profits of the lender unaffected and allows disapproved borrowers to take up the same amount of debt on the spot market. Moreover, it relaxes the constraint (1.6) and the feasibility constraint (1.7) for approved borrowers. Finally, it is easy to check that it increases the ex-ante expected utility of the borrower even if an approved borrower doesn't take up additional credit on the spot market. So (1.5) must be binding. For now, we will ignore condition (1.6) and we will check at the end that it is indeed satisfied in our candidate optimum. Substituting the zero profit condition into (1.5) gives us the relations $q^L = d^L(\alpha\theta_H + (1-\alpha)\theta_L)$ and $q^H = \hat{\theta}^H d^H - (1-\alpha)(\hat{\theta}^H - \hat{\theta}^L)d^L$. In any equilibrium we must have $r_M > 1/\hat{\theta}^H$. Otherwise, it would be optimal to set $q^L = 0$, disapproved borrowers would borrow on the spot market and transactional lenders would make negative profits. This implies that the marginal interest rate that approved borrowers have to pay when we increase the loan they receive from their relationship bank $\partial d^H/\partial q^H$ is strictly smaller than the interest rate on the spot market. Hence, approved borrowers trade exclusively with their relationship lender and borrow as much as possible. Using the feasibility constraint (1.7) we get $q^H = \frac{1}{1-\hat{\theta}^H\rho_p}\left(A\hat{\theta}^H\rho_p - (1-\alpha)(\hat{\theta}^H - \hat{\theta}^L)d^L\right)$. We can check that irrespective of whether borrowers with good projects decide to drop out of the spot market or not, the borrower's ex-ante expected utility is linear in q^L. If we abstract from the knife-edge case in which the borrower is indifferent between a continuum of contracts stipulating different levels of q^L this leaves us with two scenarios: Either a contract has $q^L = d^L = 0$, $q^H = \frac{\hat{\theta}^H\rho_p}{1-\hat{\theta}^H\rho_p}A$ and $d^H = \frac{\rho_p}{1-\hat{\theta}^H\rho_p}A$. Or the contract offers disapproved borrowers as much credit as possible without violating the feasibility constraint. In this case we get $q^L = q^H = \frac{(\alpha\theta_H+(1-\alpha)\theta_L)\rho_p}{1-(\alpha\theta_H+(1-\alpha)\theta_L)\rho_p}A$ and $d^L = d^H = \frac{\rho_p}{1-(\alpha\theta_H+(1-\alpha)\theta_L)\rho_p}A$, i.e. the terms of credit are independent of $\hat{s}$. It is easy to verify that both contracts do indeed satisfy condition (1.6). Without loss of generality we can assume that uninformed lenders expect the quality of borrowers who are in need of new funding at $t = 1$ to be $\hat{\theta}^L$. While this belief is not uniquely pinned down by rational expectations in case the spot market is inactive in equilibrium, any equilibrium that can be supported by some off-equilibrium belief can also be supported by a belief that the expected quality of borrowers in need of funds is $\hat{\theta}^L$. □

Proof of Lemma 2. Let us denote the borrower's expected utility in case he signs a contract that depends on $\hat{s}$ and in case all disapproved borrowers trade on the spot market by $W^L(e)$. If the borrower signs an ex-ante contract that does not condition on $\hat{s}$, his expected utility is given by $W^L(0)$: The utility is the same as if the signal of the relationship lender contained no information at all. In this case lenders on the spot market would infer no information from the fact that a borrower needs funding and a borrower receives funds at an interest rate of $1/(\alpha\theta_H + (1-\alpha)\theta_L)$ no matter if he borrows from his relationship lender or transactional lenders. If on the other hand a borrower signs an ex-ante contract that offers credit only to approved types and everybody else trades on the spot market we get $W^L(e) = A\rho_{np}\left[\frac{\alpha}{1-\rho_p\hat{\theta}^H} + \frac{(1-\alpha)}{1-\rho_p\hat{\theta}_M}\right]$ where $\hat{\theta}_M = \theta_L + \alpha(1-e)(\theta_H - \theta_L)$ and $\hat{\theta}^H = \theta_H - (1-\alpha)(1-e)(\theta_H - \theta_L)$, so we get

$$\frac{\partial W^L(e)}{\partial e} = A\rho_{np}\left[\frac{(1-\alpha)\alpha\rho_p}{\left(1-\rho_p\hat{\theta}^H\right)^2} - \frac{(1-\alpha)\alpha\rho_p}{\left(1-\rho_p\hat{\theta}_M\right)^2}\right](\theta_H - \theta_L) > 0 \quad \text{(A1.1)}$$

for all $e \in [0,1]$. This implies that it must be optimal to sign a contract that has $q^L = d^L = 0$ for all $e \in [0, \tilde{e}]$ since in these cases the spot market does indeed not break down. □

Proof of Proposition 2. Assume the borrower signs a contract that conditions on $\hat{s}$ and the spot market does break down. In this case a borrower with a good project who ends up on the spot market doesn't take out a loan and receives a utility equal to the level he could obtain by borrowing at a rate of $1/\tilde{\theta}_M$ since for this interest rate he is indifferent between taking up credit on the spot market and choosing autarky. So we can express the expected utility of an entrepreneur as $W^{BD}(e) = A\rho_{np}\left[\frac{\alpha}{1-\rho_p\hat{\theta}^H} + \frac{(1-\alpha)(1-\alpha_M)}{1-\rho_p\theta_L} + \frac{(1-\alpha)\alpha_M}{1-\rho_p\tilde{\theta}_M}\right]$. Since $\alpha_M(1) = 0$ and $\hat{\theta}_M(1) = \theta_L$ we get $W^{BD}(1) = W^L(1)$. Hence, if $e \to 1$ it is always optimal to sign an ex-ante contract that conditions on $\hat{s}$ instead of signing a contract that does not depend on $\hat{s}$ and receiving a utility of $W^L(0)$. Next, let us look at the marginal effect of transparency on B's expected utility in case the market breaks down:

$$\begin{aligned}\frac{\partial W^{BD}(e)}{\partial e} &= A\rho_{np}\left[\frac{\rho_p\alpha(1-\alpha)(\theta_H-\theta_L)}{\left(1-\rho_p\hat{\theta}^H\right)^2} + \frac{(1-\alpha)\alpha}{1-\rho_p\theta_L} - \frac{(1-\alpha)\alpha}{1-\rho_p\tilde{\theta}_M}\right] \\ &= \alpha(1-\alpha)A\rho_{np}\left[\frac{\rho_p(\theta_H-\theta_L)}{\left(1-\rho_p\hat{\theta}^H\right)^2} - \frac{\rho_p\left(\tilde{\theta}_M-\theta_L\right)}{(1-\rho_p\theta_L)\left(1-\rho_p\tilde{\theta}_M\right)}\right] > 0 \quad \text{(A1.2)}\end{aligned}$$

Since (A1.2)>(A1.1) there exists some $\breve{e} > 0$ such that $W^L(0) = W^{BD}(\breve{e})$. This implies that a borrower strictly prefers a contract that doesn't depend on $\hat{s}$ for all $e \in (\tilde{e}, \hat{e})$ where $\hat{e} = \max\{\tilde{e}, \breve{e}\}$. Note that as $\theta_L \to (\theta_H/(\rho_{np} + \theta_H \rho_p) - \alpha\theta_H)/(1-\alpha)$ we get $\tilde{e} \to 0$ and the interval $(\tilde{e}, \hat{e})$ is non-empty. □

Proof of Proposition 3. The proof is straightforward and follows along the lines of Proposition 2. Since a transparency level of $\tilde{e}$ can prevent a market breakdown, the borrower can always obtain a utility of $W^L(\tilde{e})$ by restricting transparency just enough to ensure market liquidity. Since $W^L(1) = W^{BD}(1)$, full information disclosure is always optimal if $\overline{E} \to 1$. By the same logic as before, there exists some $\tilde{E} > \tilde{e}$ such that $W^L(\tilde{e}) = W^{BD}(\tilde{E})$. So the borrower restricts information disclosure if $\overline{E} \leq \tilde{E}$ and discloses all information if $\overline{E} > \tilde{E}$. Since $W^L(\tilde{e}) > W^L(0)$ we know that $\tilde{E} > \hat{e}$. □

Proof of Proposition 4. Without loss of generality, we can assume that the government will buy just enough claims in order to ensure market liquidity. Otherwise the government would transfer more wealth than necessary to borrowers which can not be optimal since $\lambda > \underline{\lambda}$. So whenever the government intervenes, we have $r_G = 1/\tilde{\theta}_M$. Moreover, we can restrict attention to the case where $\overline{E} > \tilde{e}$ since otherwise the spot market never breaks down. If the government intervenes, it has to finance a fraction $\left(1 - \frac{\alpha_M(e)}{\tilde{\alpha}_M}\right)$ of the projects that do not receive funding from relationship lenders and it makes a loss of $\left(\tilde{\theta}_M - \theta_L\right) \frac{A\rho_p}{1-\rho_p\tilde{\theta}_M}$ on each borrower. Let us assume that the borrower always chooses the maximal level of transparency for which the government announces to intervene. In this case the government can choose e directly in order to maximise total surplus.

$$\max_e W = \left[\alpha \frac{1}{1-\rho_p\hat{\theta}^H(e)} + (1-\alpha)\frac{1}{1-\rho_p\tilde{\theta}_M}\right] A\rho_{np} - (1-\alpha)\lambda\left(1 - \frac{\alpha_M(e)}{\tilde{\alpha}}\right)\left[\frac{\rho_p\left(\tilde{\theta}_M - \theta_L\right)}{1-\rho_p\tilde{\theta}_M}\right] A.$$

Using the fact that $\hat{\theta}^H(e) = \theta_H - (1-\alpha)(1-e)\left(\theta_H - \theta_L\right)$ we can easily see that the optimisation problem is globally convex: While the cost of larger market interventions is linear in e, the benefit of increased transparency is convex. So the government will either not intervene at all, or it will offer market support for all $e \leq \overline{E}$. Whenever the government does offer an intervention, the borrower's expected utility is monotonically increasing in e over the interval $[0, \overline{E}]$ and the borrower does indeed choose the maximal level of transparency for which the government promises to intervene. If the borrower chooses $e \leq \tilde{e}$ the government has no reason to intervene. So the borrower's expected utility is the same as in Section 1.3 and monotonically increasing in e over the interval $[0, \tilde{e}]$. Moreover, in the interval $[\tilde{e}, \overline{E}]$ the borrower's expected utility is increasing in e, too: A larger level of transparency reduces the interest rate d^H/q^H demanded by relationship lenders while the interest rate on the spot market is $1/\tilde{\theta}_M$ and does not depend on e.

In order to derive the conditions under which the government is willing to offer market support, we have to compare the losses the government makes in case of an intervention to the change in the expected utility of borrowers. In case the government doesn't intervene and $\overline{E} \leq \tilde{E}$ borrowers would voluntarily restrict information disclosure in order to ensure that the spot market for debt does not break down. So the condition for an intervention is given by equation (1.8). If on the other hand $\overline{E} > \tilde{E}$ borrowers would choose to disclose as much information as possible in case the government announces not to intervene. In this case the condition for an intervention is given by (1.9). In either case we can solve for a threshold $\overline{\lambda}^c$ such that the government intervenes if and only if $\lambda \leq \overline{\lambda}^c$. □

Proof of Proposition 5. Assume that borrowers have chosen a level of transparency $e > \tilde{e}$, i.e. absent any intervention the spot market would freeze. The government intervenes whenever the benefit from preventing a market freeze exceeds the cost of an intervention:

$$\left[\frac{1}{1-\rho_p\tilde{\theta}_M} - \frac{1}{1-\rho_p\theta_L}\right] A\rho_{np}(1-\alpha_M(e)) \geq \lambda \left(1 - \frac{\alpha_M(e)}{\tilde{\alpha}_M}\right) \left[\frac{\rho_p\left(\tilde{\theta}_M - \theta_L\right)}{1-\rho_p\tilde{\theta}_M}\right] A. \quad \text{(A1.3)}$$

The benefit of an intervention is restricted to low types who end up in need of new funds. Approved borrowers will still receive credit from their relationship lender and high types that end up on the spot market receive the same level of utility no matter if they choose autarky or if they borrow at an interest rate of $1/\tilde{\theta}_M$. The cost of an intervention on the other hand is given by the share of claims that the government needs to buy times the loss that it makes on each claim. We can check that there exists a lower bound on $\alpha_M(e)$ that guarantees that a government will want to intervene if and only if $\alpha_M(e)$ is larger than this lower bound. Solving for $\alpha_M(e)$ gives us the expression given in the proposition. Since $\alpha_M(e)$ is decreasing in e, the lower bound on $\alpha_M(e)$ translates into an upper bound on e which we denote by $\hat{E}^{nc}$.

By the same argument as in the proof of Proposition 4, the borrowers expected utility is increasing in the interval $[0, \hat{E}^{nc}]$. Moreover, while the borrower's expected utility drops discontinuously once he chooses $e > \hat{E}^{nc}$ and no longer benefits from an intervention, his expected utility increases monotonically again over the interval $(\hat{E}^{nc}, \overline{E}]$. So the borrower will either choose $e = \hat{E}^{nc}$ or $e = \overline{E}$. He prefers to choose $e = \hat{E}^{nc}$ if and only if

$$\frac{\alpha}{1-\rho_p\hat{\theta}^H(\overline{E})} + \frac{(1-\alpha)\alpha_M(\overline{E})}{1-\rho_p\tilde{\theta}_M} + \frac{(1-\alpha)(1-\alpha_M(\overline{E}))}{1-\rho_p\theta_L} \leq \frac{\alpha}{1-\rho_p\hat{\theta}^H(\hat{E}^{nc})} + \frac{(1-\alpha)}{1-\rho_p\tilde{\theta}_M}$$

where again we use the fact that under a market freeze high types that end up on the spot market receive the same utility as if they were to borrow at a rate of $r_M = 1/\tilde{\theta}_M$.

