Three Essays in Applied Microeconomics

Inaugural-Dissertation
zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften
durch die
Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von
Christoph Wagner
aus München

Bonn 2013

Dekan: Prof. Dr. Klaus Sandmann
Erstreferent: Prof. Dr. Paul Heidhues
Zweitreferent: Prof. Dr. Daniel Krähmer

Tag der mündlichen Prüfung: 11.3.2013

Acknowledgements

First of all, I want to thank my supervisor Paul Heidhues. He provided very detailed and constructive comments for all chapters of this thesis. Without his questions and willingness to challenge my results this work would not have attained its current level. I also highly appreciate his numerous invitations and generosity after he had taken over a chair in Berlin.

I also owe a lot to Daniel Krähmer, my second supervisor. He always found time to listen to my ideas and gave me very valuable feedback. His help improved this thesis significantly and I am indebted to his encouraging support.

Moreover, this work profited from discussions with and comments from Rafael Aigner, Jörg Budde, Carsten Dahremöller, Deniz Dizdar, Markus Fels, Fabian Herweg, Eugen Kovac, Matthias Kräkel, Sebastian Kranz, Matthias Lang, Tobias Luhmer, Tymofiy Mylovanov, Ariel Pakes, Patrick Rey, Christoph Schottmüller, Christian Seel, Dora Simroth, Axel Stahmer, Philipp Strack, Deszö Szalay, Thomas Tröger, Alexander Westkamp, Venuga Yokeeswaran, Simon Zehnder, and many other people. I am very grateful to Dora, Axel, and Deszö for reading parts of this thesis. Special thanks goes to Jörg Budde whose topics course inspired Chapter 2 and to Venuga with whom I collaborated in an early stage of this chapter.

The Bonn Graduate School of Economics provided financial support and an excellent research atmosphere for over four years. In particular, I would like to thank my fellow students and the administration of the BGSE: Urs Schweizer, Silke Kinzig, Corinna Lehmann, and Pamela Mertens.

Finally, I am greatly indebted to my friends Attila, Axel, Christof, and Lorenz. Thanks for all the wise support and special time with you. I hope I was not too much a strain in the last years.

Above all I would like to thank my family. Without you I would have never been able to be where I am now. I am especially grateful to my parents for giving me so much freedom and love and for my brother's advice. In addition, I want to thank him for his hospitaltiy in the final stage of my thesis and assisting me when he had more important things to do.

Contents

List of Figures

Introduction

General equilibrium theory had traditionally a very simplified view of the firm: Inputs are transformed into outputs according to a technological process, which is summarized in a production function. Since then the economic profession has made a long step to derive a more accurate picture of the firm. On the one hand, realizing that market structure and firm behavior influence each other made it inevitable to take a closer look at the way companies operate. On the other hand this motivation was accompanied by an important progress in economic theory: The development of game theory and information economics gave academics the necessary methods to study the firm more extensively from their vantage point.

This thesis adds to this research program and consists of three chapters. Each chapter is a contribution to get a more refined perspective on a specific part of firm activity: The first chapter deals with the firm's problem to deter employee poaching by a rival firm. Optimal management compensation is considered in the second chapter, and the last chapter is devoted to the firm's distribution strategy.

Chapter 1 focuses on a firm that has to defend herself against a competitor who tries to raid parts of her workforce. This is a vital topic for many firms. Just recently, in 2011, the American Department of Justice prohibited no-employee-poaching agreements among six major high-tech enterprises in the US. There has also been a trend that further aggravates the employing firm's problem: Incentive contracts are increasingly used for worker compensation, which allows raiding firms to write contracts that are attractive to different worker types at the same time. Moreover, in some industries, e.g. in large parts of the public sector, outside offers are not matched, which makes it even harder for firms to keep their workforce.

There are two major contributions of this chapter: Startlingly, there does not exist any research so far that is directly concerned with how companies should defend themselves against employee raids. We explain how poaching can be contained if the raider uses incentive contracts to solicit employees and the current employer cannot match outside offers. We integrate in our model a realistic feature that is pivotal for a successful defense: The current employer has superior information about his worker types compared to the raider. This information asymmetry imposes a restriction on the raider's profits that he can obtain from soliciting workers. In particular, if he intends to poach low- and high-ability types, but cannot distinguish them, he has to provide the high type with an information rent. This rent payment is the employing firm's lever for keeping some workers in her company: By making a convenient contract choice to her workers, she can lift the relative burden of this payment so high that the raider prefers to solicit only a single worker type. This novel information rent argument as defense mechanism is our first contribution.

Our second contribution rests on the fact that the average pool of poached workers consists of low-ability types in most models. In reality, however, mostly high-ability types switch jobs despite their higher wages. While there are models that account for this feature, turnover heavily relies on the fact that the raiding firm has a higher valuation than the current employer for these workers. We provide an additional strategic explanation which is consistent with this observation—even when workers have the same productivity in the raiding company.

Chapter 2 deals with optimal incentive provision for one of the most important part of the firm's workforce—the management. In the standard principal-agent model of moral hazard, the agent and the principal know which distribution over profits is induced by each effort level. If the owner of a firm hires a *new* manager, however, especially in the initial phase, the owner has superior knowledge about the manager's influence on the firm. Thus, we modify the standard moral hazard model by assuming that the stochastic relationship between effort and profits is private information of the principal.

There are two important findings in this chapter: First, in the existing literature of informed principal models with moral hazard, the principal's information acts distortive, i.e. the principal's payoff differs from the payoff in the standard model. In contrast, we identify in a general model a linear independence condition under which the principal's payoff is unaffected by his private information. Roughly spoken, this condition guarantees that the agent's incentives can be sufficiently fine-tuned to provide optimal effort independent of her belief about the principal's information. A similar condition has also been employed for irrelevance results in adverse selection models. It is the first approach, however, to apply it in a moral hazard setting.

Second, we show for a special case of independent interest that the linear independence condition is also necessary for the principal's information to be irrelevant. Because this condition is essentially satisfied if and only if the number of available profit states is large enough there exists an interesting relationship: distortive effects of the principal's information directly hinge on this number. This role of contractible signal states is novel and it gives a new perspective on existing informed principal models with moral hazard.

Finally, Chapter 3 is concerned with the firm's optimal distribution strategy for its goods. More specifically, we scrutinize the emergence of so-called shopping clubs, which are frequently used by firms in the apparel industry to distribute goods for clearance sales. Surprisingly, the business model of shopping clubs, an exclusive online sales platform for a selected community, is in dire contrast to their actual business conduct. We provide a different rationale for their appearance and build a model that assigns clubs a relative cost advantage to traditional store distribution. While this assumption clarifies why shopping clubs are used for clearance sales, it raises the question why they are not also employed for regular sales.

To explain this sales pattern, we consider a monopolistic durable-good firm without commitment power, which sells a good over two periods. Each period the monopolist has to decide whether to distribute the good via his own store or use a shopping club as a distributor. Because the monopoly has all the bargaining power and the club is more cost efficient, the firm lets the club distribute the good in the second period. Regarding the first period, the firm faces a tradeoff between selling via the store and the club. While the own store involves higher distribution costs than the club, distribution via the shopping club implies a loss over price control. Because all the club's second-period profits are absorbed by the firm, the club's first-period price fails to internalize the firm's future profits and is therefore distorted from the firm's perspective. We find that under sufficiently low cost differences between store and club distribution, the monopolist prefers selling his good via the shop in the first period. This sales behavior is in line with the observation that stores are used for regular and shopping clubs for clearance sales.

Chapter 1

Defending Raids on Employees

We analyze an employee poaching model in which the raiding firm makes use of incentive contracts that can potentially attract all worker types to her. If the current employer, unlike the raider, knows his workers' types, but has committed to a no offer-matching policy, can he profitably induce at least some worker types to stay in his company?

Our answer is affirmative. We provide an information rent argument and characterize the equilibrium behavior that deters the raiding firm from poaching all employee types. Moreover, we give a novel explanation for the observation why predominantly high worker types are poached—even if employee productivity is the same in the raiding firm.

1.1 Introduction

Employee poaching is a major concern for firms in tight labor markets: Raiding companies can fill vacant positions swiftly, save on training costs and even increase their productivity by incorporating knowledge from other firms in their own. On the other hand employers take individual measures (e.g. foster social ties among co-workers, ponder their offer-matching policy, etc.), but also join forces to avoid losing parts of their workforce. For instance, DOJ-inquiries questioning the lawfullness of no employee-poaching agreements among major high-tech enterprises in the US underline the importance of this practice.[1]

We present an employee-poaching model in which the current employer has private information regarding his workers' types and the raiding firm makes use of incentive con-

[1]The final judgement concerned Adobe, Apple, Google, Intel, Intuit and Pixar (Department of Justice, Antitrust Division, case: 1:10-cv-01629).

tracts that potentially attract all worker types to her. Moreover, the current employer has comitted to not match any outside offers to his workers.

There are two major contributions of this paper: First, to our knowledge there does not exist any research on how companies should defend themselves against employee raids. How can poaching be at least partially contained without matching offers? We propose a strategy based on a new information rent argument. It shows how the current employer can enforce to profitably keep some employees within the company. In addition, we characterize the equilibrium defending contract and explain the underlying mechanism that deters the raider from poaching all worker types.

Second, in previous models of the literature the pool of poached workers consists of low-ability types on average. Lazear (1986) points out, however, that mainly high-ability types switch jobs despite higher wages. While his model accounts for this, turnover heavily relies on the fact that the raiding firm has a higher valuation than the current employer for these workers. We provide an additional strategic explanation, which is consistent with this observation—even when workers have the same productivity in the raiding company.

In Section 1.2 we discuss the literature related to our paper, before we introduce our basic model setup in Section 1.3. A company employs workers of two ability types. Ability is completely general and has no firm-specific character, as capacity for teamwork, leadership, willingness to work hard and so on. A worker's type is private information to the current employer and unobservable to outside firms.[2] The current employer assigns to each type of employee an effort level and a remuneration.[3] A task level could, for instance, determine a range of responsibilities an employee has to meet or measure the difficulty of a certain kind of work.[4] The disutility generated by a certain amount of effort depends on the employee's type, so that a choice of distinct contracts is in general desirable for the current employer.

---

[2]This need not always be the case: employees network at conferences, trade fairs, or on internet platforms and headhunters invest plenty of ressources to obtain a versed view on the workforce in a specific industry. That said, information accuracy dwindles on lower hierarchy level as the number of available employees increases. Moreover, these workers tend to be less concerned about their careers and show less endeavor to make themselves visible to outside firms.

[3]We assume that employees have no bargaining power and there do not exist any moral hazard problems.

[4]For instance, a salesman could be assigned to an easy or challenging sales region.

In our model both effort-wage combinations are observable to the outside firm.[5] Nevertheless, the raiding firm cannot observe the specific contracts of the employees that she is trying to poach—otherwise she could simply identify the workers' types.[6] After observing the current employer's offers, the raider uses incentive contracts or menus of effort-wage contracts, which can be attractive to both types of workers at the same time.[7] [8]

Finally, each employee chooses the most attractive contract available to her and possibly switches to the raiding firm. In our model no matching offers are made. This practice is quite common, e.g. in large parts the public sector, and some reasons, why a firm may pursue this policy, are listed in the literature review.

So the picture we have in mind here is a department of low-level white collar workers that share the same job title, but differ quite substantially in their skills, in the work they do, and what they earn. This is consistent with real world data: For instance, Baker et al. (1994) report a large wage variation within job levels of the same firm. An outside firm has vacancies to fill and tries to poach employees by offering incentive contracts to them. In anticipation of this event, the current employer designs effort-wage contracts to minimize the costs of this raid as much as possible, because he has previously committed to not match any outside offers.[9]

---

[5] Surveys on job salaries are regularly published and they also provide pay ranges for what a specific employee, say an engineer with three years working experience, can earn.

[6] While it is true that in reality employees often reveal their current wage at some point in the negotiation process with a new employer, there are reasons, why his potential to draw conclusions is rather limited: First, uncertainty can turn the wage into a noisy signal for him. For instance, if there's a bonus component in the worker's remuneration, then he cannot be sure whether high wages were due to the worker's skill or a good performance of her old firm. And second, considerable parts of an employee's compensation may be non-pecuniary and unobservable to outsiders. It's hard to verify a worker's prospects of promotion, authority to give directives, scope of duties, access to seniors, office space, and so on. Thus, we abstract from these intricate issues by assuming that the whole range of contracts is observable to raiding firms, but not a contract or type of specific workers.

[7] We allow for this feature because firms have been incorporating flexibility in their contracts over the last decades and do not always confine themselves to simple contract structures. Lebow et al. (1999) point out that eligibility for variable-pay programs is rapidly spreading throughout firm organizations and has been steadily rising for lower-income job classes. Moreover, they report compensation schemes to be increasingly designed in a decentralized manner, e.g. on a team, product line, or division level.

[8] We explain in Section 1.3, why it is not restrictive that the current employer can only use single effort-wage contracts.

[9] Of course, an employer's endeavors to thwart a rival's poaching activities do not stop at merely selecting a task and wage level. Wary employers take measures to figure out or maintain their employee's satisfaction and foster events where co-workers can build up social ties. We strongly believe, however, that workload and remuneration are the most important components for an employee's decision to switch jobs and a good starting point for a first investigation.

We analyze this model in Section 1.4 and begin with a series of necessary equilibrium conditions. At first, we show that the current employer has strictly positive equilibrium profits (Lemma 1), because there always exists a strategy that allows him to profitably keep all low-type workers in the company. Intuitively, if the raiding firm designs a menu that attracts both types of workers, but has to respect self-selection constraints, the high-ability type will benefit from such a contract by obtaining an information rent. By choosing effort-wage combinations with sufficiently high surplus levels for his employees, the current employer can raise the relative burden of the information rent so high that the raider refrains from incentive contracts attractive to the low-type worker. Thus, the current employer is able to retain this employee type in his workforce and, because he need not offer the full surplus of the low-type worker, obtains a strictly positive profit. As the raiding firm has always the option to slightly overbid the current employer, it is immediate that she gets positive equilibrium profits as well (Lemma 2). These two results imply that exactly one worker type is poached and one remains in the firm (Theorem 1).