From the definition of $\tilde{E}$ we know that for all $\overline{E} \leq \tilde{E}$

$$\frac{\alpha}{1-\rho_p\hat{\theta}^H(\overline{E})} + \frac{(1-\alpha)\alpha_M(\overline{E})}{1-\rho_p\tilde{\theta}_M} + \frac{(1-\alpha)(1-\alpha_M(\overline{E}))}{1-\rho_p\theta_L} \leq \frac{\alpha}{1-\rho_p\hat{\theta}^H(\tilde{e})} + \frac{(1-\alpha)}{1-\rho_p\tilde{\theta}_M}.$$

So we know that for all $\overline{E} \leq \tilde{E}$ the borrower will indeed prefer to choose $e = \hat{E}^{nc}$. However, in case $\overline{E} > \tilde{E}$ the condition is satisfied if and only if $\hat{E}^{nc}$ is sufficiently close to $\overline{E}$. Since $\hat{E}^{nc}$ is decreasing in λ this will be the case whenever the cost of public funds is sufficiently low. □

Proof of Proposition 6. Consider the case where $\overline{E} \leq \tilde{E}$. By the definition of $\tilde{E}$ we know that for all $\overline{E} \leq \tilde{E}$

$$\frac{\alpha}{1-\rho_p\hat{\theta}^H(\overline{E})} + \frac{(1-\alpha)\alpha_M(\overline{E})}{1-\rho_p\tilde{\theta}_M} + \frac{(1-\alpha)(1-\alpha_M(\overline{E}))}{1-\rho_p\theta_L} \leq \frac{\alpha}{1-\rho_p\hat{\theta}^H(\tilde{e})} + \frac{(1-\alpha)}{1-\rho_p\tilde{\theta}_M}.$$

or simply

$$\left[\frac{1}{1-\rho_p\tilde{\theta}_M} - \frac{1}{1-\rho_p\theta_L}\right](1-\alpha_M(\overline{E})) \geq \frac{\alpha}{(1-\alpha)}\left[\frac{1}{1-\rho_p\hat{\theta}^H(\overline{E})} - \frac{1}{1-\rho_p\hat{\theta}^H(\tilde{e})}\right].$$

So whenever (1.8) is satisfied, equation (A1.3) is satisfied for $e = \overline{E}$, too. Hence, whenever a government with commitment power offers (unconditional) interventions, a government that lacks commitment power will do so, too. Moreover, a government that lacks commitment power may intervene even if a government with commitment power does not find it optimal to do so.

In case $\overline{E} > \tilde{E}$ a government with commitment power will intervene if and only if condition (1.9) holds. This condition coincides with (A1.3) evaluated at $e = \overline{E}$. So a government with commitment power will offer unconditional interventions if and only if a government without commitment power will do so. While under some circumstance a government that lacks commitment will offer bailouts that are limited in size, a government with commitment power will never offer such conditional interventions.

We can easily check that the threshold $\overline{\lambda}^c(\overline{E})$ is increasing in $\overline{E}$ for all $\overline{E} \in \left(\tilde{e}, \tilde{E}\right)$. So in order to show that $\overline{\lambda}^c(\overline{E}) > \underline{\lambda}$ for all $\overline{E} \in (\tilde{e}, \tilde{E}]$ it is sufficient to show that $\lim_{\overline{E}\to\tilde{e}} \overline{\lambda}^c(\overline{E}) > \underline{\lambda}$. In the limit the intervention needed to ensure market liquidity becomes arbitrarily small and will be welfare-increasing whenever

$$\rho_{np}\frac{(\theta_H-\theta_L)}{\left(1-\rho_p\hat{\theta}^H(\tilde{e})\right)^2} - \frac{\overline{\lambda}^c(\tilde{e})}{\tilde{\alpha}_M}\frac{\left(\tilde{\theta}_M-\theta_L\right)}{1-\rho_p\tilde{\theta}_M} = 0$$

or $\overline{\lambda}^c(\tilde{e}) = \frac{\rho_{np}^2}{\left(1-\rho_p\hat{\theta}^H(\tilde{e})\right)^2}\frac{1}{(\rho_{np}+\theta_H\rho_p)}$ which is increasing in α. Moreover,

$$\lim_{\alpha\to 1}\overline{\lambda}^c(\tilde{e}) = \frac{\rho_{np}^2}{(1-\rho_p\theta_H)^2}\frac{1}{(\rho_{np}+\theta_H\rho_p)} > \underline{\lambda}.$$

So there exists an $\alpha < 1$ such that $\overline{\lambda}^c(\overline{E}) > \underline{\lambda}$ for all $\overline{E} \in (\tilde{e}, \tilde{E}]$. □

A2 Screening Contracts

In Section 1.2 we have assumed that relationship lenders do not offer screening contracts, i.e. for any report $\hat{s}$ they offer a unique contract $(q^{\hat{s}}, d^{\hat{s}})$. In this section we will show that there is indeed an equilibrium in which for a given report $\hat{s}$ all borrowers accept the same contract irrespective of i.

Without loss of generality we assume that for a given report $\hat{s}$ the relationship lender asks the borrower to choose between two options $(q_H^{\hat{s}}, d_H^{\hat{s}})$ and $(q_L^{\hat{s}}, d_L^{\hat{s}})$ at $t = 1$ where in equilibrium a borrower with $i = H$ chooses the first option and borrower with $i = L$ chooses the latter one. Whenever $r_M \leq 1/\tilde{\theta}_M$ all borrowers complement their relationship contract with credit taken up on the outside market and pledge all of their future earnings to lenders. This implies that their utility is independent of i for any given contract and there exists a pooling equilibrium in which $q_H^{\hat{s}} = q_L^{\hat{s}}$ and $d_H^{\hat{s}} = d_L^{\hat{s}}$. From Section 1.2.2 we know that for the optimal pooling contract we do indeed have $r_M \leq 1/\tilde{\theta}_M$ whenever $e \leq \tilde{e}$.

Let us now assume that $r_M = 1/\theta_L$. In this case only low types take out extra credit on the spot market, while high types receive a strictly larger expected utility by not pledging any of their future income to uninformed lenders. This implies that lenders may be able to implement separation by offering contracts that restrict the amount of credit granted to borrowers who claim to be of type $i = H$. Any such contract must satisfy the following incentive compatibility constraints:

$$(q_H^{\hat{s}} + A)\rho_{np} + \theta_H \left[(q_H^{\hat{s}} + A)\rho_p - d_H^{\hat{s}}\right] \geq (q_L^{\hat{s}} + A)\rho_{np} + \theta_H \left[(q_L^{\hat{s}} + A)\rho_p - d_L^{\hat{s}}\right]$$

and

$$\left(q_L^{\hat{s}} + A + \left[\frac{(q_L^{\hat{s}} + A)\rho_p - d_L^{\hat{s}}}{r_M - \rho_p}\right]\right)\rho_{np} \geq \left(q_H^{\hat{s}} + A + \left[\frac{(q_H^{\hat{s}} + A)\rho_p - d_H^{\hat{s}}}{r_M - \rho_p}\right]\right)\rho_{np}$$

for all $\hat{s} \in \{H, L\}$. If we simplify these constraints and take into account that a relationship lender has to make zero profits and has to have an incentive to announce his signal truthfully, we get the following system of constraints that any contract has to satisfy:

$$d_H^H - d_L^H \geq \left(q_H^H - q_L^H\right) r_M \tag{A2.1}$$

$$(\rho_{np} + \theta_H \rho_p)\left(q_H^H - q_L^H\right) \geq \theta_H \left(d_H^H - d_L^H\right) \tag{A2.2}$$

$$d_H^L - d_L^L \geq \left(q_H^L - q_L^L\right) r_M \tag{A2.3}$$

$$(\rho_{np} + \theta_H \rho_p)\left(q_H^L - q_L^L\right) \geq \theta_H \left(d_H^L - d_L^L\right) \tag{A2.4}$$

$$\begin{aligned} &\alpha_H \theta_H d_H^H + (1-\alpha_H)\theta_L d_L^H - \alpha_H q_H^H - (1-\alpha_H) q_L^H \geq \\ &\alpha_H \theta_H d_H^L + (1-\alpha_H)\theta_L d_L^L - \alpha_H q_H^L - (1-\alpha_H) q_L^L \end{aligned} \tag{A2.5}$$

$$\begin{aligned} &\alpha_L \theta_H d_H^L + (1-\alpha_L)\theta_L d_L^L - \alpha_L q_H^L - (1-\alpha_L) q_L^L \geq \\ &\alpha_L \theta_H d_H^H + (1-\alpha_L)\theta_L d_L^H - \alpha_L q_H^H - (1-\alpha_L) q_L^H \end{aligned} \tag{A2.6}$$

$$\begin{aligned} &\alpha \left[\alpha_H \theta_H d_H^H + (1-\alpha_H)\theta_L d_L^H - \alpha_H q_H^H - (1-\alpha_H) q_L^H\right] + \\ &(1-\alpha)\left[\alpha_L \theta_H d_H^L + (1-\alpha_L)\theta_L d_L^L - \alpha_L q_H^L - (1-\alpha_L) q_L^L\right] = 0 \end{aligned} \tag{A2.7}$$

In order to show that there exists a pooling equilibrium in case $r_M = 1/\theta_L$ we proceed in three steps: Step 1): Conditions (A2.1) and (A2.2) (or, equivalently, (A2.3) and (A2.4)) can be re-expressed as $(\rho_{np}+\theta_H\rho_p)(q_H^{\hat{s}}-q_L^{\hat{s}}) \geq \theta_H(d_H^{\hat{s}}-d_L^{\hat{s}})$ and $\frac{\theta_H}{\theta_L}(q_H^{\hat{s}}-q_L^{\hat{s}}) \leq \theta_H(d_H^{\hat{s}}-d_L^{\hat{s}})$. Since $(\rho_{np} + \theta_H\rho_p) < \frac{\theta_H}{\theta_L}$ these conditions can only be jointly satisfied if $q_H^{\hat{s}} \leq q_L^{\hat{s}}$ and $d_H^{\hat{s}} \leq d_L^{\hat{s}}$ for all $\hat{s} \in \{H, L\}$.

Step 2): For now, we will assume that (A2.6) is always satisfied. We will check at the very end that (A2.6) is indeed satisfied in our candidate optimum. Assume that (A2.1) is not binding. In this case the lender could reduce q_L^H by one unit and d_L^H by ρ_p units while increasing q_H^H by $\frac{(1-\alpha_H)}{\alpha_H}$ units and d_H^H by $\frac{\theta_L}{\theta_H}\frac{(1-\alpha_H)}{\alpha_H}\rho_p$ units. This change increases the borrower's expected utility while leaving (A2.5) and (A2.7) unaffected. Moreover, the proposed change relaxes condition (A2.2). This implies that (A2.1) must be binding for any optimal contract. Similarly, assume that (A2.3) is not binding. In this case the lender could reduce q_L^L by one unit and d_L^L by ρ_p units while increasing q_H^L by $\frac{(1-\alpha_L)}{\alpha_L}$ units and d_H^L by $\frac{\theta_L}{\theta_H}\frac{(1-\alpha_L)}{\alpha_L}\rho_p$ units. This change increases a lender's expected utility while leaving

(A2.7) unaffected. Moreover, the proposed change relaxes (A2.5) and (A2.4).

Step 3): Let us now reduce $q_L^{\hat{s}}$ by one unit and $d_L^{\hat{s}}$ by r_M units. This leaves conditions (A2.1) and (A2.3) unchanged and since $r_M = 1/\theta_L$ conditions (A2.5) and (A2.7) remain unaffected, too. So the change does not influence the borrower's expected utility and is feasible unless (A2.2) and (A2.4) become binding. This will be the case when $q_H^{\hat{s}} = q_L^{\hat{s}}$ and hence $d_H^{\hat{s}} = d_H^{\hat{s}}$. Moreover, in Section 1.2.2 we have shown that for any optimal contract that has $q_H^{\hat{s}} = q_L^{\hat{s}}$ and $d_H^{\hat{s}} = d_L^{\hat{s}}$ condition (A2.6) is satisfied. Moreover, if relationship lenders offer the optimal pooling contract, offering $r_M = 1/\theta_L$ is indeed a best response for uninformed lenders whenever $e > \tilde{e}$. So there always exists a pooling equilibrium in which borrowers with different types i but the same signal s receive the same conditions from their relationship lender.

Relationship lenders can always hand out less credit to low types and have them take out more credit on the spot market instead. In case of a market breakdown lenders on the spot market make zero profits when dealing with low types. Hence, the relationship lender can change the ex-ante contract in a fashion that leaves both, the utility of low types and his own profits for any given announcement $\hat{s}$ constant. So the reduction in $q_L^{\hat{s}}$ does neither affect the relationship lender's profits, nor his incentive to reveal his information truthfully. This implies that the relationship lender might as well offer the same amount of credit to high and low types. However, incentive compatibility on the side of the borrower implies that in this case the repayments have to be the same for both types, too, and we end up with pooling between different types that received the same signal s.

Chapter 2

Bonuses and Managerial Misbehaviour

In the wake of the recent financial crisis, excessive bonuses for bankers were frequently blamed as a source of irresponsible behaviour. In order to be awarded those tempting rewards, bankers were supposedly prepared to engage in behaviour that was not in the interest of their employers. In particular, bankers had incentives to take excessive risks, since they could reap the benefits in the case of success and were protected by limited liability in the case of failure. This line of argument suggests however that banks did collectively set sub-optimal incentive schemes, an idea which at least warrants some closer scrutiny.[1] In this chapter we show that high-powered incentives are in many cases not only robust to potential undesirable behaviour, but the cost of non-compliance may even increase the optimal bonus an agent is offered. Offering large bonuses hence may have been optimal even if banks were aware of their employees' ability to take up excessive risks.

Our finding that reducing pay-performance ratios may not be a suitable way to encourage compliance is consistent with the observation by Kose John and Yiming Qian (2003)

[1] An alternative justification for large bonuses would be that shareholders were confident of being bailed out in case risky investments turned bad. Given the large losses incurred by equity holders during the recent crisis however, this confidence seems at most only partially justified ex-post.

that controlling for leverage,[2] banks offer incentives that are not significantly different from the ones given in firms where misbehaviour is less costly. Comparisons of incentives in a number of different industries such as in Martin J. Conyon and Kevin J. Murphy (2000), Kevin J. Murphy (1999) and Xianming Zhou (2000) paint a similar picture: Pay-performance ratios do not seem to be noticeably lower in sectors where compliance is key, such as the Financial Services or Natural Resources industries. Furthermore, our model is in line with the finding by Fahlenbrach and Stulz (2011) that the size of cash bonuses a bank paid was not negatively correlated with performance during the recent financial crisis.

In this chapter we propose a standard moral-hazard model where the agent is risk-neutral but protected by limited liability and enlarge the agent's action space by assuming that he can not only choose how much effort to exert, but also whether (and how much) to engage in undesirable behaviour. This unwanted behaviour (or non-compliance) increases the contractible profit signal and hence the agent's variable compensation, but is nevertheless against the principal's interest. If the agent carries out an undesirable action, the principal finds out about this non-compliance with positive probability and is able to punish the manager.

Whenever bonuses are small, an increase in the bonus raises the level of misbehaviour, since any action that positively affects the profit signal now becomes more attractive. On the other hand, for large bonuses, the level of non-compliance is decreasing in the bonus. The higher the incentives for effort, the higher the bonus the agent is going to lose out on in case his misbehaviour is discovered, in which case he receives zero wage payments. Hence, with very high incentives he will be less prepared to jeopardise these expected earnings by taking undesirable actions and will be more likely to comply. Thus by offering large bonuses the principal exacerbates the maximal punishment that he can impose on a misbehaving agent. Given that misbehaviour is most pronounced for intermediate incentives, the principal optimally chooses "extreme" incentives. If effort is of

[2] Teresa A. John and Kose John (1993) offer a theoretical explanation why pay-performance ratios should depend on the debt ratio of a company.

little importance he will curtail incentives in order to reduce the level of misbehaviour. But if it is important to motivate effort, the principal will offer lavish bonuses in order to curb non-compliance. In both cases, the principal may find it optimal to complement bonuses with fixed wages that are paid regardless of the firm's profit as long as no evidence of misbehaviour surfaces in order to enhance compliance. Our model takes the somewhat extreme view that all wages the agent receives are given to ensure incentive compatibility. They are neither due to collusive practices in the determination of pay (Lucian A. Bebchuk and Jesse Fried, 2004), nor do they reflect scarcity of prospective employees (Marko Tervio, 2008; Xavier Gabaix and Augustin Landier, 2008). While this view is unlikely to fully describe the reality of executive compensation, nevertheless it offers some interesting insights. In particular, fixed wages can not only be used in order to satisfy participation constraints, but they may also be an additional disciplining device aimed against misbehaviour.

Our results imply that for top-management positions high incentives may in fact be a method to induce compliance, whereas lower ranks in a firm's hierarchy will receive very performance-inelastic pay in order to achieve the same goal. Clearly our analysis does not only apply to the remuneration of executives, but can equally explain the pay of portfolio managers, traders etc., who typically receive very high-powered incentives and are equally able to engage in undesirable behaviour, e.g. excessive risk taking that only becomes evident in adverse states of the world. Our model brings us one step closer towards understanding the strong monetary incentives in these jobs.

Undesirable behaviour not only plays a role in the banking industry, but is always an issue if agents have the possibility to game incentive schemes by engaging in actions that are harmful to their employer. Our model is sufficiently general to apply to a variety of situations and is in no way confined to incentive problems in the financial industry. Examples for undesirable behaviour in other industries may include illegal actions such as setting up a cartel or paying bribes in order to be awarded a contract. While sufficiently high expected fines guarantee that these actions are often not in the principal's interest, a manager may nevertheless be tempted to engage in such behaviour in order to be awarded a bonus.[3]

Managerial misbehaviour will often not only have negative consequences for the firm itself but will also impose negative externalities on society as a whole. In the final section of this chapter we consider the implications which our model has for a number of policies that a legislator might consider in order to increase the level of compliance within organisations. While shareholders will themselves typically give their managers some incentives to comply, a policy maker may want to reduce misbehaviour even further, and he is able to do so by different means. We conclude that caps on bonuses may be counter-productive. But even if they have positive effects, it is always more efficient to make shareholders liable for the misbehaviour of their managers.

In looking at a two-dimensional moral hazard model, our work is clearly related to the multi-tasking literature initiated by Holmström and Milgrom (1991). Yet, by assuming that the manager is protected by limited liability, we reach quite different conclusions. While in the traditional multi-tasking models the introduction of additional tasks typically reduces optimal incentives, in our setting the opposite can be true: Incentive problems in the second dimension are mitigated by *increasing* incentives in the first dimension. Precedents like the Enron case have shown that limited liability constraints are indeed an issue even for (usually rather wealthy) executives since courts are reluctant to enforce fines in excess of recent wage income.[4]

The idea that monetary incentives may trigger undesirable behaviour is not new and has for example been studied in the empirical literature on earnings management (e.g. Paul M. Healy, 1985; Beth J. Asch, 1990; Robert W. Holthausen, David F. Larcker and Richard G.

[3]Our only key assumption is that the agent doesn't fully internalize the negative consequences of his actions. There are two reasons for this: a) The agent's limited liability and b) The imperfect observability of undesirable actions. Imperfect observability may arise because often, the negative consequences of undesirable effort only emerge in the distant future. Alternatively, it may be impossible to condition the agent's remuneration on certain outcomes, e.g. a drop in a firm's reputation, due to verifiability constraints. In this case the principal can only impose punishments in cases where additional evidence of misbehaviour is found.

[4]See Lucian A. Bebchuk, Joseph E. Bachelder, Roel C. Campos, Byron S. Georgiou, Alan G. Hevesi, William Lerach, Robert Mendelsohn, Robert A.G. Monks, Toby Myerson, John F. Olson, Leo E. Strine and John C. Wilcox (2006).

Sloan, 1995; Paul Oyer, 1998; Ian Larkin, 2007) or in the context of the optimal scope of a firm (Paul Fischer and Steven Huddart, 2008). This literature does not however derive the optimal structure of management contracts. Another related strand of literature looks at the interplay between the incentives for effort, short-termism or risk-taking and a company's financial structure (e.g. Jeremy C. Stein, 1988; Jeremy C. Stein, 1989; John and John, 1993; von Thadden, 1995; Bruno Biais and Catherine Casamatta, 1999; Patrick Bolton, José Scheinkman and Wei Xiong, 2006). Moreover, a more recent literature on dynamic contracts has explored the question how deferred compensation (e.g. Alex Edmans, Xavier Gabaix, Tomasz Sadzik and Yuliy Sannikov, 2012; Gustavo Manso, 2011) or earnings- (rather than stock-) based compensation can mitigate short-termism (Effi Benmelech, Eugene Kandel and Pietro Veronesi, 2010). Finally, Giancarlo Spagnolo (2000; 2005) looks at the question whether or not incentive contracts can make collusion harder to sustain while abstracting from the effect a given incentive has on the agent's choice of effort. The most closely related work is by Inderst and Ottaviani (2009) who look at optimal contracts if sales agents must be induced to search for potential customers, but not to sell to unsuitable customers. When a sales agent is faced with an unsuitable customer, he is unable to earn any bonus by acting in the principal's interest and can *only* earn a bonus by misbehaving. This implies that unlike in our setting higher incentives will never have a disciplining effect and will always increase misbehaviour.

The rest of the chapter is organised as follows: Section 2.1 formulates the model. Section 2.2 analyses the agent's problem for a given contract and derives the key insights concerning the agent's behaviour. In Section 2.3 we explore the properties of an optimal contract. Sections 2.4 and 2.5 derive testable predictions and show that our results are robust to allowing for more general payment schemes. In Section 2.6 we explore the policy implications of our model. Section 2.7 concludes.

2.1 The Model

A risk-neutral principal employs a risk-neutral but wealth constrained agent to manage a firm. The firm can make either high profits $\overline{\pi}$ or low profits $\underline{\pi}$ with $\Delta = \overline{\pi} - \underline{\pi} > 0$. The agent's wealth is initially zero and has to be non-negative in all states of the world. He can exert unobservable effort which determines the probability $a \in [0, \overline{a})$ that high profits arise. We will denote the agent's effort cost for working in the firm by $C(a)$.

Furthermore, the agent has the possibility to unobservably increase the probability of high profits by $u \in [0, \overline{u}]$ at a private cost $K(u)$ if he engages in actions that are seen as undesirable by the principal. Since the overall probability of high profits can not exceed one we assume that $\overline{a} + \overline{u} \leq 1$. Misbehaviour imposes an expected, non-verifiable cost of $\tau(u) = \delta\gamma(u)$ on the principal where δ is some scalar, $\tau(0) = 0$ and $\tau'(u) \geq \Delta$ for all u. That is, the marginal cost of undesirable effort outweighs the benefit of an increase in the likelihood of high profits from the principal's point of view. With a small probability $p(u)$ the principal gains hard information that the agent has been engaging in undesirable behaviour and can punish him by reducing his wage payments to any level that does not violate the limited liability constraint. In what follows we impose the following assumptions:

Assumption 1.

$$i) C(0) = 0, C'(0) = 0, C''(a) > 0, \lim_{a \to \overline{a}} C'(a) = \infty$$
$$ii) K(0) = 0, K'(0) = 0, K''(u) > 0$$

Part $i)$ of the assumption says that the cost of effort is an increasing and convex function of a and ensures that the agent always finds it optimal to choose some $a < \overline{a}$. Similarly, part $ii)$ says that the cost of undesirable behaviour is increasing and convex in u.

Assumption 2.

$$p'(0) = 0, p''(u) > 0, p(\overline{u}) = 1$$

Assumption 2 guarantees that not only the explicit cost of undesirable behaviour but also the implicit cost, i.e. the risk of being caught, is convex. Moreover, the last part of the assumption makes sure that the agent will always find it optimal to choose an interior level of u. Note that we do not assume that $p(0) = 0$, i.e. there may be a positive probability that an agent who has been engaging in no misbehaviour at all is found guilty.

Assumption 3. *$C'(a)$ is convex and $\frac{C'''(a)}{[C''(a)]^2}C'(a) < f$ for some $f < 2$.*

This technical assumption ensures that the amount of effort an agent exerts is concave in the bonus he expects to get but is not too concave. Assumption 3 is for example satisfied by all power functions $C(a) = ka^r$ with $k \geq 0$ and $r \geq 2$. Moreover, log-concavity of the marginal cost of effort is sufficient for the second part of the assumption to hold.[5]

While the model is fairly general, it can in particular be applied to excessive risk-taking in the financial industry: If the principal is unable to monitor the investment decisions of his banker, he has to incentivize the agent via bonus contracts. The agent can exert effort looking for efficient projects in order to increase expected profits and hence the bonus that he can expect to earn. Additionally, he can spend some effort looking for excessively risky projects. With a large probability $(1 - p(u))$ those risky investments have consequences indistinguishable from standard projects. However, with a small probability $p(u)$ they trigger a crisis which imposes large costs on the principal. At the same time, a crisis generates hard evidence on the agent's misbehaviour and allows the principal to punish his employee. Note that the probability that the bank suffers from a crisis is endogenous and depends on the amount of risk-taking the agent has been engaging in.

2.2 The Agent's Decision Problem

Let us start by characterising the agent's decision problem. In general, an employment contract may specify different wage payments depending on whether misbehaviour has been detected and whether the firm has been making high profits. In order to simplify

[5] Log-concavity implies that $\frac{C'''(a)}{[C''(a)]^2}C'(a) \leq 1$.

our analysis we assume for now that the agent receives a payment only if no misbehaviour is observed and the firm makes high profits. This payment will be called the bonus b. We will relax the assumption that the agent is only paid in the "best" state of the world in Section 2.5 and show that our results are robust to allowing for more general compensation schemes. We will show that it is indeed always optimal to pay no wage in cases where the principal has observed non-compliance. Our assumption that the manager receives no wages in the case where the firm makes low profits allows us to concentrate on the main effects and will be relaxed in Section 2.5.

The utility an agent receives if offered a bonus of $b \geq 0$ is given by

$$U = (a+u)\Big(1 - p(u)\Big)b - C(a) - K(u).$$

We assume that the agent has an outside option of zero, so he will accept any contract he is offered and we can ignore his participation constraint. Any optimal choice of a and u will have to satisfy the following two first order conditions:

$$\frac{\partial U}{\partial a} = (1 - p(u))b - C'(a) = 0 \tag{2.1}$$

$$\frac{\partial U}{\partial u} = (1 - p(u))b - K'(u) - p'(u)b(a+u) = 0 \tag{2.2}$$

Given Assumptions 1 and 2 we can show that the optimum is always unique (see the appendix). Note that we have assumed the agent's effort cost to be additively separable in the two dimensions and hence there are no technological complementarities between the two tasks. Nevertheless, we see that the two dimensions are strongly intertwined: The principal is left with only one instrument, b, to encourage effort and discourage misbehaviour. Moreover, undesirable behaviour will itself reduce the probability with which a successful manager receives the bonus and will therefore erode incentives for effort as can be seen in equation (2.1). Effort, on the other hand, increases the expected bonus the agent loses out on in case misbehaviour is detected and will increase the level of compliance as determined by (2.2).