We then sharpen the overbidding argument to the result that the raider obtains strictly higher equilibrium profits (Lemma 3). If not, the raiding firm could offer the chosen equilibrium contracts with a slightly higher wage and win over both employee types. Because profits were equal across firms, the specific contract choice of the employees would not matter and the gain of employing both types would outweigh the costs of the wage raise—the raiding firm had a profitable deviation. This result will be helpful for Proposition 1, which refines Theorem 1 by showing that the high-type employee will be poached in equilibrium if the share of low-type workers is sufficiently small.[10] The basic idea is that otherwise the current employer necessarily received less than the low-type full surplus, because the raiding firm poached the low type. But the full low-type surplus could always be attained by the current employer by offering a complete extraction contract to both employee types. Then, due to the small share of low types, the raiding firm had no interest in attracting these employees and responded with a full extraction contract for the high type. Thus, the current employer could secure the full low-type surplus for himself and improve on his profits.

[10]To be precise, the share of low-type workers is sufficiently small if in the standard monopolistic screening problem the low type would be excluded from employment under this share.

We then turn to the equilibrium contract choice, which lies at the heart of the defending mechanism. Effort-wage combinations of the current employer are designed in such a way that the poached worker type strictly prefers the contract of the staying type (which is unavailable to her) to her own contract accepted from the raiding firm (Lemma 4). There's an implicit threat to the raiding firm in this structure: If she attempts to poach the remaining type as well, she will have to make a more attractive offer than the current employer to this type. But then not only this type will select the new offer, but also the type she has already poached and she cannot extract so much value from these employees anymore. In equilibrium, the gain of employing both worker types does not compensate the loss due to lower rent extraction. Thus, under such a contract structure, the raider favors to poach just one worker type.

Our final set of necessary conditions concerns the choice of effort levels in equilibrium. Because each firm employs exactly one worker type, there's no need for a distortion in the task assignment and firms can extract the highest profits by choosing efficient effort levels (Lemma 5 and 6). Because wages merely influence the distribution of economic profits, any equilibrium is efficient.

At last, we show that an equilibrium exists in this game (Theorem 2). The idea is to consider two maximization problems: One is designed to extract the maximum surplus from the low-type employee for the current employer under the constraint that the raider prefers to poach exclusively the high type. The other problem is similar but concerns maximum extraction from the high-type employee. One can show that both procedures have a solution (Lemma 8). The maximizer of the problem with the higher value can then be used to construct an equilibrium. Lemma 7 and 9 are auxiliary results.

Two extensions are considered in Section 1.5. First, we revisit the model under the assumption of asymmetric worker productivity in firms. If the raiding firm has a less productive use of workers, the current employer can now profitably prevent her from poaching any employees by endowing each worker type with the maximum surplus the raider could extract. The profits of this prevention strategy are close to zero for negligibly small productivity differences, but approach the full extraction value of both types for very large ones. In contrast, the current employer can also allow poaching activities. Because lower productivity of the raiding firm does not alter his range of profitable surplus levels the logic of Lemma 1 remains unimpeded. Thus, as in the symmetric case, the current employer's profits under poaching can be bounded away from zero (Lemma 10). On the other hand, since only one worker type is kept within the company, his profits under raiding are bounded above by the maximum of what can be extracted from a single type. Hence, comparing profits when raiding is prevented to when it is allowed shows: for small productivity differences employee poaching will always occur, whereas for high ones it is entirely contained (Theorem 3). The case of a more productive raiding firm is more intricate and we sketch the results nontechnically.

Second, we endogenize the current employer's choice whether to match offers or not. We argue that it is a reasonable procedure to shortcut the dynamic offer-matching process by looking at a simultaneous first-price auction. Under this mechanism the current employer and the raider obtain zero profits (Proposition 2). The raider cannot attain any positive gains, because the current employer can always typewise overbid her offer without being interfered by self-selection constraints. But then, since the raiding firm can always overbid the contract of the current employer for any specific type, he obtains zero profits as well. Thus, in this extended game version, the current employer commits to not match any outside offers and partial employee poaching occurs (Theorem 4).

Section 1.6 concludes and provides further avenues for research; omitted proofs can be found in Section 1.7.

1.2 Related Literature

Waldman (1984) is one of the earliest contributions in the literature on employee poaching. In his model outside firms cannot directly observe a worker's skill level but can draw conclusions from the current employer's job assignment decisions. Inefficient matching of employees to jobs follows as employers try to reduce the cost for retaining their workers. Greenwald (1986) dispenses with job assignments and focuses on the characteristics of the worker pool that switches to a new company. If current employers are willing to make counteroffers to their employees, they will give up on a worker precisely when the raiding offer exceeds his productivity. Thus, the group of job changers consists of a disproportionally large share of low-type workers whose productivity is below the wage offered by the outside firm and a winner's curse arises.[11] Gibbons and Katz (1991) features elements from both models above. In their work an employer's layoff decision is observable to outside firms and workers of low ability are displaced first leading to a low-wage market for layoffs and a high-wage market for workers retained. They test for this stigma of layoff by comparing postdisplacement wages and unemployment spells of regular layoffs compared with those of a plant closing (which supposedly lays off workers of all ability levels) and find their thesis confirmed: Even though predisplacement wages are roughly the same, postdisplacement wage offers are considerably lower and time until reemployment longer for layoffs that did not follow a plant closing. Our model shares with these three classic contributions the basic informational structure of the game. There are, however, at least two important differences we would like to point out: First, in our setup there's no signaling component involved. The raiding firm observes the whole set of contracts offered to both groups of workers and can infer each type's surplus, but there's no learning regarding the type itself. Second, even though poaching offers can be contingent upon the signal in these models, the contract structure is simple and does not further distinguish between employee types. In contrast, we allow the raider to use general incentive contracts, which can be attractive to all worker types at the same time and entail a more sophisticated rent extraction.

The idea of using menus of effort-wage combinations to screen different worker types on the labor market is not new to the literature: Salop and Salop (1976) model a competitive labor market where firms use an ordinary wage and a two-part wage scheme to induce self-selection among workers with different quitting probabilities. In a study about the car glass company Safelite, Lazear (2000) analyzes the impact of a switch from hourly wages to piece-rate pay after a new management has been installed. He finds productivity effects around 44% in output per worker. About half of the increase is due to improved incentives, but substantial parts result from a sorting component, i.e. the ability to hire the most productive workers and a reduction in quits among the highest output employees. Oyer and Schaefer (2005) scrutinize the practice of granting stock options to all employees of a company. Among incentive, sorting and retention strategies as explanations for this

[11] One may wonder why outside firms engage in employee poaching at all if there are only losses to be attained. In the literature the adverse selection effect is compensated by a positive probability of exogenous turnover for all ability levels of workers and all wage levels that are offered by outside firms. Perri (1995) is a critical discussion of the winner's curse issue.

behavior they observe only the two latter to be consistent with the data. In the case of options, sorting, in particular for low-level workers, does not occur along the dimension of ability. A firm rather hopes to attract potential employees that are less risk averse and more optimistic with respect to the company's development, which are supposed to be desirable employee character traits for firms. While all of the above cited works incorporate more sophisticated instruments than mere wage offers, none takes a deeper look at the firms from whom labor is taken. In our model we explicitly consider the strategic wage-setting effects of employers whose workers are prone to being raided by an outside firm.

Finally, there's a strand of literature that highlights the costs and benefits of a firm's offer-matching policy. Lazear and Gibbs (2009, Chapter 4) point out that this decision influences a worker's on-the-job search intensity. If a valuable employee receives an outside offer and the firm has decided to make counteroffers, this worker obtains a pay raise even though he might have not been willing to leave the company. Under no offer-matching these empty threats will not be brought up, at the risk of losing employees who are serious about quitting. Thus, counteroffers provide incentives to search for outside offers even for those workers who are fairly satisfied with their job. As a result employers have to tradeoff keeping workers under frequent pay raises compared to not increasing wages, but losing workers at times. A similar consideration is presented in Postel-Vinay and Robin (2004). In addition, they state conditions under which the labor market splits up in two segements, one with offer-matching firms and the other with nonmatching companies. Weiss (1990, Chapter 4) stresses the risk of eroding the personnel's morale when a firm decides to match offers. In our extension section, we analyze the current employer's behavior when he has to settle upon one of the two policies. Our main focus, however, is not to present an additional tradeoff regarding offer-matching. We rather want to add robustness to our no-counteroffer assumption, which we use throughout our main analysis.

1.3 Model

In our model there are two employers, the current employer C and the raiding employer R. The raiding firm has job positions to fill and tries to poach employees from firm C. By anticipating R's approach employer C chooses his strategic variables to minimize the damage from losing employees. In the following we will assume that the current employer will not match the raider's offer in any way and return to this issue in Section 1.5.

Workers of firm C are assumed to take on type $\theta \in \Theta = \{\theta_L, \theta_H\}$ and the share of worker type θ_L and θ_H is given by $0 < \alpha < 1$ and $1 - \alpha$ respectively.[12] We will refer to θ_L as the low-ability and θ_H as the high-ability type. The parameter α is common knowledge to both firms and the mass of workers is normalized to unity. Workers of type θ that provide effort level $e \in \mathbb{R}_+$ at wage $w \in \mathbb{R}$ have a quasi-linear utility of the form

$$u(e, w, \theta) \;=\; w - c(e, \theta),$$

with the standard properties:

$$\begin{aligned} &c'(e,\theta) > 0, \qquad c''(e,\theta) > 0 \qquad \text{for all } e > 0,\ \theta \in \Theta, \\ &c'(e,\theta_L) > c'(e,\theta_H) \qquad \text{for all } e > 0. \end{aligned}$$

Thus, providing effort is increasingly costly, but less so for high-ability workers. Moreover, $c(0,\theta) = 0$, $\lim_{e \searrow 0} c'(e,\theta) = 0$, and $\lim_{e \to \infty} c'(e,\theta) = \infty$ for all $\theta \in \Theta$. Each employee of firm C has an outside option value of zero.

Both employers compete for workers by sequentially selecting effort-wage schedules. The current employer moves first by offering $\big((e_L^C, w_L^C), (e_H^C, w_H^C)\big) \in (\mathbb{R}_+ \times \mathbb{R})^2$ specifying for each worker type a task level and a remuneration. Note that firm C knows the worker's type and can assign an effort-wage combination to an employee. Thus, for instance, the low type θ_L cannot select bundle (e_H^C, w_H^C), the high-type contract from firm C. After observing the current employer's offer the raiding firm R chooses her menu $\big((e_L^R, w_L^R), (e_H^R, w_H^R)\big) \in (\mathbb{R}_+ \times \mathbb{R})^2$. But unlike firm C, she does not know the worker's type and has to rely on self-selection constraints in order to induce employees to pick the bundle set up for them.[13] Thus, each employee type has three effort-wage combinations available. Moreover, we will assume that there are no personalized contracts: each type has the same offer from firm

[12] At times we will use the notation α_I if we want to refer to the share of type θ_I with $I \in \{L, H\}$.

[13] Two remarks are noteworthy at this point: First, the current employer's menu does not contain any signal about workers' types for the raiding firm R. But, by observing it, she can determine the utility that each worker type will obtain from working at firm C. In that sense it is without loss of generality, to look at simple effort-wage combinations of firm C. Under more complex contract structures, the raiding firm can still not infer the workers' types, but evaluate their benefits from the current employer's menu.

C, and every worker gets the same menu from firm R. For convenience, we define (e^I, w^I) as the menu of contracts offered by firm I with $I \in \{C, R\}$.

Workers have the same productivity at both firms and there are no complementarities or substitution effects across workers. Output $f(e)$ is increasing but concave in effort. More specifically, $f(0) = 0$, $f'(e) > 0$ and $f''(e) \leq 0$ for all $e \geq 0$. The price of the output good is chosen to be unity. In what follows we will assume that production and cost functions are well-behaved in the sense that in the standard monopolistic screening problem first-order conditions are sufficient for characterizing the optimal solution.[14]

The timing of the game is as follows:

1. Firm C selects its effort-wage menu (e^C, w^C).

2. Firm R observes C's menu and selects its effort-wage menu (e^R, w^R).

3. Employees select their optimal bundle or opt out.

We solve the game for pure subgame perfect equilibria.[15]

1.4 Analysis

In the following we will first develop some basic results in order to state our main equilibrium characterization theorem. After that further necessary equilibrium conditions are derived regarding the poached worker type, firm profits, the current employer's defending mechanism, and effort levels. Finally, we will show that an equilibrium exists in this game.

1.4.1 Basic results

We start by showing that the current employer C can always guarantee himself strictly positive profits by offering an appropriate menu to his employees. Thus, in any subgame

Second, given workers' utilities of the current employer's menu, the standard Revelation Principle applies and we can restrict the raider's offer to an effort-wage combination for each type. More subtle issues arise, however, if the current employer can make his offer contingent upon the proposal of the raiding firm. This leads directly to the problem of competing mechanisms. We do not further investigate this point, because optimal contracts are hard to identify in these environments. See Epstein and Peters (1999) for a general discussion of this topic.

[14] If $0 < \theta_H < \theta_L$, the specification $f(e) = e$ and $c(e, \theta) = \frac{1}{2}\theta e^2$, akin to the model of Mussa and Rosen (1978), satisfies this requirement.

[15] Under minor changes of the model, this setup could also be interpreted as sequential competition between a third-degree and a second-degree price discriminating firm. For more on oligopolistic price discrimination see Stole (2007).

perfect equilibrium firm C must attain strictly positive profits. Intuitively, after observing firm C's menu, the raiding employer R can derive the surplus of each worker type in case they select their offer from firm C. He then has to consider which type to attract or exclude in order to maximize his profits. In particular, if the surplus is equal to zero for both worker types, firm R faces the standard monopolistic screening problem. In this situation the raider can either poach exclusively the high type or attract both types, but leave an additional information rent to the high type. Hence, whenever the loss of paying the information rent exceeds the gain from serving the low type, he will focus on the high type and exclude the low type and vice versa. Now consider the case when the surplus is positive, but equal for both types of workers. As an employee's utility is separable in money this will induce firm R to adjust his wages, but will not affect the optimal effort levels chosen in the screening problem. Consequently, a positive surplus reduces the amount extractable from both types, but leaves the information rent to be paid constant. Thus, by the reasoning above, if the surplus level is sufficiently high, firm R will start to ignore low-type workers and focus only on the high type. Because the surplus level can be continuously adjusted by the current employer C, he is still able to obtain strictly positive profits from low-type employees, when firm R begins to poach high-type workers exclusively. Hence, we have found a profitable strategy for firm C (c.f. Figure 1.1).

Definition 1. Denote with e_L^* and e_H^* the efficient effort levels for type θ_L and θ_H respectively, which can be derived from

$$f'(e_I^*) \quad = \quad c'(e_I^*, \theta_I), \qquad I \in \{L, H\}.$$

Lemma 1. *In any subgame perfect equilibrium firm C obtains strictly positive profits.*

Proof: Assume there exists a subgame perfect equilibrium in which firm C gets zero profits. We will show that firm C has a profitable deviation. Consider the following menu of firm C

$$\left((e_L^C, w_L^C), (e_H^C, w_H^C)\right) \quad = \quad \left((e_L^*, c(e_L^*, \theta_L) + s), (e_H^*, c(e_H^*, \theta_H) + s)\right),$$

where $s \geq 0$ is the employees' surplus from working at firm C.

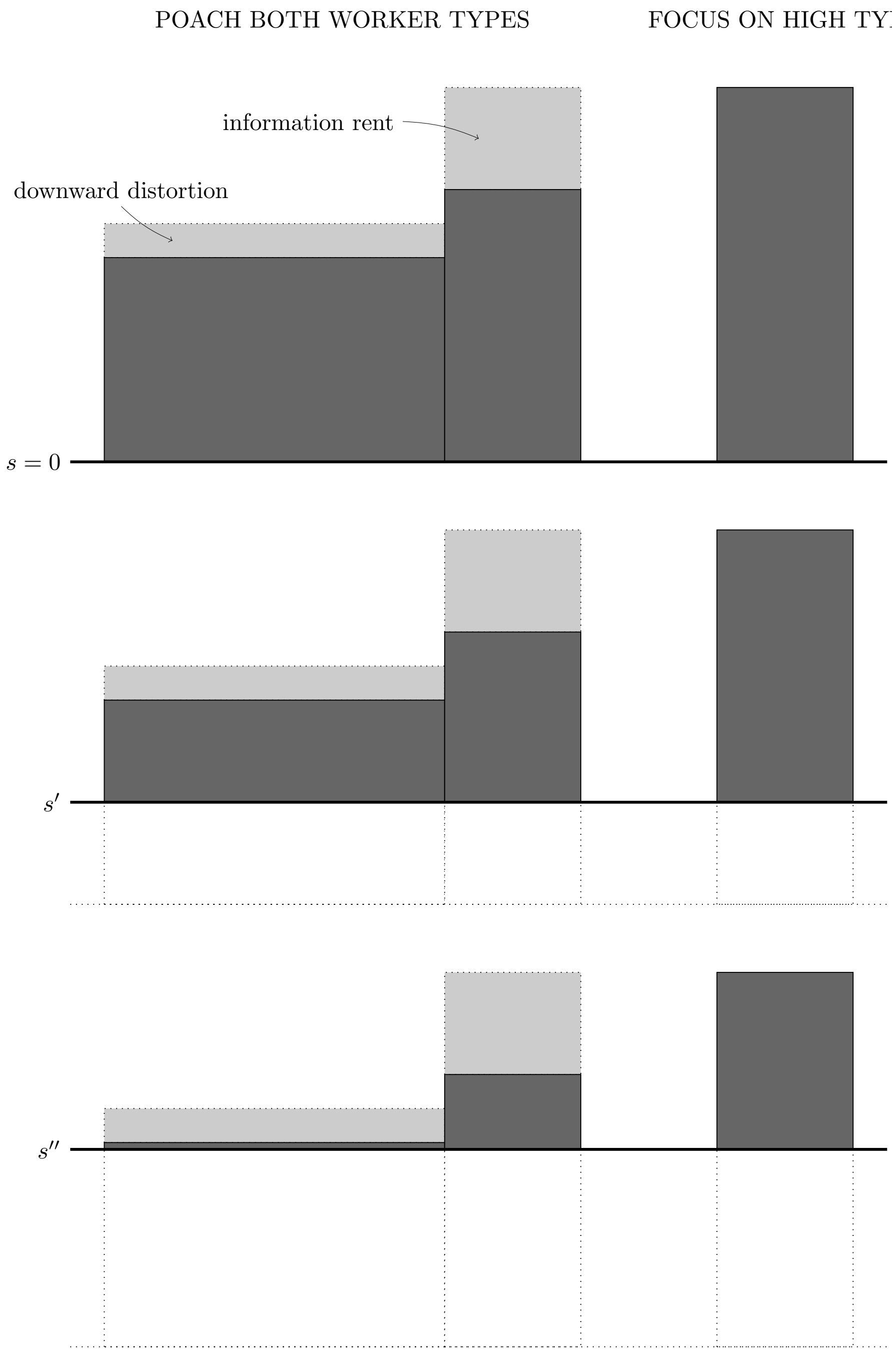

Figure 1.1: Profits (dark grey) of the raiding firm for attracting both worker types and solely the high type are depicted for different surplus levels (s, s', s''). Poaching both worker types involves a loss due to an effort downward distortion and information rent payments (light grey). For sufficiently high surplus levels the raiding firm prefers to focus on the high-type worker, even though profits with the low-type employee are still attainable.

Given this menu of firm C, if firm R is to poach both worker types, it has to solve the standard problem:

$$
\begin{aligned}
\max_{e_L, e_H, w_L, w_H} \quad & \alpha[f(e_L) - w_L] + (1-\alpha)[f(e_H) - w_H] \\
s.t. \qquad w_L - c(e_L, \theta_L) &\geq s \\
w_H - c(e_H, \theta_H) &\geq s \\
w_L - c(e_L, \theta_L) &\geq w_H - c(e_H, \theta_L) \\
w_H - c(e_H, \theta_H) &\geq w_L - c(e_L, \theta_H).
\end{aligned}
$$

By substituting $\tilde{w}_I = w_I - s$ the problem becomes

$$
\begin{aligned}
\max_{e_L, e_H, \tilde{w}_L, \tilde{w}_H} \quad & \alpha[f(e_L) - \tilde{w}_L] + (1-\alpha)[f(e_H) - \tilde{w}_H] \\
s.t. \qquad \tilde{w}_L - c(e_L, \theta_L) &\geq 0 \\
\tilde{w}_H - c(e_H, \theta_H) &\geq 0 \\
\tilde{w}_L - c(e_L, \theta_L) &\geq \tilde{w}_H - c(e_H, \theta_L) \\
\tilde{w}_H - c(e_H, \theta_H) &\geq \tilde{w}_L - c(e_L, \theta_H).
\end{aligned}
$$

Standard techniques from nonlinear pricing then show the following: the optimal effort levels $(\hat{e}_L, \hat{e}_H)$ satisfy

$$\hat{e}_L < e_L^*, \quad \hat{e}_H = e_H^*,$$

and optimal wages $(\hat{w}_L, \hat{w}_H)$ are given by

$$
\begin{aligned}
\hat{w}_L &= c(\hat{e}_L, \theta_L), \\
\hat{w}_H &= c(\hat{e}_H, \theta_H) + [\underbrace{c(\hat{e}_L, \theta_L) - c(\hat{e}_L, \theta_H)}_{\text{Inf.Rent}>0}].
\end{aligned}
$$

Hence, firm R's profit is given by

$$
\begin{aligned}
\hat{\pi}^R &= \alpha\left[f(\hat{e}_L) - c(\hat{e}_L, \theta_L) - s\right] + (1-\alpha)\left[f(\hat{e}_H) - \{c(\hat{e}_H, \theta_H) + c(\hat{e}_L, \theta_L) - c(\hat{e}_L, \theta_H)\} - s\right] \\
&= \alpha\left[f(\hat{e}_L) - c(\hat{e}_L, \theta_L) - s\right] - (1-\alpha)[c(\hat{e}_L, \theta_L) - c(\hat{e}_L, \theta_H)] \\
&\quad + (1-\alpha)\left[f(e_H^*) - c(e_H^*, \theta_H) - s\right].
\end{aligned}
$$

Note that the third term in the last equation is the profit firm R can guarantee herself if she decided to just focus on the high worker type. Moreover, as $s \nearrow f(\hat{e}_L, \theta_L) - c(\hat{e}_L)$ the first two expressions become in total strictly negative. Thus, define

$$\begin{aligned} \underline{s} &:= f(\hat{e}_L) - c(\hat{e}_L, \theta_L) - \frac{1-\alpha}{\alpha}[c(\hat{e}_L, \theta_L) - c(\hat{e}_L, \theta_H)], \\ \bar{s} &:= f(e_L^*) - c(e_L^*, \theta_L), \end{aligned}$$

and observe that $\underline{s} < \bar{s}$. Then for any firm C menu as defined above with $s \in (\max(\underline{s}, 0), \bar{s})$ firm R will respond with a menu that will not attract worker type θ_L, because $s > \underline{s}$. On the other hand the current employer attracts type θ_L (since $s > 0$) and still makes positive profits because $s < \bar{s}$. Hence, we have found a profitable deviation. □

This result shows how the current employer can avoid losing his full workforce to the raiding firm. Even though offering the same surplus to both worker types may not be optimal, it nicely illustrates how the current employer can secure strictly positive profits for himself. Once we have established that firm C receives strictly positive profits it is easy to see that firm R does so as well. As the second moving player there's always the possibility to overbid the offer of the first player. So if firm C made positive profits with one worker type, firm R could replicate firm C's offer with a slightly higher wage, win over this type to work for her and still extract positive value from this type.

Lemma 2. *In any subgame perfect equilibrium firm R obtains strictly positive profits.*

PROOF: Assume there exists an equilibrium in which firm R makes zero profits. We show that firm R has a profitable deviation. By Lemma 1 firm C obtains strictly positive profits, i.e. if (e^C, w^C) is her equilibrium menu, then there exists some $I \in \{L, H\}$ such that worker type I accepts firm C's contract and $f(e_I^C) - w_I^C > 0$. Now consider a new firm R offer on the equilibrium path with

$$\left((e_L^R, w_L^R), (e_H^R, w_H^R)\right) = \left((e_I^C, w_I^C + \epsilon), (e_I^C, w_I^C + \epsilon)\right),$$

where $\epsilon \in (0, f(e_I^C) - w_I^C)$. Given the equilibrium strategy of firm C, firm R will attract at least type θ_I and make strictly positive profits with her menu. Thus, she has a profitable deviation. □

If both firms attain strictly positive profits in a subgame perfect equilibrium, then it has to be the case that each firm employs exactly one worker. Otherwise one firm would obtain zero profits. We summarize this workforce segmentation effect in the following

Theorem 1 (Segmentation)**.** *In any subgame perfect equilibrium each firm gets strictly positive profits and employs exactly one type of worker.*

1.4.2 Profits and poached worker types

We now take a closer look at the equilibrium profits of the current employer and the raider. It is rather obvious that the current firm's profits cannot exceed those of the raider in a segmented market, because the latter could always slightly overbid the offer of his rival and secure almost all of her profits. The raider's ability to use menus of contracts even implies a gap between the two profit levels: If both were equal, the raiding firm could offer the equilibrium effort-wage combinations with a slight wage raise, so that both worker types chose a contract from firm R. For a sufficiently small raise it would not matter which specific contract was selected by each type. In any case firm R now extracts a higher value than before and has a profitable deviation.

Lemma 3. *In any subgame perfect equilibrium firm R obtains strictly higher profits than firm C.*

Proof: Assume by contradiction that $\pi^C \geq \pi^R$. By Theorem 1 both firms obtain strictly positive profits and both firms employ one worker type, say firm C employs worker type θ_I, whereas firm R employs θ_J on the equilibrium path with $I, J \in \{L, H\}$ and $I \neq J$. Denote with (e_I^C, w_I^C) and (e_J^R, w_J^R) the effort-wage combinations chosen on the equilibrium path by the workers. Because contract (e_J^R, w_J^R) was available to type θ_I but not chosen, it necessarily holds that

$$w_I^C - c(e_I^C, \theta_I) \quad \geq \quad w_J^R - c(e_J^R, \theta_I).$$

We show that firm R given firm C's menu has a profitable deviation. On the equilibrium path let firm R choose the new menu

$$\big((\hat{e}_I^R, \hat{w}_I^R), (\hat{e}_J^R, \hat{w}_J^R)\big) \quad := \quad \big((e_I^C, w_I^C + \epsilon), (e_J^R, w_J^R + \epsilon)\big)$$

with

$$\epsilon \ := \ \frac{1}{2}\min\big((1-\alpha_I)[f(e_I^C)-w_I^C],(1-\alpha_J)[f(e_J^C)-w_J^C],\pi^R\big) \ > \ 0.$$

Note that both worker types prefer working at firm R instead of choosing a contract from firm C or the outside option. There are four cases to consider: 1. both employee types choose contract $(\hat{e}_I^R, \hat{w}_I^R)$, 2. both choose $(\hat{e}_J^R, \hat{w}_J^R)$, 3. type θ_I selects $(\hat{e}_I^R, \hat{w}_I^R)$ and type θ_J selects $(\hat{e}_J^R, \hat{w}_J^R)$, or 4. the other way round.

In the first case firm R's new profits are

$$\begin{aligned}\hat{\pi}^R &= f(e_I^C) - w_I^C - \epsilon \\ &= \pi^C + (1-\alpha_I)[f(e_I^C) - w_I^C] - \epsilon \\ &> \pi^C \ \geq \ \pi^R.\end{aligned}$$

The second case is analogous. Note that in the third case firm R's new profits amount to $\pi^C + \pi^R - \epsilon > \pi^C \geq \pi^R$. And finally, in the last case firm R's profits are

$$\begin{aligned}\hat{\pi}^R &= \alpha_I\left[f(e_J^R) - w_J^R - \epsilon\right] + \alpha_J[f(e_I^C) - w_I^C - \epsilon] \\ &= \frac{\alpha_I}{\alpha_J}\pi^R + \frac{\alpha_J}{\alpha_I}\pi^C - \epsilon \\ &\geq \left(\frac{\alpha_I}{\alpha_J} + \frac{\alpha_J}{\alpha_I}\right)\pi^R - \epsilon \\ &= \left(\frac{\alpha_I^2 + 2\alpha_I\alpha_J + \alpha_J^2}{\alpha_I\alpha_J} - 2\right)\pi^R - \epsilon \\ &= \left(\frac{1}{\alpha_I\alpha_J} - 2\right)\pi^R - \epsilon \\ &\geq 2\pi^R - \epsilon \ > \ \pi^R.\end{aligned}$$

Thus, independent of workers' behavior, the raiding firm obtains strictly higher profits under this new menu, contradiction. □

The following result refines Theorem 1 by showing that for low values of α only high-type workers will be poached. It is based on an insight of the standard monopolistic screening model: if the share of low-type workers is too small, this group will be barred from employment, because their profit contribution is too insignificant to outweigh the loss through information rents for the high type. For our setup this observation implies that current employer profits are bounded below by the maximum extraction value of the low-type employee. If firm C offers a menu absorbing the whole surplus of each type, then, by construction, the raiding firm prefers to poach exclusively the high-type employee and firm C can always guarantee herself the entire low-type surplus. This effectively rules out that low-ability workers are raided in equilibrium. In that case firm R gained at most the full low-type surplus and, by Lemma 3, the current employer's profits laid strictly below this threshold. Thus, he had a profitable deviation.