Any optimum is implicitly defined by

$$F \equiv (1-p(u))b - K'(u) - p'(u)b\Big(G\Big((1-p(u))b\Big) + u\Big) = 0 \qquad (2.3)$$

where the effort level is given by $a = G((1-p(u))b)$ and where $G \equiv C'^{-1}$ is a strictly increasing, concave function. In order to determine the overall effect an increase in the bonus b will have on the agent's choice of u we have to look at

$$\frac{du}{db} = -\frac{\frac{\partial F}{\partial b}}{\frac{\partial F}{\partial u}} = \frac{(1-p\,(u)) - p'(u)\Big((a+u) + (1-p(u))bG'\Big((1-p(u))b\Big)\Big)}{-\partial F/\partial u} \qquad (2.4)$$

with the denominator being positive by strict concavity of the agent's objective function in the optimum. It will be useful to state some basic properties of the agent's choice of a and u:

Lemma 1. *Both, effort and the probability of the firm making high profits, are increasing in the bonus b:* $\frac{da}{db} > 0$ *and* $\frac{d(a+u)}{db} > 0$.

Proof. See the appendix.

By equation (2.1) we know that the agent's effort is determined by the bonus he can expect to earn in case of high profits, which we will henceforth denote by $\beta = (1-p(u))b$. The first part of Lemma 1 corresponds to an upper bound on du/db and tells us that while an increase in b may lead to more undesirable behaviour and hence a larger $p(u)$, the bonus an agent can expect to earn in case of high profits will still be strictly increasing in b. The second part of Lemma 1 establishes a lower bound on du/db. Even if larger bonuses lead to less undesirable behaviour, the overall probability of the firm making high profits (which is given by $a+u$) is still increasing in the bonus.

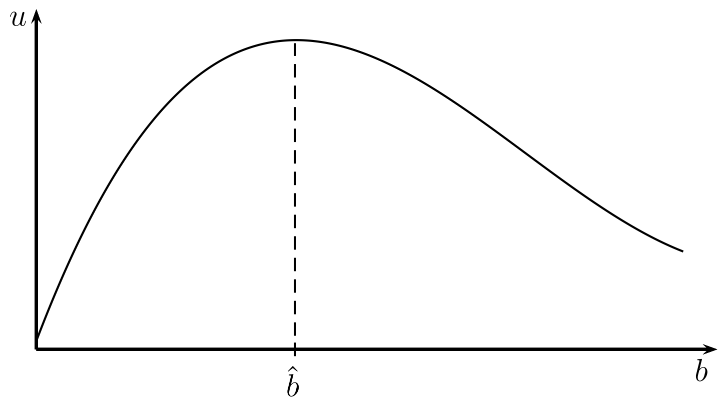

Figure 2.1: Misbehaviour as a Function of Bonus Payments.

Proposition 1. *There exists some strictly positive threshold $\hat{b}$ such that*

- *if $b < \hat{b}$ an increase in the bonus b will lead to more misbehaviour: $du/db > 0$,*
- *if $b > \hat{b}$ an increase in the bonus b will lead to less misbehaviour: $du/db < 0$.*

Proof. See the appendix.

Let us look in more detail at the numerator in equation (2.4). The first term captures the idea that an increase in b will raise the returns to undetected misbehaviour, which makes such actions more attractive. The second term corresponds to the fact that the relative harm of being caught is increasing in the bonus the agent is going to lose out on, which discourages misbehaviour. Moreover, not only is the bonus the agent misses out on in the case of detected misbehaviour increasing in b per se, but an increase in the bonus will also lead the agent to exert more effort, making the expected reward the agent jeopardises if he chooses not to comply even larger. Consider a situation in which $b = 0$ and the principal contemplates marginally increasing the bonus: In this case the second term vanishes, since there is no bonus the agent might miss out on. The marginal effect on the return to misbehaviour on the other hand is still strictly positive, which implies that misbehaviour is increasing in the size of incentives. Employees that hardly get any bonuses at all will not be disciplined by the prospect of losing them, which means that any small bonus will inevitably make misbehaviour more attractive. However, for very large bonuses this logic no longer applies. Now the agent cares about losing his high (expected) compensation and he is hence very reluctant to breach his fiduciary duties. Indeed, Assumption 3 guarantees that there always exists a threshold such that the level of misbehaviour decreases in b if and only if the bonus is larger than this threshold.

As a first approximation, let us consider how the numerator in equation (2.4) changes as we increase b and hold u constant. The marginal effect of an increase in b on the return to misbehaviour is constant. In contrast to this, the total wage payments an agent can expect to earn (or lose, in case he is caught misbehaving) are convex in b. So the second terms in the numerator of equation (2.4) will become increasingly large and the effect an increase in the bonus has on the level of non-compliance will eventually become negative.[6] The most intuitive way to think about the second effect is to say that by setting a higher bonus b the principal effectively relaxes the limited liability constraint of the agent. By using larger bonuses to augment the agent's expected wealth, the principal increases the maximal punishment that he can impose on a demonstrably misbehaving agent, which discourages misconduct.

Employees that receive low bonuses are very amenable to misbehaviour. An increase in their rewards runs the risk of encouraging misdeeds that increase the contractible profit signal. Highly incentivized executives and other high-ranked employees on the other hand can expect to earn extremely high bonuses even without misbehaving. This creates a strong concern for preserving those prospects and any policy that increases wage payments will consequently enhance compliance. Taking those two observations together, the function $u(b)$ that determines the level of misbehaviour follows an inverted-U shape.

Note that we can also express Proposition 1 in terms of the expected bonus β. By Lemma 1 any unit increase in β can be interpreted as a rise in b that is normalised as to result in a unit increase in the expected bonus. So there is a $\hat{\beta} = (1 - p(u))\hat{b}$ such that whenever $\beta < \hat{\beta}$ we have $du/d\beta > 0$ and vice versa.

[6] To see this note that Assumption 3 implies that $-G''(\beta)\beta < 2G'(\beta)$ for all β. Holding u constant, convexity of the expected wage payments an agent receives requires that $G''(\beta)\beta + 2G'(\beta) > 0$ for all $\beta = (1 - p(u))b$, which is indeed the case.

2.3 The Choice of a Bonus

The principal has two objectives that he pursues when setting a bonus: On the one hand he wants to give the agent incentives to exert effort; on the other hand he does not want him to engage in undesirable behaviour. When choosing the bonus b he will have to strike a balance between those two goals. For notational simplicity we assume that instead of choosing b, the principal (equivalently) chooses the expected bonus $\beta = \Big(1 - p(u)\Big)b$. The agent's choice of effort depends not only on b per se but also on the probability with which he believes he will lose the bonus due to misbehaviour. So the effect any change in b has on the manager's choice of effort will be amplified or dampened by the effect it has on undesirable behaviour u. Assuming that the principal chooses the effective bonus β directly allows us to ignore this issue and is without loss of generality.[7]

The principal's objective function is given by

$$\Pi\left(\beta, u(\beta)\right) = \Big(u + G(\beta)\Big)\Big(\overline{\pi} - \beta\Big) + \Big(1 - u - G(\beta)\Big)\underline{\pi} - \tau(u) \tag{2.5}$$

where u is a function of β and $a = G(\beta)$. From the principal's point of view there are two negative effects of undesirable behaviour: First, undesirable behaviour creates an efficiency loss of $\tau(u)$ which more than off-sets the positive effects undesirable behaviour has on firm profits. Secondly, undesirable behaviour allows the agent to appropriate additional wage payments. The latter effect explains why a principal wants to reduce misbehaviour even if it is costless from an efficiency point of view, i.e. if $\tau'(u) = \Delta$.

The necessary condition for an optimum with $\beta > 0$ is given by

$$\frac{d\Pi\left(\beta, u(\beta)\right)}{d\beta} = \frac{\partial \Pi}{\partial \beta} + \frac{\partial \Pi}{\partial u}\frac{du}{d\beta} = 0. \tag{2.6}$$

The first important observation is that the optimisation problem the shareholders face is not necessarily concave. The cost of misbehaviour makes very small and very large

[7] Lemma 1 implies that β is a strictly increasing function of b.

bonuses more attractive, since those bonuses guarantee high levels of compliance. So for large costs of undesirable behaviour, there will always exist multiple local maxima: (At least) one in which the principal pays small bonuses and $du/d\beta$ is positive, and (at least) one in which he offers lavish incentives and $du/d\beta$ is negative. Henceforth, we will just assume that the problem has a unique optimum, abstracting from the non-generic case where two local optima generate exactly the same level of profits.

Although it is interesting to look at the question of how the optimal bonus changes with the ease with which the agent can game his incentive scheme, we take as given that the agent has the opportunity to engage in undesirable behaviour at some cost. The question we will chiefly be interested in is a different one: How does the optimal bonus change with the damage an agent causes by misbehaving? Consider a salesman who harms a firm's reputation by misselling an excessively sophisticated product to a customer and a fund manager who invests in overly risky assets to increase his expected bonus. Even if the private cost of misbehaviour is the same for both the salesman and the fund manger, the potential damage caused by misconduct is presumably much larger in the second case and the principal will have much stronger incentives to reduce non-compliance.

Before we proceed, let us briefly recall that we have defined the cost of undesirable behaviour to the principal as $\tau(u) = \delta\gamma(u)$. So it seems natural to model a rise in the damage misbehaviour causes by an increase in the scaling parameter δ. Moreover, note that the value of (productive) effort is described by $\Delta = \overline{\pi} - \underline{\pi}$.

Proposition 2. *There exists a strictly positive threshold $\hat{\Delta}$ such that as long as $\beta > 0$ we have $\frac{d\beta}{d\delta} \gtrless 0$ if $\Delta \gtrless \hat{\Delta}$: Whenever productive effort is sufficiently important, a marginal increase in the cost of undesirable behaviour will increase the bonus a principal optimally offers. Conversely, whenever effort is not very important, a marginal increase in the cost of undesirable behaviour will reduce the bonus a principal optimally offers.*

Where $\hat{\Delta} = \sup\left\{\Delta \mid \hat{\beta} \geq \max\left\{\arg\max_{\beta} \Pi\right\}\right\}$ and where $\hat{\beta} = (1 - p(u))\hat{b}$ is defined by Proposition 1.[8]

Proof. See the appendix.

At the heart of this proposition lies a straightforward intuition: Since compliance is highest for very small or very large bonuses, the principal will generally set more "extreme" incentives as his concern for compliance increases. A principal that does not value productive effort very highly offers low incentives anyway and will depress bonuses further in order to enhance compliance. Firms that rely heavily on productive effort and offer incentives above $\hat{\beta}$ on the other hand can exploit the negative effects large bonuses have on the level of misbehaviour. Faced with an increase in the cost of non-compliance they will offer more lavish bonuses as to enhance compliance.

While Proposition 2 describes how the optimal bonus changes with a marginal change in the cost of misbehaviour, the same does not necessarily hold true for more radical changes in the damage a misbehaving agent causes. Assume that the cost of misbehaviour increases drastically. Even if a principal found it optimal to set high bonuses beforehand, after a large hike in the cost of non-compliance he may decide to scrap incentives altogether and offer very small bonuses, which results in negligible levels of misbehaviour. Since the principal's optimisation problem is not globally concave, any large change in the cost of misbehaviour may render a different local optimum more attractive and lead to a discontinuous change in the optimal bonus.

2.4 Industry Implications

Unfortunately, neither the importance of effort, nor the precise damage that a misbehaving agent causes in a particular firm are likely to be observable in reality. We would however expect the cost of misbehaviour to be industry specific, while even within an industry firms exhibit considerable heterogeneity with respect to the importance of CEO effort (e.g.

[8]Note that we do not assume that the principal's problem has a unique optimum at $\hat{\Delta}$.

because of different organisational choices). This allows us to characterise any industry by a distribution of incentive plans. We have seen that the harm done by non-compliance will result in more "extreme" incentives. Those firms that offer very low pay-performance ratios will depress incentives even further as the cost of undesirable behaviour increases, while companies with high-powered incentives will increase bonuses. Hence, the bonus distribution will be more spread out in an industry where misbehaviour is more costly than in one where it is comparatively harmless.

Consider an industry where the returns to effort Δ are distributed with strictly positive density in the interval $(0, \infty)$. Thus there is always a positive measure of firms which offer small bonuses and will react to an increase in the cost of misbehaviour by reducing incentives. But we will also have firms that offer substantial incentives and choose to raise bonuses if misbehaviour becomes increasingly costly. In order to ensure that even for arbitrarily large values of Δ misbehaviour is indeed not in the principal's interest let us redefine the cost of undesirable behaviour as $\hat{\tau}(u) = \Delta u + \delta\hat{\gamma}(u)$ where $\delta\hat{\gamma}'(u) > 0$. This specification keeps the *net* cost of misbehaviour constant for different values of Δ: In expectation the principal will lose any benefits that have accrued from misbehaviour plus some additional (net) cost of $\delta\hat{\gamma}(u)$.[9]

The fact that there are always some firms that strengthen incentives in response to an increase in δ allows us to derive clear-cut predictions concerning the interquantile ranges of the bonus distribution.[10] For the purpose of this proposition we do not assume that the optimal bonus is unique for all values of Δ. Instead, in order to guarantee that the bonus distribution is uniquely defined we simply assume that whenever the principal is indifferent he chooses the smallest optimal bonus.[11]

[9] While the specification of $\hat{\tau}(u)$ is mainly for technical convenience, it may represent a situation where the cost of misbehaviour consists of legal fines. Typically, such fines are set as to claim back any benefits the principal may have had from misbehaviour plus an additional deterrent $\delta\hat{\gamma}(u)$.

[10] The interquantile range denotes the difference in value between two given quantiles of a distribution.

[11] In this section we look at b instead of the expected bonus β since we believe b to be more easily observable in reality. We could equivalently look at β without changing any of our results.