Proposition 1. *Assume that the share of low-type workers α is small enough so that in the standard monopolistic screening problem only high types are employed.*[16] *Then, in any subgame perfect equilibrium, firm R (C) will employ the high (low) type.*

PROOF: Asssume there exists a subgame perfect equilibrium in which firm R employs the low type and firm C the high type. By Lemma 3 we must have $\pi^C < \pi^R \leq \alpha[f(e_L^*) - c(e_L^*, \theta_L)]$. Define $\epsilon := f(e_L^*) - c(e_L^*, \theta_L) - \pi^C/\alpha > 0$ and consider firm C menu

$$\left((\hat{e}_L^C, \hat{w}_L^C), (\hat{e}_H^C, \hat{w}_H^C)\right) = \left((e_L^*, c(e_L^*, \theta_L) + s), (e_H^*, c(e_H^*, \theta_H) + s)\right),$$

with $s \in (0, \epsilon)$. Given this offer, firm R's best response is to choose a menu that only attracts high-type workers: By assumption the raiding firm focuses on the high type for $s = 0$. And similar to the proof of Lemma 1, a lump sum wage increase s strengthens the effect to concentrate on the high type. Thus, because $s > 0$, firm C attracts all low-type workers and, since $s < \epsilon$, obtains a profit of

$$\begin{aligned} \hat{\pi}^C &= \alpha[f(e_L^*) - c(e_L^*, \theta_L) - s] \\ &> \pi^C. \end{aligned}$$

Hence, firm C has a profitable deviation, contradiction. □

For a small fraction of low-type workers, any equilibrium is characterized by the prediction that only high-type employees are poached by raiding firms. Lazear (1986) points out that in reality, typically, the most able workers switch jobs and builds a model that can account for this. For staff turnover to occur in his analysis, however, it is a necessary

[16] Such an interval of α always exists. To see this use the proof of Lemma 1 with $s = 0$ and compare firm R's profits from employing both worker types with those from poaching only the high type as $\alpha \searrow 0$.

condition that a raided worker has a higher firm-specific human capital at the new firm. But that may not always be the case: For instance, consider a sales representative who has switched to a rival employer but continues to market a very similar product. It is more likely that the employee's productivity is connected to his personal general skills, e.g. powers of persuasion, presentation of product, etc.) rather than firm-specific ones (mostly related to product characteristics). In our setup, at least for a certain parameter range, we can rationalize high type poaching in spite of no differences in worker productivity. Hence, we can provide a complementary explanation for the prevalence of high-skill turnover.

1.4.3 Defending mechanism

In this section we further explore how the current employer can counteract a complete loss of his workforce. In Lemma 1 we saw how firm C could guarantee herself positive profits by making it sufficiently unattractive for the raider to poach both employee types. It does not, however, provide us with further information on the specific structure of the equilibrium menu. The following result characterizes the equilibrium contract design of the current employer that prevents firm R from overbidding him and taking on both worker types.

Lemma 4. *Consider a subgame perfect equilibrium in which worker type θ_J is poached by firm R, and type θ_I remains at firm C. Then worker type θ_J strictly prefers the contract of type θ_I to her own, i.e.*

$$w_I^C - c(e_I^C, \theta_J) \quad > \quad w_J^R - c(e_J^R, \theta_J).$$

Proof: Consider such a subgame perfect equilibrium. By Theorem 1 both firms obtain strictly positive profits π^C and π^R respectively. We prove the statement by contradiction and assume that type θ_J weakly prefers her own contract from firm R to the contract of firm C offered to type θ_I, i.e.

$$w_I^C - c(e_I^C, \theta_J) \quad \leq \quad w_J^R - c(e_J^R, \theta_J).$$

We will show that firm R given firm C's menu has a profitable deviation. On the equilibrium path she now selects the menu

$$\big((\hat{e}_I^R, \hat{w}_I^R), (\hat{e}_J^R, \hat{w}_J^R)\big) \quad := \quad \big((e_I^C, w_I^C + \epsilon/2), (e_J^R, w_J^R + \epsilon)\big)$$

with

$$\epsilon \quad := \quad \frac{1}{2} \min\left((1-\alpha_J)[f(e_J^R) - w_J^R], \frac{\pi^C}{\alpha_J + \alpha_I/2}\right) \quad > \quad 0.$$

Note that both worker types prefer working at firm R instead of choosing the outside option or a contract from firm C. By the above inequality worker type θ_J will choose the bundle $(e_J^R, w_J^R + \epsilon)$ from firm R. Type θ_I will either select the same bundle as type θ_J or choose the bundle $(e_I^C, w_I^C + \epsilon/2)$ from firm R.

In the former case firm R's new profits are

$$\begin{aligned}
\hat{\pi}^R &= f(e_J^R) - w_J^R - \epsilon \\
&= \pi^R + (1-\alpha_J)[f(e_J^R) - w_J^R] - \epsilon \\
&> \pi^R.
\end{aligned}$$

In the latter case firm R's new profits are

$$\begin{aligned}
\hat{\pi}^R &= \alpha_J\left[f(e_J^R) - w_J^R - \epsilon\right] + \alpha_I\left[f(e_I^C) - w_I^C - \frac{\epsilon}{2}\right] \\
&= \pi^R + \pi^C - \left(\alpha_J + \frac{\alpha_I}{2}\right)\epsilon \\
&\geq \pi^R + \pi^C - \frac{\pi^C}{2} \\
&> \pi^R.
\end{aligned}$$

Thus, independent of worker type θ_I's behavior, the raiding firm obtains strictly higher profits under this new menu. This gives a contradiction. □

This result is very helpful to understand the threat that the current employer conveys with his contract design: If the raiding firm attempts to poach both worker types, she has to overbid firm C's offer. But then she will not only attract her rival's workers with this contract, but also the worker types she had previously poached with a different contract. In equilibrium, firm C's contract provides a relatively high surplus to the worker, and the raiding firm will restrain herself to overbid it, because the loss due to the lower profit extraction is not compensated by the effect of employing both worker types.

The current employer knows that employee poaching cannot be entirely deterred, unless he's willing to make zero profits. In this view attractive offers to some workers in anticpation of an employee raid receive an extended interpretation. Their purpose is not only to keep these workers profitably within the company. They rather prevent the raider from poaching all worker types, because overbidding the high-surplus offers of some employees would induce all employees to select these contracts. To our knowledge this role of contract design in order to partially defend employee raids is novel.

1.4.4 Effort levels

In this last part on necessary equilibrium conditions we analyze the effort levels exerted by workers in equilibrium. Because each firm employs exactly one worker type, intuition suggests that there's no need for a distortion in the task assignment and firms can extract the highest profits by choosing efficient effort levels. This idea turns out to be true. For the raiding firm as the last mover this result is rather immediate. Contracts with inefficient effort choices on the equilibrium path can always be replaced with an efficient effort-wage combination and generates higher profits.

Lemma 5. *In any subgame perfect equilibrium, the worker type poached by firm R selects a contract with efficient effort level.*

Proof: Let firm R poach worker type θ_J with contract (e_J^R, w_J^R) on the equilibrium path. Then it must be true that

$$w_J^R - c(e_J^R, \theta_J) \quad \geq \quad \max(0, w_J^C - c(e_J^C, \theta_J)).$$

This equation has to hold with equality in equilibrium—otherwise firm R could improve her profits by slightly lowering her wage and still attract type θ_J. Now assume $e_J^R \neq e_J^*$. We will show that firm R has a profitable deviation. Since by definition $f(e_J^*) - c(e_J^*, \theta_J) - [f(e_J^R) - c(e_J^R, \theta_J)] > 0$, there exists $\epsilon \in \left(c(e_J^*, \theta_J) - c(e_J^R, \theta_J), f(e_J^*) - f(e_J^R)\right)$. Consider the following deviation menu on the equilibrium path for firm R:

$$\left((\hat{e}_I^R, \hat{w}_I^R), (\hat{e}_J^R, \hat{w}_J^R)\right) \quad = \quad \left((e_J^*, w_J^R + \epsilon), (e_J^*, w_J^R + \epsilon)\right).$$

Under this new menu at least type θ_J will work for firm R, because

$$\begin{aligned}\hat{w}_J^R - c(\hat{e}_J^R, \theta_J) &= w_J^R + \epsilon - c(e_J^*, \theta_J) \\ &> w_J^R - c(e_J^R, \theta_J) \\ &= \max\left(0, w_J^C - c(e_J^C, \theta_J)\right).\end{aligned}$$

Since at least worker type θ_J is attracted by this firm R contract, we can bound the profits of this menu as follows:

$$\begin{aligned}\hat{\pi}^R &\geq \alpha_J[f(\hat{e}_J^R) - \hat{w}_J^R] \\ &= \alpha_J[f(e_J^*) - w_J^R - \epsilon] \\ &> \alpha_J[f(e_J^R) - w_J^R] \\ &= \pi^R.\end{aligned}$$

Thus, firm R has a profitable deviation, contradiction. □

A similar result can be established for firm C and the proof follows the same logic. It is, however, more involved, because the current employer has to anticipate how firm R will respond when he offers a different menu. The following lemma shows that the current employer can always make a transition to a menu with an efficient effort level and increase his profit without changing the qualitative behavior of the raiding firm.

Lemma 6. *In any subgame perfect equilibrium the worker type employed by firm C selects a contract with efficient effort level.*

Proof: Because the proof relies on notation introduced in the next subsection, it is relegated to the appendix. □

Both worker types are employed at equally productive firms and there are no poaching costs in our model. This implies with Lemma 5 and 6 that each equilibrium outcome is efficient—equilibrium wages determine the distribution, but not the size of the total interaction surplus. We will see in Section 1.5 that this result does not carry over to the case of asymmetric productivity.

1.4.5 Equilibrium existence

In this rather technical part we address the question whether a subgame perfect equilibrium exists in this game. We start by introducing some more notation.

The surplus attained by worker type θ_L and θ_H if they accept the contract from firm C's menu (e^C, w^C) is denoted by s_L and s_H respectively, i.e. $s_I := w_I^C - c(e_I^C, \theta_I)$ for $I \in \{L, H\}$.[17] Similarly, we define the maximum surplus that can be generated in a firm by employing worker type θ_I by s_I^*, i.e. $s_I^* := f(e_I^*) - c(e_I^*, \theta_I)$ for $I \in \{L, H\}$.

Finally, the function $V(s_L, s_H)$ denotes the maximum profit of the raiding firm R if she attracts both worker types and workers receive a surplus of s_L and s_H, respectively, from working at firm C, i.e.

$$\begin{aligned} V(s_L, s_H) \quad := \quad & \max_{e,w} \ \alpha[f(e_L) - w_L] + (1-\alpha)[f(e_H) - w_H] \\ s.t. \quad & w_L - c(e_L, \theta_L) \ \geq s_L \\ & w_H - c(e_H, \theta_H) \ \geq s_H \\ & w_L - c(e_L, \theta_L) \ \geq w_H - c(e_H, \theta_L) \\ & w_H - c(e_H, \theta_H) \ \geq w_L - c(e_L, \theta_H). \end{aligned}$$

Endowed with this notation, we need to define two problems: Problem L is designed to implement that firm C extracts from the low-type employee the highest profit under the constraint that firm R prefers poaching the high-type worker to poaching the low type or both types. A similar definition applies to Problem H, where firm C employs the high type.

We know that both firms offer their prospective workers equilibrium contracts with efficient effort levels by Lemma 5 and 6. Thus, for our purpose we will restrict ourselves to these sorts of menus. But then there's a one-to-one mapping between firm C's chosen wage and surplus levels.[18] Moreover, the choice of s_L and s_H fully determines both firm's profits. Consider first the case when one worker type is employed: for firm C it is directly the difference between maximum surplus and chosen surplus of this type. But the same

[17] We drop the explicit dependence of s_L and s_H on (e^C, w^C) as long as no confusion can arise.

[18] The optimal wages can be retrieved as the sum of optimal surplus and the effort cost at the efficient level for the corresponding type.

is true for firm R, because she will extract the whole employee value up to the surplus provided by firm C. Second, if firm R poaches both both worker types, her profits are given by $V(s_L, s_H)$. This allows us to represent Problem I with $I \in \{L, H\}$ by

Problem I

$$\begin{array}{lrcl} \max\limits_{s_I, s_J \geq 0} & \alpha_I[s_I^* - s_I] & & \\ & \alpha_J[s_J^* - s_J] & \geq & V(s_I, s_J) \\ & \alpha_J[s_J^* - s_J] & \geq & \alpha_I[s_I^* - s_I]. \end{array}$$

The idea of the equilibrium construction is as follows: We will first show that for both maximization problems a solution exists (not necessarily unique), which determines the value of each problem. Then we will select a maximizer from the problem with the higher value (if both values are equal any maximizer will be fine). The menu associated with this maximizer will be our equilibrium candidate strategy for firm C. We will establish a further technical result before we finally verify that a subgame perfect equilibrium can always be constructed with this candidate strategy.

Lemma 7. *The function $V(s_L, s_H)$ is continuous on $S := [0, s_L^*] \times [0, s_H^*]$.*

Proof: The proof is relegated to the appendix. □

Lemma 8. *Problem L and Problem H have a solution.*

Proof: Consider Problem I with $I \in \{L, H\}$. Note that firm C can guarantee herself zero profits by selecting $(s_I, s_J) = (s_I^*, 0)$. This choice lies in the constraint set, because $\alpha_J s_J^* \geq V(s_I^*, 0)$ and $\alpha_J s_J^* > 0$. The first inequality is a consequence of the additional incentive constraints firm R has to respect when she intends to attract both types, but cannot make any profits with type θ_I.

This implies, first, that the constraint set is nonempty. Second, we can restrict attention to $s_I \leq s_I^*$, because a higher surplus would incur a loss for firm C. By looking at the second constraint of the maximization problem one concludes that also $s_J \leq s_J^*$.

Summing up, the solution to the maximization problem above coincides with the solution of this problem under the additional constraint $(s_I, s_J) \in S := [0, s_I^*] \times [0, s_J^*]$. Due to this extra restriction the constraint set is bounded. Weak inequalities and, by Lemma 7, continuity of the constraints in (s_I, s_J) on S imply it is also closed. Hence, we maximize a continuous objective function over a nonempty and compact set. By the Extreme Value Theorem this problem has a solution. □

Definition 2. Denote with (s_L^L, s_H^L) and (s_L^H, s_H^H) solutions to Problem L and Problem H respectively. The associated values of these problems are defined by W^L and W^H.

Finally, we state a necessary condition for a solution of the maximization problem. It basically says that the worker type that stays at firm C strictly prefers her own contract to the contract accepted by the poached type. In case of indifference firm R could slightly increase the wage of her offered contract and attract both types profitably. Thus, poaching both types was more attractive than poaching just a single type and the first constraint of Problem I was violated.[19] This result will be needed to show that the worker type staying at firm C does not have a profitable deviation in our equilibrium construction.

Lemma 9. *Let $W^I > 0$ and (s_I^I, s_J^I) be a solution of Problem I with $I \in \{L, H\}$. Then*

$$s_I^I \quad > \quad s_J^I - [c(e_J^*, \theta_I) - c(e_J^*, \theta_J)].$$

PROOF: First, note that the second constraint of Problem I together with $W^I > 0$ implies that $s_J^* - s_J^I > 0$. Now assume $s_I^I \leq s_J^I - [c(e_J^*, \theta_I) - c(e_J^*, \theta_J)]$. If firm R intends to poach both worker types and offers contract $((e_I, w_I), (e_J, w_J)) = \big((e_J^*, c(e_J^*, \theta_J) + s_J^I), (e_J^*, c(e_J^*, \theta_J) + s_J^I)\big)$, it is easy to check that both individual rationality constraints are satisfied and the incentive constraints hold trivially. This establishes a lower bound for firm R's profits when she poaches both worker types:

$$\begin{aligned} V(s_I^I, s_J^I) &\geq f(e_J^*) - [c(e_J^*, \theta_J) + s_J^I] \\ &= s_J^* - s_J^I \\ &> \alpha_J [s_J^* - s_J^I]. \end{aligned}$$

This violates the first constraint of Problem I, contradiction. □

We're now ready to state the main result of this section.

Theorem 2. *There exists a subgame perfect equilibrium.*

[19] The statement is complementary to Lemma 4, which considers the poached type's preferences regarding equilibrium contracts. The similarity is hard to recognize due to different notation.

Proof: Without loss of generality assume $W^I \geq W^J$ with $I, J \in \{L, H\}, I \neq J$. By Lemma 1 there always exists a strategy to attain positive profits with the low worker type, so that $W^L > 0$, and hence $W^I > 0$. Let

$$\begin{aligned}\left((e_I^C, w_I^C), (e_J^C, w_J^C)\right) &= \left((e_I^*, c(e_I^*, \theta_I) + s_I^I), (e_J^*, c(e_J^*, \theta_J) + s_J^I)\right), \\ \left((e_I^R, w_I^R), (e_J^R, w_J^R)\right) &= \left((e_J^*, c(e_J^*, \theta_J) + s_J^I), (e_J^*, c(e_J^*, \theta_J) + s_J^I)\right)\end{aligned}$$

be the firms' menus on the equilibrium path and type θ_I accepts her contract from firm C and type θ_J chooses one of the (identical) contracts from firm R.

Off the equilibrium path, both worker types are assumed to choose the alternative with the highest utility given menus $\left((e^C, w^C), (e^R, w^R)\right)$: 1. contract designed for her by firm R, 2. contract designed for other type by firm R, 3. contract designed for her by firm C, 4. outside option. In case of any indifference between the four options, a worker chooses the option with the lowest number. In particular, she always prefers a contract of firm R to one of firm C when indifferent.

Similarly, off the equilibrium path, given any menu (e^C, w^C), firm R is assumed to select the alternative with the highest profits: 1. poach worker θ_J, 2. poach both worker types, 3. poach worker θ_I, 4. poach no worker. Naturally, the profit of each alternative of firm R depends on surplus (s_L, s_H) of contract (e^C, w^C) and the workers' choices. Note that by the indifference rule above there always exists a maximum value for each option. In case of any indifference between the four options, firm R chooses the option with the lowest number.

We will now check whether one player has a profitable deviation after any history of the game. By construction both worker types and firm R behave optimally off the equilibrium path. On the equilibrium path type θ_I obtains a utility of $s_I^I \geq 0$ if she adheres to her strategy. Deviating to firm R's contract will give her a payoff of $c(e_J^*, \theta_J) + s_J^I - c(e_J^*, \theta_I)$. Since $W^I > 0$, by Lemma 9, this is strictly lower than s_I^I and she has no incentive to deviate. Similarly, worker type θ_J obtains a payoff of $s_J^I \geq 0$ at both firms and behaves optimally. If firm R sticks to her strategy on the equilibrium path, she will obtain strictly positive profits with worker type θ_J. This follows from $\alpha_I[s_I^* - s_I^I] = W^I > 0$ and the second constraint of Problem I. Thus, poaching no worker type and getting zero profits is not a profitable deviation. Moreover, because (s_I^I, s_J^I) is the solution to Problem I, raiding worker type θ_J is weakly better than poaching type θ_I or both types. Finally, note that by choosing the efficient effort level and extracting every surplus up to s_J^I firm R cannot poach worker type θ_J in a more profitable way. Hence, firm R cannot improve on her equilibrium behavior. It remains to be checked that firm C has no profitable deviation. If she adheres to her strategy, she obtains profits $W^I > 0$ and keeps type θ_I in her company. Clearly, firm C cannot do better with contracts that will leave her with both worker types or no workers at all. In either case she attains at most zero profits. Firm C cannot improve if she offers a menu that will keep worker type θ_J exclusively and firm R focuses on type θ_I. To see this, note that firm C can receive with this method profits of at most $W^J \leq W^I$. And finally, because firm C chooses the efficient task level and extracts the maximum surplus from this type for given s_I^I, there's no room for improvement by rearranging her effort-wage schedule. Hence, firm C has no profitable deviation and the proposed strategy profile forms a subgame perfect equilibrium. □

1.5 Extensions

We now turn to two modifications of our model. Our first extension is concerned with asymmetric productivity of workers at firms. We study the qualitative nature of equilibrium outcomes and their efficiency in dependence on productivity differences.

Second, in the basic setup the current employer committed to not match any offers from the raiding firm. In another extension we endogenize the choice of the current employer whether to make counteroffers or not.

1.5.1 Asymmetric firm productivity

Up to this point, strategies of the current employer that prevent the raiding firm from poaching were of little interest to him. In order to hold down firm R to zero profits, firm C would have to hand over the maximum extractable surplus s_L^* and s_H^* to both worker types, thereby attaining zero profits as well. But by letting some of his workers leave the company, the current employer can secure strictly positive profits with workers that stay.

This rationale heavily relies on the fact that both firms have identical production functions. In general, one would expect that firms' technologies differ and one firm can use workers more efficiently than the other. So, let's assume that firms' production functions are given by $f^C(e) = f(e)$ and $f^R(e) = f(e)/\beta$. For the moment let $\beta \geq 1$, i.e. the current employer has a more productive use for his workers. We will proceed as follows: We start by generalizing some earlier notation in order to define profits of firm C if she decides to prevent the raiding firm from poaching. We then compare these profits with those if poaching is permitted. It turns out that for large productivity differences raiding is deterred, but cannot be contained for small ones and we discuss this result in terms of efficiency. At the end of this section we will briefly consider the case $0 < \beta < 1$.

First, we extend some notation we introduced earlier: Denote with $e_I^*(\beta)$ the efficient effort level for type θ_I for any $\beta > 0$, implicitly defined by $f'(e_I^*(\beta)) = \beta c'(e_I^*(\beta), \theta_I)$ for $I \in \{L, H\}$. Moreover, define with $s_I^*(\beta)$ the maximum surplus extractable from worker type θ_I for any $\beta > 0$, i.e. $s_I^*(\beta) = \frac{1}{\beta} f(e_I^*(\beta)) - c(e_I^*(\beta), \theta_I)$ for $I \in \{L, H\}$. Some properties follow immediately from these definitions: Obviously, for $I \in \{L, H\}$, we have $e_I^*(1) = e_I^*$ and $\lim_{\beta \to \infty} e_I^*(\beta) = 0$. Furthermore, $s_I^*(1) = s_I^*$ and $\lim_{\beta \to \infty} s_I^*(\beta) = 0$.

We now pin down the current employer's profit when he wishes to deter firm R from poaching any employees. Note that he has to provide necessarily both worker types with the maximum surplus $(s_L^*(\beta), s_H^*(\beta))$ they could generate in firm R. If he endowed any worker type with a surplus strictly below this threshold, the raiding firm can profitably poach at least one worker type. On the other hand, he provides both employee types with efficient effort levels, because he can identify each of them and need not distort effort to accrue a higher gain. Thus, his no-raiding profit is simply the difference of the maximum surplus in his firm and firm R for each type:

$$\pi_{nr}^C(\beta) \quad := \quad \alpha[s_L^* - s_L^*(\beta)] + (1 - \alpha)[s_H^* - s_H^*(\beta)].$$

Note that $\pi_{nr}^C(\beta)$ starts at $\pi_{nr}^C(1) = 0$, is continuous and strictly increasing in β with limit $\alpha s_L^* + (1 - \alpha) s_H^*$.

This profit has to be compared with $\pi^C(\beta)$, the profit when firm C allows some of her employees to be raided. The relative size of $\pi_{nr}^C(\beta)$ and $\pi^C(\beta)$ determine whether poaching activities will happen or not. We will see that for small β employee raiding will always take place, while there will be no raiding for sufficiently high β. We show this result by establishing two bounds for the current employer profits $\pi^C(\beta)$, when poaching is permitted: a lower bound strictly above zero and an upper bound strictly below the maximum surplus, a monopsonic firm C would extract from both worker types. The result then immediately follows due to the qualitative properties of $\pi^C_{nr}(\beta)$.

Finding an upper bound for the profit function $\pi^C(\beta)$ if employee poaching occurs is easy. Because we know firm C to keep at most one worker type, her profits are bounded by $\bar{\pi}^C := \max(\alpha s^*_L, (1-\alpha)s^*_H)$. Proving the existence of a strictly positive lower bound for the function $\pi^C(\beta)$ is a bit harder. The idea of the proof uses an insight from Lemma 1: There exists an interval of surplus levels for workers that induces the raiding firm to poach exclusively high-type employees and allows firm C to obtain positive profits with the low type. The following lemma shows that this interval is getting larger as β grows, so that any of the profits from the $\beta = 1$ case is a lower bound for the function $\pi^C(\beta)$. Intuitively, all what is needed to induce firm R to poach solely the high type is to push the surplus levels of both types high enough, so that the gain from the low-type employee lies below the information rent for the high type. But a decrease in productivity has a first-order effect on the maximum surplus of the low type and only a second-order effect on the information rent. Thus, firm R will begin to focus exclusively on high types for increasingly smaller surplus levels as productivity falls.

Lemma 10. *There exists $0 < \underline{\pi}^C < \bar{\pi}^C$ such that $\pi^C(\beta) \geq \underline{\pi}^C$ for all $\beta \geq 1$.*

PROOF: The proof is relegated to the Appendix. □

It is now easy to prove the following

Theorem 3. *There exist $1 < \underline{\beta} < \bar{\beta}$ such that in any subgame perfect equilibrium with*

- *$\beta \in [1, \underline{\beta})$ employee raiding will occur,*
- *$\beta \in (\bar{\beta}, \infty)$ no raiding will occur.*

PROOF: We know from Lemma 10 and the preceding discussion that $0 < \underline{\pi}^C < \bar{\pi}^C < \alpha s^*_L + (1-\alpha)s^*_H$ and the profit function if firm C permits poaching satisfies $\pi^C(\beta) \in [\underline{\pi}^C, \bar{\pi}^C]$. Moreover, the profit function $\pi^C_{nr}(\beta)$, when firm C suppresses raiding activities of firm R, is continuous, strictly increasing and has the properties $\pi^C_{nr}(1) = 0$ and $\lim_{\beta\to\infty} \pi^C_{nr}(\beta) = \alpha s^*_L + (1-\alpha)s^*_H$. Hence, by the Intermediate Value Theorem there exist $1 < \underline{\beta} < \bar{\beta}$ such that $\pi^C(\beta) > \pi^C_{nr}(\beta)$ for all $\beta \in [1, \underline{\beta})$ and $\pi^C(\beta) < \pi^C_{nr}(\beta)$ for all $\beta \in (\bar{\beta}, \infty)$. □

In order to deter firm R from raiding any of her workers, firm C has to compensate each employee type with an additional surplus that corresponds to the maximum amount firm

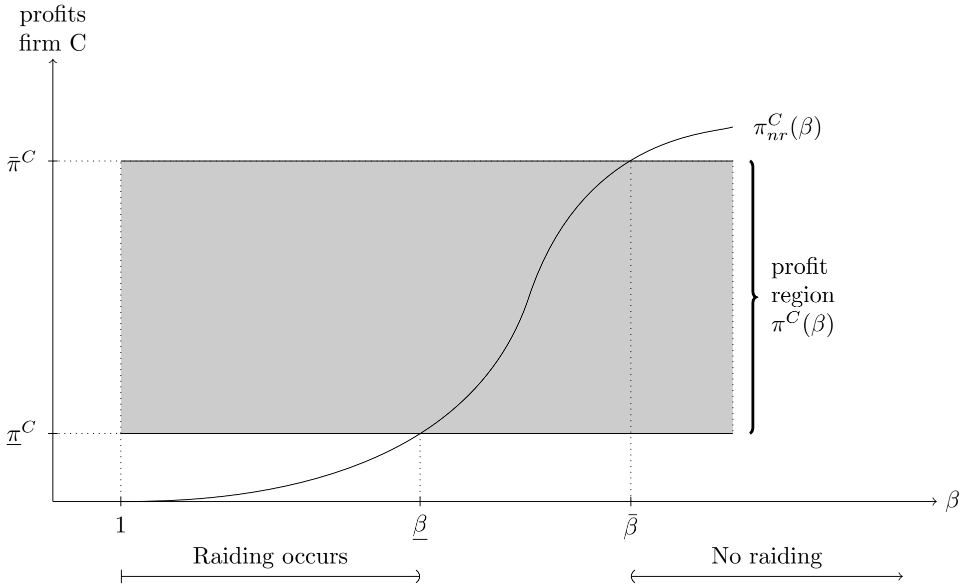

Figure 1.2: Two different firm C profits (depending on firm R productivity parameter β) are depicted: If firm C intends to prevent raiding activities of her opponent, she attains a profit of $\pi_{nr}^C(\beta)$. Profits $\pi^C(\beta)$, when firm C permits raiding, can be estimated within a lower bound $\underline{\pi}^C$ and an upper bound $\bar{\pi}^C$. The relative position of both profit functions shows that raiding activities will occur for $1 \leq \beta < \underline{\beta}$, but not for $\beta > \bar{\beta}$. No conclusion can be drawn for β in between the two productivity cutoffs.