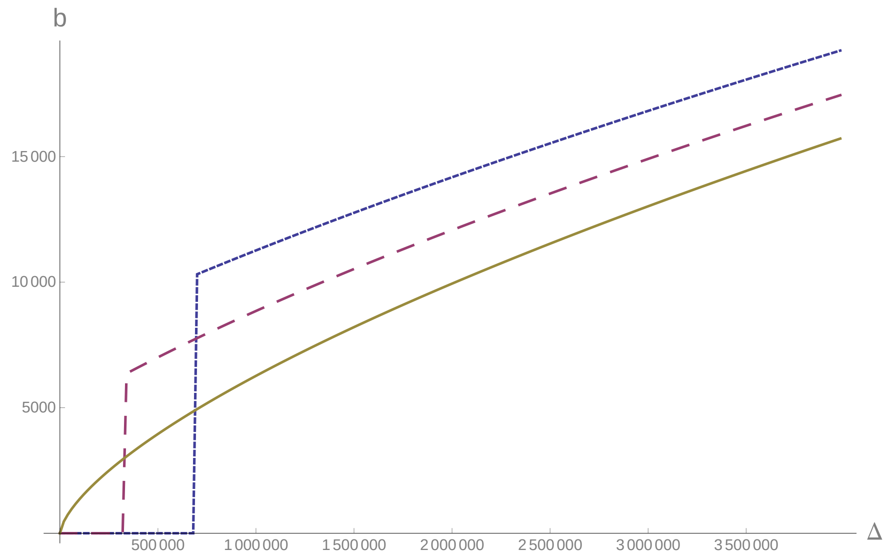

Figure 2.2: The Optimal Bonus as a Function of Δ.

Proposition 3. *Assume there is a continuum of firms which employ one manager each and returns to effort Δ are distributed in the interval $(0, \infty)$. Then the interquantile range $Q_{(1-\epsilon)} - Q_\epsilon$ of the bonus distribution $H(b)$ is increasing in δ for all small values of ϵ: The larger the harm caused by misbehaviour in a particular industry, the more spread out the bonus distribution will be.*

Proof. See the appendix.

If a rise in the marginal cost of non-compliance prompts principals to reduce bonuses that are small already and to increase those bonuses that are generous, this implies that the tails of the bonus distribution will grow further apart.

To illustrate said properties, we can simulate the optimal bonus b as a function of Δ for different costs of undesirable behaviour. Consider the example where an agent can exert effort or engage in risky investments which increase the company's chance of making high profits, but with a small probability of $p(u) = (u/0.25)^2$ a crisis occurs. In expectation, a crisis will not only destroy any value the investment generates in good states of the world, but it will also lead to a net loss of of T=US\$12m (dashed line) or US\$25m (dotted line). The cost of undesirable behaviour is hence given by $\hat{r}(u) = \Delta u + Tp(u)$. We take the private costs of effort and misbehaviour incurred by the agent to be $C(a) = a^2/(0.75 - a)$

	$\underline{\pi}$	$\overline{\pi}$
No Misbehaviour	w	$w + b$
Misbehaviour	w_p	$w_p + b_p$

Figure 2.3: General Contracts.

and $K(u) = u^2/(1-u)$. If we compare the optimal bonuses to a case where misbehaviour does not entail any efficiency loss (solid line), we see that bonuses get more spread out as the cost of misbehaviour increases. It should be noted that in this example, the firm's profit is not concave in the bonus for most values of Δ, which explains why the optimal bonus b is discontinuous in Δ.

Proposition 3 relies on the fact that distorting incentives is always the marginal instrument to enhance compliance. We will qualify this view in the next section where we allow for more general contracts. However, if a principal is constrained in increasing fixed wages, any increase in compliance will still (at least partially) be achieved by adjusting the bonus. Section 2.6 discusses in more detail why we might expect such constraints to play a role in reality.

2.5 General Contracts

So far we have assumed that the manager only receives a positive wage payment in case the firm makes high profits and no undesirable behaviour is detected. However, shareholders can of course decide to employ more sophisticated compensation schemes. In our model the most general contract a principal can offer is characterised by the tuple (w, b, w_p, b_p) where the agent gets a base salary w regardless of profits whenever no evidence of non-compliance is found and he receives an additional bonus b in the case of high profits. Similarly, w_p and b_p denote the respective payments in case misbehaviour is detected.

We will now relax the assumption that a manager is only paid in one state of the world and show that it is indeed never optimal to pay positive wages in the case of detected misbehaviour. While it may be optimal to pay the agent in the case of low profits, the main insights that we have obtained so far still apply if we allow for more general contracts.

Optimal Punishments

Let us start by considering the optimal wages in case the principal receives evidence on misbehaviour.

Lemma 2. *Any optimal contract that implements some $a > 0$ will pay no wages if undesirable behaviour is detected: $w_p = b_p = 0$. Furthermore, the wage payment in case of high profits will never be smaller than the wage payment in case of low profits: $b \geq 0$.*

Proof. See the appendix.

Any optimal contract will pay no bonus in case of observed misconduct. The intuition for this result is straightforward: Reducing b_p will always reduce the level of undesirable behaviour the agent chooses, which is beneficial. At the same time a decrease in b_p will lessen the incentives for effort a. However, this effect can be offset by adjusting b in a way that leaves the expected wage payment *given high profits* constant, which leaves us only with the negative impact on u. Also, $w_p > 0$ is never part of an optimum, since it incentivizes misbehaviour without having any effect on the choice of effort a. Finally, it is never optimal to punish the agent in the case of high profits: The same a and u could also be obtained by offering the agent no wages at all, strictly reducing the expected wage cost.

Fixed Wages

We have seen that the principal may decide to pay the agent large bonuses in order to increase the wage payments the agent might lose out on if he misbehaves. Alternatively, the principal could achieve the same by offering a fixed wage component w which will be paid out whenever no undesirable behaviour is observed, irrespective of the firm's profits. Paying the agent only in the case of high profits has the advantage of motivating effort and increasing the expected punishment for misbehaviour at the same time; yet it also increases the returns to undetected misbehaviour. This explains why it may be optimal to pay fixed wages in order to discourage non-compliance. In fact, whenever non-compliance is very costly, the principal will decide to make use of this alternative way of discouraging misbehaviour, since the positive effect of a larger bonus on the agent's effort choice becomes increasingly small.[12] The following proposition characterises the relationship between w and β when the principal decides to implement a given compliance level u:[13]

Proposition 4. *Assume that $w > 0$. Then the principal chooses β such that*

$$G'\Big(\beta\Big)\Delta = \frac{(1-p(u))}{p'(u)}. \tag{2.7}$$

Proof. See the appendix.

Let us compare the optimal bonus as characterised by Proposition 4 with the bonus a principal would choose if the agent did not have the possibility to misbehave. In that case, the optimal incentive $\tilde{\beta}$ would be characterised by $G'(\tilde{\beta})\Delta = \tilde{\beta}G'(\tilde{\beta}) + G(\tilde{\beta})$ where the left hand side represents the benefit of increased effort. The right hand side represents the marginal increment in the agent's compensation and is strictly increasing in $\tilde{\beta}$. Whenever Δ is large, $\tilde{\beta}$ is large and the right hand side of equation (2.7) will be smaller than $\tilde{\beta}G'(\tilde{\beta}) + G(\tilde{\beta})$. In this case the principal optimally chooses a bonus above $\tilde{\beta}$. If, on the other hand, Δ is small then the optimal bonus lies below $\tilde{\beta}$.

So even if we allow for fixed wages as a way to encourage compliance, the effects we have shown so far will still be present. Since distorting the bonus is costless at the margin, the principal will always choose to do so. Adjusting β to increase compliance is not only a last resort if, for some exogenous reason, no fixed wages can be paid, but it is indeed

[12] Similar reasoning applies if the principal contemplates lowering the bonus to increase compliance. While it is now possible to reduce misbehaviour by paying *lower* rewards, reducing bonuses is costly since the agent chooses less effort. Due to the concavity of $G(\beta)$ this effect gets larger the smaller the bonus.

[13] In this section we treat u as a constant since we are concerned with the question how a principal optimally implements a given u. Any change in β will be accompanied by a change in w that leaves the incentives for misbehaviour unchanged.

part of any optimal contract. In particular, as long as effort is sufficiently important the principal will still set generous bonuses to reduce misbehaviour.

2.6 Policy Implications

In many cases, undesirable behaviour has negative externalities on society as a whole. To cite our motivating example, excessive risk taking may create systemic risk and require public bail-outs. The same holds true for many other forms of misbehaviour: Cartel agreements reduce consumer surplus, bribes may undermine the rule of law etc. A natural question to ask is how a social planner may want to discourage such behaviour. In order to answer this question, we use our previous results to compare three policy instruments: The first one is a legal cap on bonuses. Such a measure is vividly discussed and therefore merits some closer theoretical examination. An instrument that is already applied in many areas is the imposition of pecuniary fines on firms if evidence of their employees' misbehaviour surfaces. Finally, we discuss whether a policy maker who is free to impose arbitrarily large punishments on misbehaving agents should always choose to do so.

In the following, we will compare the different policy instruments under the assumption that while bonuses can be freely chosen within the legal limits, the principal is constrained in increasing fixed wages. This assumption seems to be warranted given what many have called the "outrage constraint", i.e. the fact that shareholders are unlikely to accept very high levels of fixed wages, which are typically more visible ex-ante than variable compensation. Alternatively, we can think of the constraint on fixed wages as a result of current US tax legislation which treats expenditure on fixed wages unfavourably once it exceeds a certain threshold.

Caps on Bonuses

The previous discussion shows that the negative consequences of legal restrictions on the size of bonus payments are potentially twofold: Besides the obvious effect of reducing the incentives for effort, such a policy may even have adverse effects on compliance and encourage misbehaviour.

Corollary 1. *Assume that $\Delta > \hat{\Delta}$ and absent any regulation, the principal sets an expected bonus of β^*. Then a legal cap $\overline{\beta} < \beta^*$ on bonuses will $i)$ increase misbehaviour and $ii)$ decrease effort as long as $\beta^* - \overline{\beta} \leq \epsilon$ for all small ϵ.*

Corollary 1 tells us that what would seem to be a cautious regulation may in fact be very harmful. By imposing caps on bonuses that are close to the level of bonuses paid in an unregulated labour market, we may destroy incentives for effort *and* increase managerial misbehaviour. This does not however hold for large interventions: If the legal maximum on bonus payments is very small, this will have positive effects on compliance since any cap that is sufficiently close to zero will result in negligible levels of misbehaviour. In fact, even a less stringent cap on bonuses with $\overline{\beta} > \hat{\beta}$ can potentially increase compliance: Although setting a bonus of $\beta = \overline{\beta}$ would result in more misbehaviour, the principal may now find it optimal to set a much smaller bonus that results in less non-compliance. This observation implies that even if caps on bonuses are non-binding in equilibrium, they may nevertheless be effective since they induce a shift to a different local optimum. However, such interventions erode incentives for managers to work hard and may not increase social welfare.

Corporate Liability

So far we have taken the cost of undesirable behaviour for the principal to be exogenously determined, although in many cases this cost comprises legal fines. This suggests that a policy maker who wants the principal to discourage undesirable behaviour may choose to increase these fines. The policy maker has access to a monitoring technology $\mathcal{P}(u)$ with $\mathcal{P}'(u) \geq 0$ and can impose some punishment $\mathcal{T}$ on the principal if he observes illegal behaviour. For simplicity we assume that the states of the world in which the policy maker receives evidence of misbehaviour are a subset of the states in which the principal does so, which implies that $\mathcal{P}(u) \leq p(u)$.

Corollary 2. *Punishing the principal for observed misbehaviour always results in weakly higher social welfare than a cap on bonuses that implements the same u.*

Where social welfare is defined as firm profits before wage payments and fines minus any costs incurred by the agent or the general public. So a policy maker can not do worse by imposing fines on the principal rather than by capping bonuses. Furthermore, a marginal increase in the corporate fine will unambiguously increase social welfare whenever $\Delta > \hat{\Delta}$.

From the principal's point of view, an increase in the corporate fine $\mathcal{T}$ is simply a rise in the marginal cost of misbehaviour. We have already seen that whenever $\Delta > \hat{\Delta}$ the principal will respond to a marginal increase in corporate liability by increasing incentives. Since the principal will always implement effort that is too low from the point of view of a social planner, this is good news.[14] By punishing the principal, the policy maker not only reduces unwanted behaviour, but he also reduces the distortions created by the non-observability of effort. While the reverse is true if $\Delta < \hat{\Delta}$, a policy of punishing the principal is still weakly better in terms of social surplus than a cap on bonuses that implements the same u. A principal that is left to his own devices to achieve a given reduction in u will always do so by choosing weakly larger fixed wages and weakly larger bonuses, which is socially beneficial.

Personal Liability

Up to now, we have not considered the possibility that a policy maker might decide to punish the agent directly in the case of observed misbehaviour. Indeed, such punishments will usually be ruled out by the fact that the principal already chooses to punish the agent as fiercely as possible in case that undesirable behaviour is detected. So there is only a role for a policy maker in punishing the agent if the legislator disposes of additional options to penalise the manager, e.g. by imposing prison sentences. But even if we allow for the policy maker to impose arbitrarily harsh punishments on the agent, it is unclear

[14] As usual, the first-best effort level would be implemented if the agent would reap the full social benefits from effort, i.e. $\beta = \Delta$. Instead of choosing $\beta \geq \Delta$, the principal can always do strictly better by setting $\beta = 0$.

whether a welfare-maximising policy maker will choose to do so. While extremely harsh punishments can ensure that the agent chooses $u = 0$, they also crowd out any compliance-enhancing incentives the principal might otherwise set. In particular, the principal may decide to pay smaller bonuses, which is socially harmful since it aggravates the efficiency loss that is due to the non-observability of effort. Put differently, the policy maker has some incentive to preserve the moral hazard problem in the second dimension in order to reduce the distortions in the first dimension.