R can extract from this type. This is very costly for firm C for relatively low values of β, because workers are almost identically productive in both companies and little profits can be obtained from keeping both types. Thus, she rather prefers to lose some of her workers and make a considerable gain with those employees that stay at the company. The opposite is true for relatively high values of β. Keeping both worker types in the company is of little cost for firm C, because workers are very unproductive at the other firm. In this case a substantial fraction of surplus can be extracted from both types. But if firm C permits her opponent to raid one of her employee groups, she can get at most the maximum surplus from only one worker type. Preventing an employee raid is now the more profitable alternative. For average differences in productivity no clear prediction can be made. Figure 1.2 illustrates the result graphically.

Lemma 5 and 6 show that in equilibrium both firms choose contracts with efficient effort levels. Thus, for $\beta = 1$, it does not matter, in terms of efficiency, at which firm workers are employed. Things change, however, for values of the productivity parameter above unity.

As the current employer can use both worker types more productively than his opponent, efficiency requires that all employees stay at his company. In fact, since he knows both worker types, there is no need for effort distortion and he will employ them under efficient effort provision. Theorem 3 shows that for small $\beta > 1$ there's a misallocation of workers to firms. The employer's incentive to reap additional profits, make him unresponsive, when parts of his workforce are raided by the rival firm. In the intermediate range, we have no clear prediction for the equilibrium outcome. But when productivity at the raiding firm is very low, efficiency is reinstalled, because the current employer prefers to endow both worker types with sufficiently high wages to keep them both within his company.

Finally, we give a verbal description of the case when the current employer is relatively unproductive compared to the raiding firm, i.e. $0 < \beta < 1$. For productivity values close to unity, firm C can still induce firm R to poach exclusively high-type workers (as in Lemma 1), while profitably keeping the low types. Intuitively, note that the maximal surplus level for which firm C can still profitably employ low types has just marginally dropped. But the threshold surplus, which induces firm R to poach exclusively high-type workers, lies strictly below her maximum surplus level of the low type. Thus, the idea of Lemma 1 still goes through. For lower values of β we have in general no clear prediction whether partial or full poaching occurs. There is, however, a special case where a clear result can be maintained: Let the share of low-type workers be so small that this group will be excluded in the standard monopolistic screening problem. Then, even for values of β very close to zero, the current employer will still be able to profitably keep low-type employees in his company. He can simply offer full extraction contracts to both employee types and the raiding firm R will solely focus on the high type. Thus, partial poaching prevails in this case for all $0 < \beta < 1$.

Efficiency is an even more intricate issue. Whenever partial raiding occurs the outcome is inefficient, because both employee types should work for the raiding firm. If there is full poaching, however, inefficiencies may also arise because the raiding firm tends to distort effort, when she employs both worker types. Since we cannot rule out equilibria where firm C offers contracts that are responded by the raider with efficient effort levels to both worker types, we do not further address these problems here.

1.5.2 Endogenous offer-matching

We extensively analyzed the model under the assumption that the current employer does not match any raiding offers. He anticipates, however, poaching activities and chooses the effort and wage structure in order to minimize the damage to his firm. Yet, in some industries, like the financial sector, counteroffers are quite common. Even though we do not want to provide an explanation why offer-matching is widespread in some industries and in others not, we feel that it is within the firm's discretion which policy to follow. For that reason, we introduce before the raiding game an additional stage, where the current employer can decide, whether to make counteroffers or not.

In order to pin down the current employer's equilibrium policy, his profits under offer-matching have to be determined. Our main focus is on the situation when rival offers are unobservable. In this case a firm cannot respond to her opponent's contract by slightly overbidding it and she effectively has to pay, what her own contract prescribes. We will therefore use the First Price Auction as a model for the raiding process with counteroffers.

Despite having an informational advantage, the current employer does not attain positive profits in the First Price Auction. His ability to overbid the raiding firm's offer for each type separately, without danger that his employees switch contracts, holds firm R's gain down to zero. But this, in turn, implies that his own profits cannot be strictly positive, because the raiding firm can always specialize in poaching one of the current employer's profitable worker types.

Proposition 2. *In any equilibrium of the First Price Auction both firms obtain zero profits.*

Proof: Assume firm R obtains strictly positive profits π^R in an equilibrium of the First Price Auction. Then there exists for some $J \in \{L, H\}$ a worker type θ_J that profitably accepts contract (e_J^R, w_J^R). An overbidding argument immediately establishes that firm C also obtains strictly positive profits π^C and necessarily keeps worker type $\theta_I \neq \theta_J$ with some contract (e_I^C, w_I^C) in her company. We will show that firm C menu

$$\big((\hat{e}_I^C, \hat{w}_I^C), (\hat{e}_J^C, \hat{w}_J^C)\big) \quad := \quad \big((e_I^C, w_I^C + \epsilon), (e_J^R, w_J^R + \epsilon)\big)\,,$$

with $\epsilon \in (0, \pi^R)$, is a profitable deviation for her.

Note that by offering menu $(\hat{e}^C, \hat{w}^C)$, both worker types strictly prefer to stay at firm C, and select the contract designed for them. Thus, under the First Price Auction rule, firm C obtains new profits

$$\begin{aligned} \hat{\pi}^C &= \alpha_I[f(e_I^C) - w_I^C] + \alpha_J[f(e_J^C) - w_J^C] - \epsilon \\ &= \pi^C + \pi^R - \epsilon \\ &> \pi^C. \end{aligned}$$

This yields a contradiction. Hence, in any equilibrium of the First Price Auction firm R necessarily obtains zero profits.

Now assume firm C receives strictly positive profits in equilibrium and firm R gets zero profits. So there exists for some $I \in \{L, H\}$ a worker type θ_I that accepts contract (e_I^C, w_I^C) and $f(e_I^C) - w_I^C > 0$. But if firm R offers menu

$$\left((\hat{e}_I^R, \hat{w}_I^R), (\hat{e}_J^R, \hat{w}_J^R)\right) \quad := \quad \left((e_I^C, w_I^C + \delta), (e_I^C, w_I^C + \delta)\right),$$

with $\delta := (f(e_I^C) - w_I^C)/2$, she attracts at least worker type θ_I. Thus, we can bound her profits from this deviation by

$$\hat{\pi}^R \;\geq\; \alpha_I[f(e_I^C) - w_I^C - \delta] \;>\; 0,$$

and firm R has a profitable deviation. □

Finally, we briefly consider the case, when rival offers are observable. In such a situation solicited employees can swing hence and forth between the firms with new counteroffers until one firm stops the bidding process. Thus, the ascending English auction is the most salient modeling alternative. If the worker's type is revealed in such an auction, firms can overbid their rival's offer in neglectably small steps, and this mechanism can be represented by a Second Price Auction. Even though multiple equilibria exist for this auction type, we believe that the equilibrium in weakly dominant strategies in which both firms bid for each type their willingness to pay, is the most persuasive solution of the game. Both firms obtain zero profits in this equilibrium. More subtle issues arise, however, if no type revelation occurs in the bidding process and we do not investigate this case. Thus, assuming that this last eventuality does not occur, the current employer's profits are zero, whenever counteroffers, observable or not, can be made. For that reason, he will commit to not match any offers from outside firms and secure strictly positive profits by permitting partial raids. We summarize this result in the following

Theorem 4. *In any subgame perfect equilibrium firm C commits to no offer-matching. Firm C and R obtain strictly positive profit and precisely one worker type will stay in firm C, wheras the other type will be raided by firm R.*

1.6 Conclusion

In our model a firm employs two types of workers, which differ in their cost for exerting effort and whose type is known to the employer. An outside firm tries to poach his employees, but is constrained by her inability to distinguish the workers' types and knows only their distribution. She can, however, use menus of contracts in order to attract both worker types to her firm. If the current employer has committed to not match any outside offers, can he choose effort levels and wages in a suitable way to prevent a total loss of his workforce to the raider?

We show that the current employer can always induce the raiding firm to poach exclusively high-type workers and make low-type employees stay in his company. The intuition of this result is based on information rent payments, which the raiding firm has to leave to the high-type worker, when she wants to hire both employee groups. This is, however, not the case if she specializes on the high type. By making the low-type group sufficiently unattractive through surplus payments, the current employer can increase the relative burden of the information rent so high that the raiding firm focuses solely on the high type. As a consequence, in any equilibrium partial raiding occurs (current employer and raider employ exactly one type) and both firms obtain strictly positive profits (raider's profits are strictly higher). In addition, for small shares of low-type workers poached employees are always high types. In equilibrium, the current employer deters poaching of his remaining employees by selecting a peculiar contract structure. If the raiding firm slightly overbids his contract, not only the current employer's workers will select this new offer, but also employees she already attracted with a different contract. In general, the additional gain from employing both worker types cannot compensate the loss of high extraction profits from employees she already had. Finally, we establish that an equilibrium exists.

In any equilibrium, both worker types select contracts with efficient effort levels. Thus, when employees are equally productive in both firms, the equilibrium outcome is efficient. Under small differences in productivity, however, poaching activities still occur and lead to inefficiencies. If labor is highly unproductive in the raiding firm, the current employer manages to suppress raiding activities and efficiency is reinstalled.

At last, we endogenize the current employer's decision, whether to make counteroffers or not. We find that the ability to match offers makes him worse off and, hence, in any subgame perfect equilibrium, he will commit to a policy of not matching offers.

There are some remaining issues in this paper that have not been entirely covered by us. First of all, we show in Proposition 1 that in any equilibrium only high-type workers will be subject to a raid if the share of low-type workers is small in the population. For arbitrary shares of low-type employees we do not provide an equilibrium characterization of the poached type. Similarly, we cannot determine which workers receive a positive surplus in equilibrium. Do raided employees get a surplus, or their peers that stay in the company? In both cases some technical challenges remain that need to be overcome in order to progress in these directions.[20] Finally, we restrict the analysis to only two worker types. For more employee types the logic of Lemma 1 still holds true and the current employer can always guarantee himself profits from workers of the lowest type. Thus, even though we will also obtain a partial poaching result with strictly positive profits for both firms, there seems to be room for further interaction. For instance, is there a threshold type, above or below which poaching occurs or has the set of raided types less structure? And what is the role of bunching from the standard monopolistic screening problem in this setup?

We leave these questions open for further research.

[20] In particular, the main obstacle is to obtain more information about firm R's profit function $V(s_L, s_H)$ if both types are poached. Unfortunately, $V(s_L, s_H)$ is not well-behaved: it is piece-wise defined and, even for simplifying examples, a rather complicated object.

1.7 Appendix

PROOF OF LEMMA 6: Assume that on the equilibrium path firm C offers menu (e^C, w^C), employs worker type θ_I, but $e_I^C \neq e_I^*$. Then we can define

$$\delta \;:=\; f(e_I^*) - c(e_I^*, \theta_I) - \left[f(e_I^C) - c(e_I^C, \theta_I)\right] \;>\; 0.$$

Let (s_I, s_J) be the surplus for each worker type associated with the menu of firm C, i.e. the difference between the wage and the effort cost when working for firm C. Hence, equilibrium profits of firm C are given by

$$\pi^C \;=\; \alpha_I \left[f(e_I^C) - c(e_I^C, \theta_I) - s_I\right],$$

and firm R's profits from poaching worker type θ_J are

$$\begin{aligned}\pi^R \;&=\; \alpha_J \left[f(e_J^*) - c(e_J^*, \theta_J) - s_J\right] \\ &=\; \alpha_J [s_J^* - s_J],\end{aligned}$$

because firm R chooses the efficient effort level by Lemma 5. Moreover, since firm R decides to poach worker θ_J it must be true that

$$\alpha_J [s_J^* - s_J] \;\geq\; V(s_I, s_J) \tag{1.7.1}$$

and

$$\alpha_J [s_J^* - s_J] \;\geq\; \alpha_I [s_I^* - s_I]. \tag{1.7.2}$$

We will now construct a profitable deviation for firm C with menu

$$\left((\hat{e}_I^C, \hat{w}_I^C), (\hat{e}_J^C, \hat{w}_J^C)\right) \;=\; \left((e_I^*, c(e_I^*, \theta_I) + s_I + \epsilon_I), (e_J^*, c(e_J^*, \theta_J) + s_J + \epsilon_J)\right),$$

where $\epsilon_I := \frac{\delta}{2}$ and $\epsilon_J := \frac{1}{2} \min\left(\epsilon_I, \frac{\alpha_I}{\alpha_J}\epsilon_I\right)$. Note that by construction $\epsilon_I, \epsilon_J > 0, \delta > \epsilon_I$ and

$$\epsilon_I \;>\; \max\left(\epsilon_J, \frac{\alpha_J}{\alpha_I}\epsilon_J\right). \tag{1.7.3}$$

We will proceed as follows: First, we will show that firm R will still poach worker type θ_J exclusively.[21] Then, we will establish that firm C makes strictly higher profits under this new menu.

Firm R prefers attracting worker type θ_J to getting both types, because

$$\begin{aligned}
\alpha_J[s_J^* - s_J - \epsilon_J] &\overset{1.7.1}{\geq} V(s_I, s_J) - \alpha_J \epsilon_J \\
&> V(s_I, s_J) - \epsilon_J \\
&= V(s_I + \epsilon_J, s_J + \epsilon_J) \\
&\overset{1.7.3}{\geq} V(s_I + \epsilon_I, s_J + \epsilon_J).
\end{aligned}$$

But firm R also favors poaching exclusively worker type θ_J over type θ_I:

$$\begin{aligned}
\alpha_J[s_J^* - s_J - \epsilon_J] &\overset{1.7.2}{\geq} \alpha_I[s_I^* - s_I] - \alpha_J \epsilon_J \\
&\overset{1.7.3}{>} \alpha_I[s_I^* - s_I] - \alpha_I \epsilon_I \\
&= \alpha_I[s_I^* - s_I - \epsilon_I].
\end{aligned}$$

Now consider firm C: By offering this menu she attracts worker type θ_I, because she did so under the old menu, and receives profits

$$\begin{aligned}
\hat{\pi}^C &= \alpha_I[s_I^* - s_I - \epsilon_I] \\
&> \alpha_I[s_I^* - s_I - \delta] \\
&= \alpha_I[f(e_I^C) - c(e_I^C, \theta_I) - s_I] \\
&= \pi^C.
\end{aligned}$$

Thus, firm C has a profitable deviation, contradiction. □

[21] For the remainder of the proof it is implicitly assumed that firm R chooses efficient effort levels and extracts all surplus up till s_I or s_J as a best response, when she poaches just one worker type. One could modify Lemma 5 to show that this response is optimal for her.

Preparation for Proof of Lemma 7

The value function $V(s_L, s_H)$ is given by

$$\begin{aligned} V(s_L, s_H) \ &:= \ \max_{e,w} \quad \alpha[f(e_L) - w_L] + (1-\alpha)[f(e_H) - w_H] \\ & s.t. \qquad (e,w) \in C(s_L, s_H) \end{aligned}$$

with

$$\begin{aligned} C(s_L, s_H) \ := \ \{ \ (e,w) \in (\mathbb{R}_+ \times \mathbb{R})^2 : & \\ w_L - c(e_L, \theta_L) \ &\geq \ s_L \\ w_H - c(e_H, \theta_H) \ &\geq \ s_H \\ w_L - c(e_L, \theta_L) \ &\geq \ w_H - c(e_H, \theta_L) \\ w_H - c(e_H, \theta_H) \ &\geq \ w_L - c(e_L, \theta_H) \ \} \, . \end{aligned}$$

For the proof we will make use of Berge's Theorem of the Maximum, which we state here for the sake of convenience.

Berge's Theorem of the Maximum. *Given sets $X \subset \mathbb{R}^n$ and $S \subset \mathbb{R}^p$, let $g : X \times S \longrightarrow \mathbb{R}$ be a continuous function, and $C : S \Longrightarrow X$ a compact-valued and continuous correspondence, and consider the parameterized maximization problem*

$$\max_{x \in C(s)} \quad g(x,s).$$

Then the value function

$$V(s) \ := \ \max_{x \in C(s)} \quad g(x,s)$$

is continuous.

In our case $X = (\mathbb{R}_+ \times \mathbb{R})^2$ is the set of possible contracts, $S = [0, s_L^*] \times [0, s_H^*]$, and $C(s)$ defines for each surplus level $s = (s_L, s_H) \in S$ a subset of admissible contracts determined by the participation and incentive constraints. The objective function g corresponds to the raiding firm's profit function if she poaches both employee types.

Showing $C(s)$ to be a continuous correspondence is the main difficulty of applying Berge's Theorem. One key idea will be to artificially bound the set of admissible contracts and will lead to a new constraint correspondence $\tilde{C}(s)$ and associated value function $\tilde{V}(s)$.

Conveniently, as we will see, $\tilde{V}(s) = V(s)$ for all $s \in S$, and it is sufficient for our purpose to establish the continuity of $\tilde{V}(s)$. But now we can restrict ourselves to proving the continuity of $\tilde{C}(s)$ which turns out to be feasible. We begin by defining lower and upper bounds for the set of admissible contracts.

Definition 3. Denote with $\bar{e}_L$ and $\bar{e}_H$ the upper bounds for the effort level for each worker type, implicitly defined by

$$\begin{aligned} \alpha[f(\bar{e}_L) - c(\bar{e}_L, \theta_L)] + (1-\alpha)s_H^* &= -s_H^*, \\ \alpha s_L^* + (1-\alpha)[f(\bar{e}_H) - c(\bar{e}_H, \theta_H)] &= -s_H^*. \end{aligned}$$

Each effort level is chosen in such a way that firm R makes a loss of s_H^*, even though she compensates one worker type just for her effort cost and extracts the maximum amount from the other type. Note that both bounds are well-defined, strictly positive, and unique. Furthermore, $\bar{e}_I > e_I^*$ for $I \in \{L, H\}$. This follows from the assumptions on the production and cost function.

Definition 4. Denote with $\bar{w}_L$ and $\bar{w}_H$ the upper bounds for the wage level for each worker type, implicitly defined by

$$\begin{aligned} \alpha[f(\bar{e}_L) - \bar{w}_L] + (1-\alpha)s_H^* &= -2s_H^*, \\ \alpha s_L^* + (1-\alpha)[f(\bar{e}_H) - \bar{w}_H] &= -2s_H^*. \end{aligned}$$

Similarly, at each wage level firm R obtains negative profits of $2s_H^*$, while having one worker exert the upper bound effort level and extracting the maximum surplus from the other type. Again, both bounds are well-defined, strictly positive and unique. Moreover, by construction, $s_H^* < \bar{w}_I$ for all $I \in \{L, H\}$.

Now we can define an associated maximization problem with additional bounds on the contract possibilites. Let

$$B \quad := \quad [0, \bar{e}_L] \times [0, \bar{w}_L] \times [0, \bar{e}_H] \times [0, \bar{w}_H]$$

and define

$$\begin{aligned}\tilde{V}(s_L, s_H) &:= \max_{e,w} \quad \alpha[f(e_L) - w_L] + (1-\alpha)[f(e_H) - w_H] \\ & \quad s.t. \qquad (e, w) \in \tilde{C}(s_L, s_H) := C(s_L, s_H) \cap B.\end{aligned}$$

Next, we will show that these extra bounds are not restrictive as long as $(s_L, s_H) \in S$:

Lemma 11. *For all $(s_L, s_H) \in S$ it holds $\tilde{V}(s) = V(s)$.*

PROOF: First, note that menu $(e_L, w_L, e_H, w_H) = (0, s_H^*, 0, s_H^*)$ lies in $C(s_L, s_H)$ for all $(s_L, s_H) \in S = [0, s_L^*] \times [0, s_H^*]$ because $s_L^* < s_H^*$. But these contracts also lie in B as $s_H^* < \bar{w}_I$ for all $I \in \{L, H\}$. Thus, $(0, s_H^*, 0, s_H^*) \in \tilde{C}(s_L, s_H)$ for all $(s_L, s_H) \in S$ and firm R can always guarantee herself profits of $-s_H^*$ for all surplus levels in S.

Obviously, we then have $V(s_L, s_H) \geq \tilde{V}(s_L, s_H) \geq -s_H^*$ for all $(s_L, s_H) \in S$. Now assume there exists $(\hat{s}_L, \hat{s}_H) \in S$ with $V(\hat{s}_L, \hat{s}_H) > \tilde{V}(\hat{s}_L, \hat{s}_H)$. Hence, we can find a contract $(\hat{e}, \hat{w}) \in C(\hat{s}_L, \hat{s}_H) \backslash \tilde{C}(\hat{s}_L, \hat{s}_H)$ that generates strictly higher profits. Assume $\hat{e}_I > \bar{e}_I$ for some $I \in \{L, H\}$. As both participation constraints have to be satisfied and $\hat{e}_I > \bar{e}_I > e_I^*$, firm R obtains profits $V(\hat{s}_L, \hat{s}_H) < -s_H^*$.[22] Thus, $\hat{e}_I \in [0, \bar{e}_I]$. So we must have $\hat{w}_I \notin [0, \bar{w}_I]$ for some $I \in \{L, H\}$. Due to the participation constraints it holds $\hat{w}_I \geq 0$ and it must be true that $\hat{w}_I > \bar{w}_I$. But this, together with individual rationality and $\hat{e}_I \leq \bar{e}_I$, implies $V(\hat{s}_L, \hat{s}_H) < -2s_H^*$, contradiction. Thus, for all $(s_L, s_H) \in S$ it holds $V(s_L, s_H) = \tilde{V}(s_L, s_H)$. □

We now have the tools at hand to prove the continuity of $V(s_L, s_H)$ on S.

PROOF OF LEMMA 7: By Lemma 11 it suffices to show continuity of $\tilde{V}(s_L, s_H)$ in (s_L, s_H) on S. We establish this by verifying that all conditions of Berge's Theorem are met. Evidently, firm R's profit function is continuous in (e, w).

Next, we prove that our constraint correspondence $\tilde{C}(s_L, s_H)$ is compact-valued. First, as was shown in the proof of Lemma 11, menu $(e_L, w_L, e_H, w_H) = (0, s_H^*, 0, s_H^*) \in \tilde{C}(s_L, s_H)$ for all $(s_L, s_H) \in S$. This implies that $\tilde{C}(s_L, s_H)$ is nonempty for all $(s_L, s_H) \in S$. Second, by the continuity of $c(e, \theta)$ and weak inequalities in the constraints, $C(s_L, s_H)$ is closed. As B is closed by definition, and intersections of closed sets are closed, so

[22] By definition, $f(e) - c(e, \theta_I)$ is strictly decreasing for $e > e_I^*$.

is $\tilde{C}(s_L, s_H)$. And third, because B is bounded, this holds true for $\tilde{C}(s_L, s_H)$. Thus, $\tilde{C}(s_L, s_H)$ is a compact-valued correspondence.

Finally, it remains to be shown that $\tilde{C}(s_L, s_H)$ is a continuous correspondence. We split up the proof by establishing upper- and lower-hemicontinuity separately and begin with the former.

Recall that a compact-valued correspondence $\tilde{C} : S \Longrightarrow B, (s_L, s_H) \longrightarrow \tilde{C}(s_L, s_H)$ is upper-hemicontinuous at (s_L, s_H) if and only if for every sequence $(s_L^n, s_H^n)_{n\in\mathbb{N}}$ converging to (s_L, s_H), every companion sequence $(e_L^n, w_L^n, e_H^n, w_H^n)_{n\in\mathbb{N}}$ with $(e_L^n, w_L^n, e_H^n, w_H^n) \in \tilde{C}(s_L^n, s_H^n)$ for all $n \in \mathbb{N}$, has a convergent subsequence with limit in $\tilde{C}(s_L, s_H)$. So take such a sequence and a corresponding companion sequence. Because $\tilde{C}(s_L^n, s_H^n) \subset B$ we can consider $(e_L^n, w_L^n, e_H^n, w_H^n)_{n\in\mathbb{N}}$ as a sequence in the bounded set B. By the Bolzano-Weierstrass Theorem there exists a converging subsequence with

$$\lim_{k\in\mathbb{N}} (e_L^{n_k}, w_L^{n_k}, e_H^{n_k}, w_H^{n_k}) = (e_L, w_L, e_H, w_H).$$

Obviously, $(e_L, w_L, e_H, w_H) \in B$ since B is closed. Furthermore, since $\lim_{k\in\mathbb{N}}(s_L^{n_k}, s_H^{n_k}) = (s_L, s_H)$, it is easy to check that $(e_L^{n_k}, w_L^{n_k}, e_H^{n_k}, w_H^{n_k}) \in C(s_L^{n_k}, s_H^{n_k})$ for all $k \in \mathbb{N}$ implies $(e_L, w_L, e_H, w_H) \in C(s_L, s_H)$.[23] Hence, $(e_L, w_L, e_H, w_H) \in \tilde{C}(s_L, s_H)$, which establishes upper-hemicontinuity in (s_L, s_H). Because $(s_L, s_H) \in S$ was arbitrary, $\tilde{C}(s_L, s_H)$ is upper-hemicontinuous on S.

Finally, note that a correspondence $\tilde{C} : S \Longrightarrow B, (s_L, s_H) \longrightarrow \tilde{C}(s_L, s_H)$ is lower-hemicontinuous at (s_L, s_H) if and only if for every sequence $(s_L^n, s_H^n)_{n\in\mathbb{N}}$ converging to (s_L, s_H) and every point $(e_L, w_L, e_H, w_H) \in \tilde{C}(s_L, s_H)$ there exists a companion sequence $(e_L^n, w_L^n, e_H^n, w_H^n)_{n\in\mathbb{N}}$ with $(e_L^n, w_L^n, e_H^n, w_H^n) \in \tilde{C}(s_L^n, s_H^n)$ for all $n \in \mathbb{N}$, that converges to (e_L, w_L, e_H, w_H). Again, take such a sequence and such a point and let's construct the desired companion sequence. We consider three cases separately.

1. $w_L < \bar{w}_L$ and $w_H < \bar{w}_H$: Define for each $n \in \mathbb{N}$

$$\delta^n := \max(|s_L^n - s_L|, |s_H^n - s_H|).$$

[23]This follows immediately from the fact that the constraint functions are continuous in $(e_L, w_L, e_H, w_H, s_L, s_H)$, and the weak inequalities of the constraints.

By assumption $\lim_{n\in\mathbb{N}} \delta^n = 0$. Thus, we can find $N \in \mathbb{N}$ such that $\delta^n < \min(\bar{w}_L - w_L, \bar{w}_H - w_H)$ for all $n \geq N$. Now define

$$(e_L^n, w_L^n, e_H^n, w_H^n) \quad \begin{cases} \in \tilde{C}(s_L^n, s_H^n) & \text{if } n < N \\ := (e_L, w_L + \delta^n, e_H, w_H + \delta^n) & \text{if } n \geq N. \end{cases}$$

Note that this assignment is well-defined, because we know $\tilde{C}(s)$ is nonempty for all $s \in S$. This sequence obviously converges to (e_L, w_L, e_H, w_H). By construction $(e_L^n, w_L^n, e_H^n, w_H^n) \in B$ for all $n \in \mathbb{N}$. Finally, observe that $(e_L^n, w_L^n, e_H^n, w_H^n) \in C(s_L^n, s_H^n)$ for all $n \in \mathbb{N}$. This is evident for $n < N$. For $n \geq N$, because $(e_L, w_L, e_H, w_H) \in C(s_L, s_H)$, the incentive constraints are satisfied, whereas the individual rationality constraints are potentially relaxed. Thus, we found a companion sequence with $(e_L^n, w_L^n, e_H^n, w_H^n) \in \tilde{C}(s_L^n, s_H^n)$ that converges to (e_L, w_L, e_H, w_H).