2.7 Conclusion

Our model shows that while bonus schemes generally open up ways for a manager to game them, very large bonus payments may discourage this kind of misbehaviour. When applied to the financial industry, this implies that large bonuses should not be mistaken for conclusive evidence that "too big to fail" created significant moral hazard problems and motivated banks to readily accept excessive risk taking by their employees. Nor are high bonuses necessarily a feature of sub-optimal contracts. According to our model, offering large bonuses may in fact have been an optimal strategy to limit risky behaviour - even though such conduct could apparently not be prevented entirely.

Moreover, the finding that there does not exist a monotonic relationship between bonuses and the incentives for misbehaviour is consistent with the observation by Fahlenbrach and Stulz (2011) that the size of previous cash bonuses paid by a bank did not correlate with bad performance during the recent crisis.

Finally, our results shed some light on the proposal to legally restrict the size of bonuses. We have shown that this may have counterproductive effects and may reduce compliance, while at the same time diluting incentives for managers to work hard. In order to maximise social surplus, it is always more attractive to give the principal financial incentives to increase compliance. This may even induce the principal to implement higher levels of effort, which reduces the distortions that are caused by the non-observability of effort.

A3 Mathematical Appendix

Proof of Existence and Uniqueness. Let us show that there will always be a unique optimum to the agent's problem and that any local maximum is a global one. In this proof we allow for the possibility that in addition to a bonus b the principal chooses to pay some fixed wage w irrespective of profits whenever no misbehaviour is observed. We only consider situations where $b > 0$: The principal will never choose any $b < 0$ as we will show in Lemma 2. Moreover, if $b = 0$, the unique optimum has $a, u = 0$. Any optimum satisfies $(1 - p(u))b - C'(a) = 0$ and there is a unique optimal a for any choice of u. This allows us to look at a one-dimensional optimisation problem where the agent chooses u and $a(u)$ is given by the above first order condition. Since the agent's utility is continuous on the closed interval $[0, \overline{u}]$, a maximum always exists. Moreover, we can easily check that any optimum must be interior. Now consider the largest (potentially locally) optimal $\hat{u} = \max\{\arg\max_u \{(a(u) + u)(1 - p(u))b - C(a(u)) - K(u) + w(1 - p(u))\}\}$. By the necessary condition (2.2) we must have $(1 - p(\hat{u}))b - bp'(\hat{u})a(\hat{u}) > 0$ at the optimum.

Let us now show that for a given b the optimum is unique on the interval $[0, \hat{u}]$. In order to do so, we will need to find a lower bound for $a'(u)$. First, note that $G \equiv C'^{-1}$ is a concave function, since we can use the inverse function rule to show that $G''(C'(a)) = -\frac{C'''(a)}{[C''(a)]^3}$ is negative for all a. As $a(u) = G((1 - p(u))b)$ this implies that $a(u)$ is concave in u. So it suffices to look at $a'(\hat{u}) = -bp(\hat{u})G'((1 - p(\hat{u}))b)$. Using the fact that $(1 - p(\hat{u}))b - bp'(\hat{u})a(\hat{u}) > 0$ and that $G((1-p(u))b) \geq G'((1-p(u))b)(1-p(u))b$ by concavity of G we get $1 > p'(\hat{u})bG'((1 - p(\hat{u}))b)$. It follows that $a'(u) > -1$ for all $u \in [0, \hat{u}]$

Since the agent's utility is given by $U(u) = (a(u) + u)(1 - p(u))b - C(a(u)) - K(u) + w(1 - p(u))$ we get $U''(u) = -p''(u)b(a(u) + u) - (2 + a'(u))\,p'(u)b - K''(u) - p''(u)w$, which is negative for all $u \leq \hat{u}$. So $U(u)$ is strictly concave over the interval $[0, \hat{u}]$ and the optimum is unique. □

For notational convenience we will henceforth write $\beta = (1 - p(u))b$ wherever possible. Moreover, we will no longer stress that $p = p(u)$ in the interest of brevity.

Proof of Lemma 1. First, let us show that $\frac{da}{db} > 0$ which is clearly the case if $b = 0$. If $b > 0$ the condition is equivalent to $\frac{du}{db} < \frac{(1-p)}{bp'}$ since $\frac{da}{db} = (1-p)G'(\beta) - bp'G'(\beta)\frac{du}{db}$. Again, we allow for a positive fixed wage w that is paid out whenever no misbehaviour is detected, even in case of low profits. We want to show that

$$\frac{du}{db} = \frac{(1-p) - p'(G(\beta)+u) - p'\beta G'(\beta)}{2bp' + b(G(\beta)+u)p'' + wp'' + K''(u) - [bp']^2 G'(\beta)} < \frac{(1-p)}{bp'}.$$

A sufficient condition for this inequality to hold is that

$$\frac{(1-p) - p'\beta G'(\beta)}{2bp' + b(G(\beta)+u)p'' + wp'' + K''(u) - [p'b]^2 G'(\beta)} < \frac{(1-p)}{bp'}$$
$$\Leftrightarrow \quad 0 < bp' + K''(u) + p''b(G(\beta)+u) + wp''$$

which is true by Assumptions 1 and 2.

Now we need to show that $\frac{d(a+u)}{db} > 0$, which is equivalent to $(1-p)G'(\beta) + (1 - bp'G'(\beta))\frac{du}{db} > 0$. Again, it is sufficient to consider situations where $b > 0$. From the proof of uniqueness we know that $1 - bp'G'(\beta) > 0$, so the condition can only be violated if $\frac{du}{db} < 0$. Let us hence look for a lower bound for $\frac{du}{db}$:

$$\begin{aligned}\frac{du}{db} =& \frac{(1-p) - p'(G(\beta)+u) - p'\beta G'(\beta)}{2bp' + b(G(\beta)+u)p'' + wp'' + K''(u) - [bp']^2 G'(\beta)} > \\ & -\frac{p'\beta G'(\beta)}{2bp' + b(G(\beta)+u)p'' + wp'' + K''(u) - [bp']^2 G'(\beta)} > \\ & -\frac{(1-p)G'(\beta)}{1 - bp'G'(\beta)}\end{aligned} \tag{A3.1}$$

where the first inequality follows from the fact that $(1-p) > p'(G(\beta)+\beta)$ by equation (2.2) and the second inequality is implied by $1 - bp'G'(\beta) > 0$. Plugging (A3.1) into our initial condition shows that indeed $\frac{d(a+u)}{db} > 0$. □

Proof of Proposition 1. We want to show that there is a bonus $\hat{b}$ such that $b > \hat{b} \Leftrightarrow \frac{du}{db} < 0$ and $b < \hat{b} \Leftrightarrow \frac{du}{db} > 0$. Again, we will allow for the possibility that the principal pays some positive wage $w \geq 0$ irrespective of profits. It can easily be seen that $\frac{du}{db}|_{b=0} > 0$.

Now, let us consider the case where $b > 0$ and show that $\frac{du}{db}$ will become negative for very large bonuses. First, note that $0 < u < \overline{u}$ and $0 < \lim_{b\to\infty} u < \overline{u}$. Since the denominator in (2.4) will always be positive, we just have to show that the numerator will eventually turn negative. This in turn is equivalent to showing that $b^{(1-\alpha)}\frac{\partial F}{\partial b}$ will be strictly negative in the limit for some α. Using equation (2.2) we get that for all $b > 0$

$$\frac{\partial F}{\partial b} = \frac{K'(u) + p'w}{b} - p'\beta G'(\beta).$$

Multiplying both sides by $b^{(1-\alpha)}$ and taking the limit we get

$$\lim_{b\to\infty}\left(b^{(1-\alpha)}\frac{\partial F}{\partial b}\right) = \lim_{b\to\infty}\frac{K'(u) + p'w}{b^{\alpha}} - \lim_{b\to\infty}\frac{p'}{(1-p)^{1-\alpha}}\lim_{b\to\infty}\beta^{(2-\alpha)}G'(\beta).$$

We can see that for $\alpha > 0$ the first term is zero and that $\lim_{b\to\infty}(p'/(1-p)^{1-\alpha})$ is strictly positive. Finally, by Assumption 3 we can show that $\beta^{(2-\alpha)}G'(\beta)$ is increasing in b for $2 > 2-\alpha > f$. To see this note that Assumption 3 implies that $-G''(\beta)\beta < 2G'(\beta)$. So $\lim_{b\to\infty}\beta^{(2-\alpha)}G'(\beta)$ must be strictly positive and we get $\lim_{b\to\infty} b^{(1-\alpha)}\frac{\partial F}{\partial b} < 0$ for all $2 - f > \alpha > 0$. This means that for all sufficiently large values of b we have $\frac{du}{db} < 0$.

By continuity of $\frac{du}{db}$ we know that $\frac{du}{db}$ has at least one root. In order to show that it has exactly one root, we are now going to show that $\frac{du}{db}$ is strictly decreasing in b whenever $\frac{du}{db} = 0$, which is sufficient since $\frac{du}{db}$ is differentiable. First, we can show that at any point where $\frac{du}{db} = 0$ we have

$$\frac{d^2u}{db^2} = -\frac{\frac{d\left(\frac{\partial F}{\partial b}\right)}{db}}{\frac{\partial F}{\partial u}}.$$

Since $\frac{\partial F}{\partial u}$ is always negative we just need to show that $\frac{d\left(\frac{\partial F}{\partial b}\right)}{db} < 0$:

$$\frac{d\left(\frac{\partial F}{\partial b}\right)}{db} = -p'(1-p)G'(\beta)\left(2 + \frac{\beta G''(\beta)}{G'(\beta)}\right)$$

which is strictly negative by Assumption 3. So $\frac{du}{db} = 0$ implies that $\frac{d^2u}{db^2} < 0$. This concludes the proof. □

Proof of Proposition 2. By assumption, the optimum of the principal's problem is unique. So a marginal change in the cost undesirable behaviour imposes on the principal will never lead to a discontinuous change in the optimal bonus and we can restrict attention to the neighbourhood around β where the principal's objective function is strictly concave. By the implicit function theorem the change in the bonus a principal chooses to offer in response to a marginal increase in the cost of misbehaviour is given by

$$\frac{d\beta}{d\delta} = \frac{\gamma'(u)}{\Theta}\frac{du}{d\beta}$$

where $\Theta = \frac{d^2\Pi}{d\beta^2} < 0$ by local concavity.

So $\frac{d\beta}{d\delta}$ takes the sign of $-\frac{du}{d\beta}$ and we have $\frac{d\beta}{d\delta} \gtrless 0 \Leftrightarrow -\frac{du}{d\beta} \gtrless 0 \Leftrightarrow \beta \gtrless \hat{\beta}$. Using the second part of Lemma 1 we can show that for any $\Delta' > \Delta$ *any* optimal β is strictly larger than *any* β that is optimal for Δ. This allows us to express the relation above as $\frac{d\beta}{d\delta} \gtrless 0$ if $\Delta \gtrless \hat{\Delta}$ where $\hat{\Delta}$ is defined by $\hat{\Delta} = \sup\left\{\Delta \mid \hat{\beta} \geq \max\left\{\arg\max_{\beta} \Pi\right\}\right\}$. Note that we have allowed for the fact that for some values of Δ the optimal incentive β may not be unique.
□

Proof of Proposition 3. Analogous to Proposition 2 we can show that for any given δ there exists some strictly positive threshold $\hat{\Delta}$ such that whenever $\Delta < \hat{\Delta}$ any optimal incentive must be smaller than $\hat{b}$ and vice versa. So the bonus distribution will always have strictly positive mass on both sides of the threshold $\hat{b}$.

For two bonus distributions that result from different costs of misbehaviour, δ and $\hat{\delta}$, we can always find some $\tilde{\epsilon} \in (0, 0.5)$ such that for any $\epsilon \leq \tilde{\epsilon}$ the ϵ-quantile will have $b < \hat{b}$ and the $(1-\epsilon)$-quantile will have $b > \hat{b}$ for both, δ and $\hat{\delta}$.

An increase in the marginal cost of misbehaviour from δ to $\hat{\delta}$ will leave the identity of the firm offering the $(1-\epsilon)$-quantile bonus unchanged. Now consider the problem faced by that firm: Any bonus that is optimal for δ must satisfy $\frac{\partial \Pi}{\partial \beta} + \frac{\partial \Pi}{\partial u}\frac{du}{d\beta} = 0$. Clearly, this bonus can no longer be optimal for $\hat{\delta} > \delta$ since it would now pay to marginally increase the bonus. Moreover, no other bonus in $[\hat{b}, b)$ can be optimal for $\hat{\delta}$ since those bonuses are associated with a strictly larger level of misbehaviour. Finally, by design the optimal bonus can not be less than $\hat{b}$. So the $(1-\epsilon)$-quantile bonus must be strictly increasing in δ. By a similar argument we can show that the ϵ-quantile bonus is weakly decreasing in δ, which concludes the proof. □

Proof of Lemma 2. Since we can't make sure that $b_p, b \geq 0$, simply reducing the fix wage components to zero may not be possible without violating limited liability constraints. For this proof it will hence be useful to redefine the wage that is paid in state $[i, j]$ as $w_{i,j}$ where $i \in \{h, l\}$ and $j \in \{p, n\}$ denote whether high profits (h) have been made or not (l) and whether misbehaviour has been detected (p) or not (n) . The utility of the agent is then given by

$$\begin{aligned}
U =& (a+u)\Big((1-p)w_{h,n} + pw_{h,p}\Big) + (1-a-u)\Big((1-p)w_{l,n} + pw_{l,p}\Big) - C(a) - K(u) \\
\frac{\partial U}{\partial a} =& (1-p)w_{h,n} + pw_{h,p} - (1-p)w_{l,n} - pw_{l,p} - C'(a) = 0 \\
\frac{\partial U}{\partial u} =& (1-p)w_{h,n} + pw_{h,p} - (1-p)w_{l,n} - pw_{l,p} - K'(u) - p'\Big((a+u)(w_{h,n} - w_{h,p}) \\
& + (1-a-u)(w_{l,n} - w_{l,p})\Big) = 0
\end{aligned}$$

In the general case, the expected bonus β is given by $\beta = (1-p)w_{h,n} + pw_{h,p} - (1-p)w_{l,n} - pw_{l,p}$. It can never be optimal to have $\beta < 0$ since the same a and a weakly lower u can be implemented by offering a contract $(0, 0, 0, 0)$ that has a strictly lower wage cost.

Suppose a contract $(w_{l,n}, w_{h,n}, w_{l,p}, w_{h,p})$ implements some $a > 0$ and has $w_{l,p} > 0$ or $w_{h,p} > 0$. Instead, we can choose a contract $(\hat{w}_{l,n}, \hat{w}_{h,n}, 0, 0)$ that has $\hat{w}_{l,n} = w_{l,n} + \frac{p}{1-p}w_{l,p}$ and

$\hat{w}_{h,n} = w_{h,n} + \frac{p}{1-p} w_{h,p}$ and would implement the same a if the agent were to choose the same level of u (which he is not). Now consider a contract $(\tilde{w}_{l,n}, \tilde{w}_{h,n}, 0, 0)$ that does indeed have the same expected wage payments as the initial contract conditional on profits being either high or low. Clearly, this contract implements the same a. Furthermore, assume that it has a probability of detecting misbehaviour $\tilde{p} \geq p$. In this case $(1-\tilde{p})\tilde{w}_{h,n} = (1-p)\hat{w}_{h,n}$ and $(1-\tilde{p})\tilde{w}_{l,n} = (1-p)\hat{w}_{l,n}$ implies that $\tilde{w}_{h,n} \geq \hat{w}_{h,n}$ and $\tilde{w}_{l,n} \geq \hat{w}_{l,n}$. However, this contradicts the assumption that $\tilde{p} \geq p$. So there exists a contract $(\tilde{w}_{l,n}, \tilde{w}_{h,n}, 0, 0)$ that has *i)* the same a, *ii)* a strictly lower u and *iii)* the same expected wage payments conditional on the realisation of profits. This leads us to conclude that any contract involving $w_{l,p} > 0$ or $w_{h,p} > 0$ is dominated. The assumption that $a > 0$ is only needed to ensure that any contract with $w_{l,p} > 0$ or $w_{h,p} > 0$ results in the agent choosing some $u > 0$. □

Proof of Proposition 4. The principal can either adjust the bonus b or increase fixed wages w in order to increase compliance. Allowing for $w > 0$ the agent's optimal choice of u satisfies

$$\beta - K'(u) - p'b\left(G(\beta) + u\right) - p'w = 0.$$

So we can show the marginal rate of substitution between b and w which leaves u constant to be

$$\frac{dw}{db} = \frac{(1-p)}{p'} - (G(\beta) + u) - \beta G'(\beta).$$

A contract can only be optimal if for any given u the principal can not increase profits by choosing a different combination of b and w that implements this u. If $w > 0$ (which can only be the optimal if $b > 0$) this implies that in an optimum we must have

$$
\begin{aligned}
& \frac{\partial \Pi}{\partial b} + \frac{\partial \Pi}{\partial w}\frac{dw}{db} = 0 \\
\Leftrightarrow \quad & (1-p)G'(\beta)\left(\overline{\pi} - \beta - \underline{\pi}\right) - (1-p)\left(G(\beta)+u\right) - (1-p)\frac{dw}{db} = 0 \\
\Leftrightarrow \quad & G'(\beta)\Delta = \frac{(1-p)}{p'}.
\end{aligned}
$$

□

Chapter 3

Optimal Tolerance for Failure*

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 manager 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

*This chapter is based on joint work with Piers Trepper.

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. In particular, this is 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 (Ivan E. Brick, Oded Palmon and John K. Wald, 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 facing 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".*

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 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 unfavourable 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 - 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 does turn 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 unsuccessful agent 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. So 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 (Bengt Holmström and Paul 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 Stephen E. Spear and Cheng 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, Bruno Biais, Thomas Mariotti, Jean-Charles Rochet and Stéphane Villeneuve (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 strongly 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 favourably 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, Xuan Tian and Tracy Y. Wang (2012) show empirically that start-ups financed by more failure-tolerant venture capital firms are more innovative. Augustin 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, Sanford J. Grossman and Oliver D. Hart (1983) emphasise 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.[1] Moreover, the responsiveness of CEO turnover to changes in performance does not seem to be systematically higher for firms with good corporate governance. Warner, Watts and Wruck (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 Dirk Jenter and Katharina 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".[2] 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 3.1 we propose our main model and characterise the optimal contract in a static, one-period setting. In Section 3.2 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 3.3 we turn to an example where the effort choice is binary. This allows us to characterise the optimal policies more closely. The main insights from the general case carry over to the

[1] See, e.g., Jerold B. Warner, Ross L. Watts and Karen H. Wruck (1988), David J. Denis, Diane K. Denis and Atulya Sarin (1997), Charles J. Hadlock and Gerald B. Lumer (1997), Murphy (1999), Mark R. Huson, Robert Parrino and Laura T. Starks (2001) and Steven N. Kaplan and Bernadette A. Minton (2006).

[2] 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.

example and we can additionally derive necessary and sufficient conditions for optimality of different employment regimes. Section 3.4 concludes.

3.1 The Model

We consider the 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 conditioned 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 is 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.[3] More generally, it is satisfied whenever the agent's coefficient of absolute

[3]Applying the inverse function theorem yields $h''' = 1/(u')^3(2(A(W))^2 + A'(W))$ where $A(W) =$

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 Hector Chade and Virginia N. Vera de Serio (2011).

3.1.1 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 maximise his expected utility $U(e_1|w_H, w_L) = (\theta + e_1)u(W + w_H) + (1 - \theta - e_1)u(W + w_L) - C(e_1)$. Therefore, 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 minimises 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 whenever the principal offers a contract that implements some $e_1 > 0$. Since the principal only chooses two wages, the wage scheme is fully pinned down by the two constraints and

$-u''(W)/u'(W)$ is the measure of absolute risk aversion. In case of constant relative risk aversion with risk aversion parameter r the condition simplifies to $r(2r - 1)/(W^2(u')^3) \geq 0$.

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 wage payments associated with a given level of effort, we can characterise 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 $E(h'(v_k)v_k'(e^*)) \geq 0$.[4]

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.

[4] 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 $e > 0$. So 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; 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 profit depends 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 expected ability anticipate that they are more likely to be successful than less able agents. Since $v_H \geq v_L$ 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

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

where σ^2 is the variance of the prior distribution $F(\tilde{\theta})$.[5] 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 unfavourable 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 realised profits. If the principal did, however, have a precise idea about the agent's ability beforehand, there is little additional 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 maximised. If the principal implements no effort at all, a high outcome is very informative about the agent's ability since it can not 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.

3.2 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 3.2.1 we will see that, in order to derive our key results, we can ignore the change in period-one contracts since

[5]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 3.2.2 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. So 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 the exposition without affecting our results.[6] Similarly, in Appendix A5 we show that the the results by Drew Fudenberg, Bengt Holmström and Paul Milgrom (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 maximises $E\left(u(W+w_k)\right)-C(e_1)$.

[6] 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 maximise total surplus, the distribution of bargaining power does not affect actions in period two and the principal's overall profits remain unchanged.

3.2.1 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 as well as the 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 maximises 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 3.3 we analyse 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. So 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 either of the two policies 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.

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.*

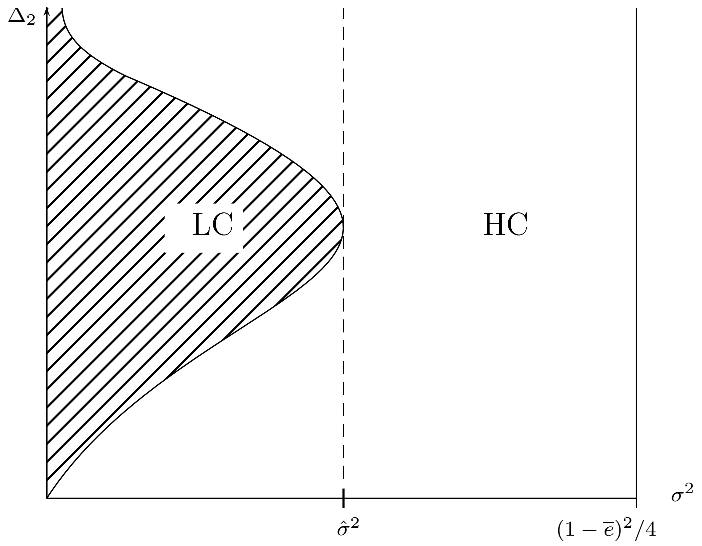

Figure 3.1: Comparison between HC and LC.

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.[7]

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 favour 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 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.

[7]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 $\tilde{\theta}$.

Unfortunately, the problem faced by the principal 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.[8] 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. In this case, 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 3.3.

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 capitalises 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. So the expected quality of agents is constant over time and independent of σ^2. The results for AC, however, are ambiguous: From an ex-ante perspective, 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.

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

This explains why profits under AC can be increasing or decreasing in σ^2.

3.2.2 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

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

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.[9] 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 the effort in the static setting 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

[9]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. Consequently, the indirect effect will be negative under HC and ambiguous under LC and 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 characterise the optimal employment policies more closely while preserving the key trade-offs of the more general model.

3.3 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 characterising 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.[10] 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 induce the agent to exert 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.[11] 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

[10] We discuss the key implications of square root utility at the end of Section 3.3.1.

[11] 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 costs 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 a high period-one outcome, 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 $\hat{\Delta}_0 \geq \hat{\Delta}_H$.

3.3.1 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 for the value of success such that the former is 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 his higher expected ability, a new manager is 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.

So 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 interval 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. So 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. Hence, 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

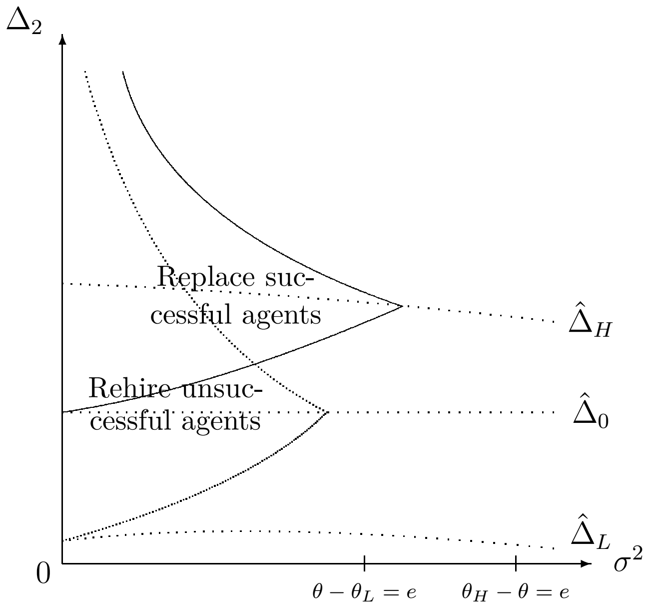

Figure 3.2: Optimal Continuation Decisions in the σ^2-Δ_2-space.

that is available. We can now translate these regions into optimal employment policies 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 in Figure 3.2 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.

Characterising 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. Also, 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

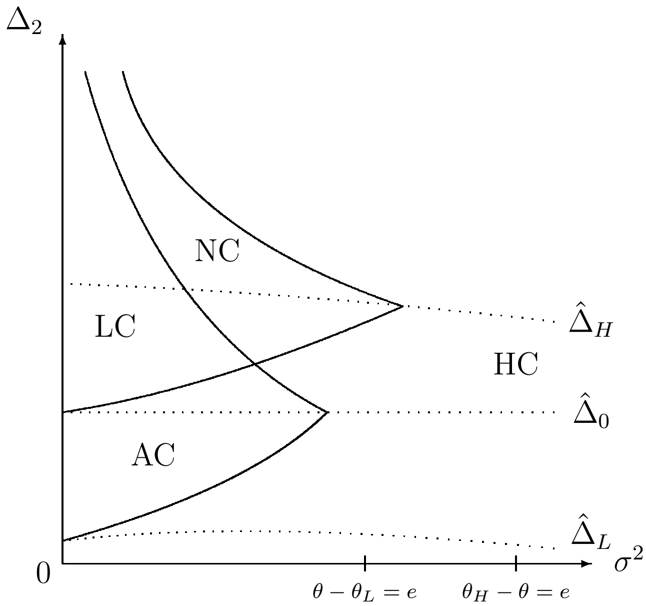

Figure 3.3: Optimal Policies in the σ^2-Δ_2-space.

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 may 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. □

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 favours the policy conditioning the reemployment decision most severely on an agent's expected 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 for some levels of uncertainty the ability advantage of a successful agent can be off-set by the effort advantage of a new agent, while the ability 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.[12] The opposite situation will usually appear if $\theta_H - \theta > \theta - \theta_L$: The expected ability of successful agents is so high that it always pays to reemploy these agents, while unsuccessful agents can make up for their low ability by exerting effort. 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 the sufficient condition 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

[12]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.

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 $(\theta+e)\hat{\Delta}_H+(1-\theta-e)\hat{\Delta}_L \geq \hat{\Delta}_0$.

3.3.2 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 that involve hiring agents of inferior ability are to be optimal.

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

Lemma 4. *1. 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 $e > \theta - \theta_L$.*

2. 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 $e > \theta_H - \theta$.

3. 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 $e > \max\{\theta - \theta_L, \theta_H - \theta\}$.

4. *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 costs 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 costs 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 values of success. 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.

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 3.2 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. So 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.

3.4 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 value of success is very high or very low, but he may be lenient vis-à-vis failure if the importance of effort lies between those 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 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.

A4 Mathematical Appendix

Proof of Proposition 3. 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

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

For $\Delta_2 = 0$ the principal implements zero effort regardless of the period one outcome $k \in \{H, 0, L\}$ and does not need to pay any wages in order to so so. By the envelope theorem, the principal's profits increase most strongly in Δ_2 when hiring a successful manager ($k = H$) and least strongly when hiring an unsuccessful manager ($k = L$). Hence, for any given period-one contract, the principal will want to rehire a manager if and only if he has been successful if Δ_2 is close to zero. 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. □

Proof of Lemma 1. 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 $e > 0$. □

Proof of Lemma 4.

1. 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 can not 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.

2. Necessity: *NC* includes rehiring unsuccessful agents. By Lemma 3 this can only be optimal if it is optimal at $\hat{\Delta}_{\text{H}}$. 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 $\hat{\Delta}_{\text{H}}$ *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.

3. Follows from the proof of 1. The necessary condition is the stricter of the two

conditions (3.7) and (3.10) for reemploying inferior agents.

4. The first part is a direct consequence of Lemma 2 and 3. The remainder follows from 1, 2, and 3.

□

Proof of Proposition 6. 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 $\theta_L \to \theta$, 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

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

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. So we can again 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., $\theta_H \to \theta$. Again, this implies that $\underline{C}^{NC} \to 0$ 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 can not 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. □

A5 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 characterised 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, the principal does indeed make the same effort and employment decisions under short-term contracts. The principal can replicate *any* long-term contract at the same cost. So 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 once period one payments have been made.

Bibliography

Akerlof, George A. 1970. "The Market for Lemons: Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, 84(3): 488–500.

Asch, Beth J. 1990. "Do Incentives Matter? The Case of Navy Recruiters." *Industrial and Labor Relations Review*, 43(3): 89–106.

Attar, Andrea, Thomas Mariotti, and François Salanié. 2011. "Nonexclusive Competition in the Market for Lemons." *Econometrica*, 79(6): 1869–1918.

Bebchuk, Lucian A., and Jesse Fried. 2004. *Pay without Performance: The Unfulfilled Promise of Executive Compensation.* Cambridge, MA: Harvard University Press.

Bebchuk, Lucian A., Joseph E. Bachelder, Roel C. Campos, Byron S. Georgiou, Alan G. Hevesi, William Lerach, Robert Mendelsohn, Robert A.G. Monks, Toby Myerson, John F. Olson, Leo E. Strine, and John C. Wilcox. 2006. "Director Liability." *Delaware Journal of Corporate Law*, 31(3): 1011–1045.

Benmelech, Effi, Eugene Kandel, and Pietro Veronesi. 2010. "Stock-Based Compensation and CEO (Dis)Incentives." *Quarterly Journal of Economics*, 125(4): 1769–1820.

Biais, Bruno, and Catherine Casamatta. 1999. "Optimal Leverage and Aggregate Investment." *Journal of Finance*, 54(4): 1291–1323.

Biais, Bruno, Thomas Mariotti, Jean-Charles Rochet, and Stéphane Villeneuve. 2010. "Large Risks, Limited Liability, and Dynamic Moral Hazard." *Econometrica*, 78(1): 73–118.

Bolton, Patrick, José Scheinkman, and Wei Xiong. 2006. "Executive Compensation and Short-Termist Behaviour in Speculative Markets." *Review of Economic Studies*, 73(3): 577–610.

Boot, Arnoud W.A. 2000. "Relationship Banking: What Do We Know?" *Journal of Financial Intermediation*, 9(1): 7–25.

Boot, Arnoud W.A., Anjan V. Thakor, and Gregory F. Udell. 1987. "Competition, Risk Neutrality and Loan Commitments." *Journal of Banking & Finance*, 11(3): 449–471.

Brick, Ivan E., Oded Palmon, and John K. Wald. 2012. "Too Much Pay-Performance Sensitivity?" *Review of Economics and Statistics*, 94(1): 287–303.

Bureau of Economic Analysis. 2008-2011. "Annual Industry Accounts of the United States." *http://www.bea.gov/industry*, United States Department of Commerce, (accessed May 31, 2012).

Campbell, Tim S. 1978. "A Model of the Market for Lines of Credit." *Journal of Finance*, 33(1): 231–244.

Chade, Hector, and Virginia N. Vera de Serio. 2011. "Wealth Effects and Agency Costs." *Unpublished.*

Conyon, Martin J., and Kevin J. Murphy. 2000. "The Prince and the Pauper? CEO Pay in the United States and United Kingdom." *Economic Journal*, 110(467): 640–671.

Dang, Tri Vi, Gary Gorton, and Bengt Holmström. 2009. "Opacity and the Optimality of Debt for Liquidity Provision." *Unpublished.*

Denis, David J., Diane K. Denis, and Atulya Sarin. 1997. "Ownership Structure and Top Executive Turnover." *Journal of Financial Economics*, 45(2): 193–221.

Detragiache, Enrica, Paolo Garella, and Luigi Guiso. 2000. "Multiple versus Single Banking Relationships: Theory and Evidence." *Journal of Finance*, 55(3): 1133–1161.

Edmans, Alex, Xavier Gabaix, Tomasz Sadzik, and Yuliy Sannikov. 2012. "Dynamic CEO Compensation." *Journal of Finance*, forthcoming.

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

Fahlenbrach, Rüdiger, and René M. Stulz. 2011. "Bank CEO Incentives and the Credit Crisis." *Journal of Financial Economics*, 99(1): 11–26.

Fischer, Paul, and Steven Huddart. 2008. "Optimal Contracting with Endogenous Social Norms." *American Economic Review*, 98(4): 1459–1475.

Fudenberg, Drew, Bengt Holmström, and Paul Milgrom. 1990. "Short-Term Contracts and Long-Term Agency Relationships." *Journal of Economic Theory*, 51(1): 1–31.

Gabaix, Xavier, and Augustin Landier. 2008. "Why Has CEO Pay Increased So Much?" *Quarterly Journal of Economics*, 123(1): 49–100.

Gibbons, Robert, and Kevin J. Murphy. 1992. "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence." *Journal of Political Economy*, 100(3): 468–505.

Grossman, Sanford J., and Oliver D. Hart. 1983. "An Analysis of the Principal-Agent Problem." *Econometrica*, 51(1): 7–45.

Hadlock, Charles J., and Gerald B. Lumer. 1997. "Compensation, Turnover, and Top Management Incentives: Historical Evidence." *Journal of Business*, 70(2): 153–187.

Healy, Paul M. 1985. "The Effect of Bonus Schemes on Accounting Decisions." *Journal of Accounting and Economics*, 7(1-3): 85–107.

Holmström, Bengt, and Jean Tirole. 1998. "Public and Private Supply of Liquidity." *Journal of Political Economy*, 106(1): 1–40.

Holmström, Bengt, and Paul Milgrom. 1987. "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica*, 55(2): 303–328.

Holmström, Bengt, and Paul Milgrom. 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, and Organization*, 7: 24–52.

Holthausen, Robert W., David F. Larcker, and Richard G. Sloan. 1995. "Annual Bonus Schemes and the Manipulation of Earnings." *Journal of Accounting and Economics*, 19(1): 29–74.

Huson, Mark R., Robert Parrino, and Laura T. Starks. 2001. "Internal Monitoring Mechanisms and CEO Turnover: A Long-Term Perspective." *Journal of Finance*, 56(6): 2265–2297.

Inderst, Roman, and Marco Ottaviani. 2009. "Misselling through Agents." *American Economic Review*, 99(3): 883–908.

International Monetary Fund. 2010. "Transcript of a Press Briefing on the International Monetary Fund's Global Financial Stability Report." *http://www.imf.org/external/np/tr/2010/tr042010.htm*, April 20.

Jenter, Dirk, and Katharina Lewellen. 2010. "Performance-Induced CEO Turnover." *Unpublished.*

John, Kose, and Yiming Qian. 2003. "Incentive Features in CEO Compensation in the Banking Industry." *Economic Policy Review*, 9(1): 109–121.

John, Teresa A., and Kose John. 1993. "Top-Management Compensation and Capital Structure." *Journal of Finance*, 48(3): 949–974.

Kahn, Charles, and Gur Huberman. 1988. "Two-Sided Uncertainty and Up-or-Out Contracts." *Journal of Labor Economics*, 6(4): 423–444.

Kaplan, Steven N., and Bernadette A. Minton. 2006. "How has CEO Turnover Changed? Increasingly Performance Sensitive Boards and Increasingly Uneasy CEOs." *National Bureau of Economic Research Working Paper 12465.*

Landier, Augustin. 2006. "Entrepreneurship and the Stigma of Failure." *Unpublished.*

Larkin, Ian. 2007. "The Cost of High-Powered Incentives: Employee Gaming in Enterprise Software Sales." *Unpublished.*

Malherbe, Frédéric. 2010. "Self-Fulfilling Liquidity Dry-Ups." *National Bank of Belgium Working Paper 185.*

Manso, Gustavo. 2011. "Motivating Innovation." *Journal of Finance*, 66(5): 1823–1860.

Martin, J. Spencer, and Anthony M. Santomero. 1997. "Investment Opportunities and Corporate Demand for Lines of Credit." *Journal of Banking & Finance*, 21(10): 1331–1350.

Murphy, Kevin J. 1999. "Executive Compensation." In *Handbook of Labor Economics.* Vol. 3B, , ed. Orley Ashenfelter and David Card, 2485–2563. Amsterdam: Elsevier.

Ongena, Steven, and David C. Smith. 2000. "What Determines the Number of Bank Relationships? Cross-Country Evidence." *Journal of Financial Intermediation*, 9(1): 26–56.

Oyer, Paul. 1998. "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality." *Quarterly Journal of Economics*, 113(1): 149–185.

Pagano, Marco, and Paolo Volpin. 2008. "Securitization, Transparency and Liquidity." *Center for Economic Policy Research Discussion Paper 7105.*

Paribas, BNP. 2007. “BNP Paribas Investment Partners temporaly suspends the calculation of the Net Asset Value of the following funds : Parvest Dynamic ABS, BNP Paribas ABS EURIBOR and BNP Paribas ABS EONIA.” *http://www.bnpparibas.com/en/pressroom*, August 9.

Philippon, Thomas, and Vasiliki Skreta. 2012. “Optimal Interventions in Markets with Adverse Selection.” *American Economic Review*, 102(1): 1–28.

Rajan, Raghuram G. 1992. “Insiders and Outsiders: The Choice between Relationship and Arm’s-Length Debt.” *Journal of Finance*, 47(4): 1367–1400.

Securities Industry and Financial Markets Association. 2002-2011. “US ABS Issuance and Outstanding.” *http://www.sifma.org/research/statistics.aspx*, (accessed May 31, 2012).

Sharpe, Steven A. 1990. “Asymmetric Information, Bank Lending and Implicit Contracts: A Stylized Model of Customer Relationships.” *Journal of Finance*, 45(4): 1069–1087.

Spagnolo, Giancarlo. 2000. “Stock-Related Compensation and Product-Market Competition.” *RAND Journal of Economics*, 31(1): 22–42.

Spagnolo, Giancarlo. 2005. “Managerial Incentives and Collusive Behavior.” *European Economic Review*, 49(6): 1501–1523.

Spear, Stephen E., and Cheng Wang. 2005. “When to Fire a CEO: Optimal Termination in Dynamic Contracts.” *Journal of Economic Theory*, 120(2): 239–256.

Stein, Jeremy C. 1988. “Takeover Threats and Managerial Myopia.” *Journal of Political Economy*, 96(1): 61–80.

Stein, Jeremy C. 1989. “Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior.” *Quarterly Journal of Economics*, 104(4): 655–669.

Sufi, Amir. 2009. “Bank Lines of Credit in Corporate Finance: An Empirical Analysis.” *Review of Financial Studies*, 22(3): 1057–1088.

Tervio, Marko. 2008. “The Difference That CEOs Make: An Assignment Model Approach.” *American Economic Review*, 98(3): 642–668.

TheCityUK. 2011. “Bond Markets.” *http://www.thecityuk.com/research/our-work/reports-list/*, (accessed May 31, 2012).

Thiele, Henrik, and Achim Wambach. 1999. “Wealth Effects in the Principal Agent Model.” *Journal of Economic Theory*, 89(2): 247–260.

Tian, Xuan, and Tracy Y. Wang. 2012. “Tolerance for Failure and Corporate Innovation.” *Review of Financial Studies*, forthcoming.

Tirole, Jean. 2012. “Overcoming Adverse Selection: How Public Intervention Can Restore Market Functioning.” *American Economic Review*, 102(1): 29–59.

U.S. Securities and Exchange Commission. 2010. “Proposed Rule: Asset-Backed Securities.” *http://www.sec.gov/rules/proposed/proposedarchive/proposed2010.shtml*, April 7.

von Thadden, Ernst-Ludwig. 1995. “Long-Term Contracts, Short-Term Investment and Monitoring.” *Review of Economic Studies*, 62(4): 557–575.

Warner, Jerold B., Ross L. Watts, and Karen H. Wruck. 1988. “Stock Prices and Top Management Changes.” *Journal of Financial Economics*, 20: 461–492.

Zhou, Xianming. 2000. “CEO Pay, Firm Size, and Corporate Performance: Evidence from Canada.” *Canadian Journal of Economics*, 33(1): 213–251.

Curriculum Vitae

seit Dezember 2007	Doktorandenstudium der Ökonomie *Munich Graduate School of Economics*
November 2011 - Dezember 2011	Visiting Scholar *Kellogg School of Management, USA*
Oktober 2009 - April 2010	Visiting PhD Student *Toulouse School of Economics, Frankreich*
August 2007	Diplom in Volkswirtschaftslehre *Ludwig-Maximilians-Universität München*
September 2006 - April 2007	Austauschstudent *Université Aix-Marseille II, Frankreich*
September 2004	Vordiplom in Volkswirtschaftslehre *Eberhard Karls Universität Tübingen*
Juni 2003	Abitur *Erzbischöfliches Suitbertus-Gymnasium, Düsseldorf*
September 1983	Geboren in Neuss, Deutschland