2. $w_I = \bar{w}_I$ for both $I \in \{L, H\}$: We will show that the constant sequence $(e_L^n, w_L^n, e_H^n, w_H^n) = (e_L, w_L, e_H, w_H)$ is an appropriate companion sequence. Obviously, she converges to (e_L, w_L, e_H, w_H). Because $(e_L, w_L, e_H, w_H) \in \tilde{C}(s_L, s_H)$, the sequence lies in B, but it also implies that the incentive constraints are always satisfied for all (s_L^n, s_H^n). Finally, note that both types of workers receive a surplus above s_H^* from menu (e_L, w_L, e_H, w_H):

$$\begin{aligned} w_I - c(e_I, \theta_I) &= \frac{2}{\alpha_I} s_H^* + \frac{\alpha_J}{\alpha_I} s_J^* + f(\bar{e}_I) - c(e_I, \theta_I) \\ &\geq \frac{2}{\alpha_I} s_H^* + \frac{\alpha_J}{\alpha_I} s_J^* + f(\bar{e}_I) - c(\bar{e}_I, \theta_I) \\ &= \frac{2}{\alpha_I} s_H^* + \frac{\alpha_J}{\alpha_I} s_J^* - \frac{1}{\alpha_I} s_H^* - \frac{\alpha_J}{\alpha_I} s_J^* \\ &> s_H^*. \end{aligned}$$

Since $(s_L, s_H) \in S$ the individual rationality constraints will always be satisfied and $(e_L, w_L, e_H, w_H) \in C(s_L^n, s_H^n)$ for for all $n \in \mathbb{N}$. Thus, for our constant sequence $(e_L, w_L, e_H, w_H) \in \tilde{C}(s_L^n, s_H^n)$ for all $n \in \mathbb{N}$ as desired.

3. $w_I = \bar{w}_I$ for precisely one $I \in \{L, H\}$: First, note that for $e_I = 0$ we're already done, because it follows from the incentive constraint of type θ_J that $w_J - c(e_J, \theta_J) \geq \bar{w}_I - c(0, \theta_J) > s_H^*$. Both types obtain a surplus above s_H^* and we could proceed as in 2. So let $e_I > 0$. Define $\delta^n := |s_J^n - s_J|$, and by construction $\lim_{n\to\infty} \delta^n = 0$. So we can find

$N_w \in \mathbb{N}$ with $w_J + \delta^n < \bar{w}_J$ for all $n \geq N_w$. We now consider three subcases.
(3.a) Type θ_I strictly prefers contract (e_I, w_I) to contract (e_J, w_J): In that case we can find $N_{\delta_I} \in \mathbb{N}$ such that $w_J + \delta^n - c(e_J, \theta_I) < w_I - c(e_I, \theta_I)$ for all $n \geq N_{\delta_I}$. Define $N_a := \max(N_w, N_{\delta_I})$ and

$$(e_I^n, w_I^n, e_J^n, w_J^n) \quad \begin{cases} \in \tilde{C}(s_I^n, s_J^n) & \text{if } n < N_a \\ := (e_I, w_I, e_J, w_J + \delta^n) & \text{if } n \geq N_a. \end{cases}$$

This assignment is well-defined and the sequence obviously converges to (e_I, w_I, e_J, w_J). By construction $(e_I^n, w_I^n, e_J^n, w_J^n) \in B$ for all $n \in \mathbb{N}$. Finally, observe that $(e_I^n, w_I^n, e_J^n, w_J^n) \in C(s_L^n, s_H^n)$ for all $n \in \mathbb{N}$. Thus, we found a companion sequence with $(e_I^n, w_I^n, e_J^n, w_J^n) \in \tilde{C}(s_I^n, s_J^n)$ that converges to (e_I, w_I, e_J, w_J).
(3.b) Type θ_I is indifferent between contract (e_I, w_I) and contract (e_J, w_J), but type θ_J strictly prefers (e_J, w_J): We construct a sequence $(\epsilon^n)_{n \in N}$ that is implicitly defined by $w_J + \delta^n - c(e_J, \theta_I) = w_I - c(e_I - \epsilon^n, \theta_I)$ whenever possible ($e_I - \epsilon^n \geq 0$ restriction), and $\epsilon^n = -1$ else. Because $e_I > 0$ and $\lim_{n \to \infty} \delta^n = 0$, we can find $N_\epsilon \in \mathbb{N}$ with $\epsilon^n \geq 0$ for all $n \geq N_\epsilon$ and $\lim_{n \to \infty} \epsilon^n = 0$. Because type θ_J strictly prefers contract (e_J, w_J), we can find $N_{\delta_J} \in \mathbb{N}$ such that $w_J + \delta^n - c(e_J, \theta_J) \geq w_I - c(e_I - \epsilon^n, \theta_J)$ for all $n \geq N_{\delta_J}$. Define $N_b := \max(N_w, N_\epsilon, N_{\delta_J})$ and

$$(e_I^n, w_I^n, e_J^n, w_J^n) \quad \begin{cases} \in \tilde{C}(s_I^n, s_J^n) & \text{if } n < N_b \\ := (e_I - \epsilon^n, w_I, e_J, w_J + \delta^n) & \text{if } n \geq N_b. \end{cases}$$

As before one can easily verify that $(e_I^n, w_I^n, e_J^n, w_J^n) \in \tilde{C}(s_I^n, s_J^n)$ for all $n \in \mathbb{N}$ and converges to (e_I, w_I, e_J, w_J).
(3.c) Both types are indifferent between contract (e_J, w_J) and (e_I, w_I). In this case both contracts are identical, which easily follows by looking at both incentive constraints and using the single-crossing property. Let's write (e, w) for this contract. Construct a sequence $(\gamma^n)_{n \in \mathbb{N}}$ that is implicitly defined by $w - c(e - \gamma^n, \theta_J) = s_J + \delta^n$ whenever possible

($e-\gamma^n \geq 0$ restriction), and $\gamma^n = -1$ else. Because $e = e_I > 0$ and $\lim_{n\to\infty} \delta^n = 0$, we can find $N_c \in \mathbb{N}$ with $\gamma^n \geq 0$ for all $n \geq N_c$ and $\lim_{n\to\infty} \gamma^n = 0$. Define

$$(e_I^n, w_I^n, e_J^n, w_J^n) \qquad \begin{cases} \in \tilde{C}(s_I^n, s_J^n) & \text{if } n < N_c \\ := (e-\gamma^n, w, e-\gamma^n, w) & \text{if } n \geq N_c. \end{cases}$$

Again, one can verify that $(e_I^n, w_I^n, e_J^n, w_J^n) \in \tilde{C}(s_I^n, s_J^n)$ for all $n \in \mathbb{N}$ and converges to (e_I, w_I, e_J, w_J).

Hence, we can find the desired sequence when $w_I = \bar{w}_I$ for precisely one $I \in \{L, H\}$.

By taking all three cases together, we have proved that our correspondence is also lower-hemicontinuous in (s_L, s_H). Because $(s_L, s_H) \in S$ was arbitrary it is also lower-hemicontinuous on S. Together with upper-hemicontinuity it follows that the correspondence ist continuous on S. Thus, all conditions of the Theorem of the Maximum are satisfied and the proof is complete. □

Proof of Lemma 10: Following the proof of Lemma 1, firm R's profit from poaching both worker types if firm C offers menu

$$\left((e_L^C, w_L^C), (e_H^C, w_H^C)\right) \;=\; \left((e_L^*, c(e_L^*, \theta_L) + s), (e_H^*, c(e_H^*, \theta_H) + s)\right),$$

with $s \geq 0$ is

$$\begin{aligned}
\pi^R(\beta) = {} & \alpha\left[f(\hat{e}_L(\beta))/\beta - c(\hat{e}_L(\beta), \theta_L) - s\right] + (1-\alpha)\left[f(\hat{e}_H(\beta))/\beta\right. \\
& \left. - \{c(\hat{e}_H(\beta), \theta_H) + c(\hat{e}_L(\beta), \theta_L) - c(\hat{e}_L(\beta), \theta_H)\} - s\right] \\
= {} & \alpha\left[f(\hat{e}_L(\beta))/\beta - c(\hat{e}_L(\beta), \theta_L) - s\right] - (1-\alpha)[c(\hat{e}_L(\beta), \theta_L) - c(\hat{e}_L(\beta), \theta_H)] \\
& + (1-\alpha)\left[f(e_H^*(\beta))/\beta - c(e_H^*(\beta), \theta_H) - s\right].
\end{aligned}$$

Note that the third term in the last equation is the profit firm R can guarantee herself if she decided to just focus on the high worker type. Moreover, as $s \nearrow f(\hat{e}_L(\beta))/\beta - c(\hat{e}_L(\beta), \theta_L)$ the first two expressions become in total strictly negative.

Thus, define

$$\begin{aligned}
\underline{s}(\beta) \;&:=\; f(\hat{e}_L(\beta))/\beta - c(\hat{e}_L(\beta), \theta_L) - \frac{1-\alpha}{\alpha}[c(\hat{e}_L(\beta), \theta_L) - c(\hat{e}_L(\beta), \theta_H)] \\
\bar{s} \;&:=\; f(e_L^*) - c(e_L^*, \theta_L).
\end{aligned}$$

Note that $\underline{s}(\beta) < \bar{s}$, because $\hat{e}_L(\beta) < e_L^*(\beta) \leq e_L^*$ for all $\beta \geq 1$. Then for any firm C menu as defined above with $s \in (\max(\underline{s}(\beta), 0), \bar{s})$ firm C makes strictly positive profits with the low type as in Lemma 1.

The proof is complete when $\underline{s}(\beta) \leq \underline{s}(1)$ for all $\beta \geq 1$, because firm C could then always offer the above menu with $\tilde{s} := [\max(0, \underline{s}(1)) + \bar{s}]/2$ and obtain profits $\underline{\pi}^C := \alpha[f(e_L^*) - c(e_L^*, \theta_L) - \tilde{s}] \in (0, \bar{\pi}^C)$.

We will prove this by showing that $\underline{s}(\beta)$ is a decreasing function of β. Taking the derivative of $\underline{s}(\beta)$ with respect to β gives

$$\begin{aligned}
\frac{d\underline{s}(\beta)}{d\beta} &= \left[f'(\hat{e}_L(\beta))/\beta - c'(\hat{e}_L(\beta), \theta_L) - \frac{1-\alpha}{\alpha}[c'(\hat{e}_L(\beta), \theta_L) - c'(\hat{e}_L(\beta), \theta_H)]\right] \frac{d\hat{e}_L(\beta)}{d\beta} \\
&\quad - f(\hat{e}_L(\beta))/\beta^2 \\
&= -f(\hat{e}_L(\beta))/\beta^2 \\
&\leq 0.
\end{aligned}$$

The term in square brackets vanishes, because it corresponds to the derivative of firm R's profit function with respect to $e_L(\beta)$ evaluated at the optimum. □

Chapter 2

An Informed Principal Problem with Moral Hazard and Irrelevance of the Principal's Information

We consider a discrete principal-agent model with moral hazard in which the principal has private information regarding the stochastic relationship between effort and profits. With both parties risk neutral, we provide sufficient conditions for this relationship to render the principal's information inconsequential and which are novel in informed principal problems with moral hazard. A full characterization is derived for a special case of the model. This establishes an unstudied connection between distortive effects of the principal's information and the number of available profit states. Finally, if the agent is risk averse, more subtle issues arise and equilibrium payoffs can merely be bounded by the second-best profit from the standard model of moral hazard.

2.1 Introduction

In a standard principal-agent model there is a set of observable profit levels that are imperfectly correlated with the agent's hidden effort choice. The principal attempts to utilize these observables to design an incentive scheme maximizing his expected profits. A classic example is a firm owner (the principal) who wants to hire a manager (the agent) to run his firm. We will refer to this framework as the *standard model*.

If both parties are risk neutral, a particularly salient solution to the principal's problem exists. By selling the firm to the agent, the agent becomes the residual claimant of firm profits. Because she has no concerns regarding risk, the efficient action is chosen, which maximizes the difference between expected firm profits and her effort cost. The principal anticipates this behavior and asks for a firm price that makes the agent just indifferent between accepting and rejecting the offer. He can therefore extract the same payoff as in the situation when effort is observable. We will refer to the selling-the-firm contract as the *standard solution* of the principal's problem.

A natural extension of the standard model is the case when the stochastic relationship between the agent's effort and firm profits is private information of the principal. For instance, the firm owner has privileged access to information that is designed for in-house use only (e.g. quality of staff), and information that is to become public, but has not yet been published (e.g. quarterly reports). It is easy to see that in general selling-the-firm contracts do not provide a solution to this problem: A principal type with a more profitable technology can ask for a higher firm price than a type with a less profitable one. But then it is always in the latter's interest to pretend to be a principal with a more profitable technology in order to obtain the higher firm price. Thus, after observing the high type's selling-the-firm contract, the agent rejects the contract because the profits she expects are strictly below the offered price of the firm, i.e. the standard solution fails.

There are two major contributions of our paper: First, we show that the principal can still obtain his payoff from the standard model despite this additional layer of private information. More precisely, we identify a linear independence condition for the stochastic relationship between effort and profits under which the principal's payoff is unaffected by his private information. To our knowledge this is the first approach to study informational irrelevance in an informed principal setup with moral hazard.

Second, we show for a special case of the model that these distributional assumptions are also necessary for the irrelevance of the principal's information. Because these conditions can now be directly related to the number of available profit states in the model, distortive effects of the principal's information hinge on this number. This role of contractible signal states in a principal-agent setup is novel. Moreover, it allows comparisons to a branch of the literature that traces specific distortions back to the principal's information and qualifies these results for our setup.

In Section 2.2, we review related literature before we present the model in Section 2.3. In our principal-agent framework profit levels and actions available to the agent are both finite and each action choice induces a specific probability distributions over profits. In the standard model of moral hazard, the principal and the agent know these distribution for a given effort level of the agent. In our setup this knowledge is private information of the principal, and the relationship translating effort into stochastic profits is altered by different types of the principal. Therefore, a distribution over profits for each action-type combination exists. In what follows the set of these distributions will play key role in our analysis. Note also that the principal's wage offer to the agent does not only provide incentives to choose a specific effort level, but it also signals her which type of principal she is facing.

Sections 2.4-2.6 contain all results of the paper. In our analysis, we focus on the relationship between the set of distributions over profits and the distortive power of the principal's information. We start assuming that the agent is risk neutral. Theorem 1 shows that linear independence of the set of distributions over profits is a sufficient condition for the irrelevance of the principal's information. More precisely, in any Perfect Bayesian Equilibrium the principal obtains his profits from the standard moral-hazard model. Intuitively, a given type of principal can only fall short of his profits from the standard model if the agent responds to his optimal contracts with a different action choice, which in turn requires the agent to believe that she faces a different principal. This event can be ruled out, however, if there exists a contract that induces the agent to choose the *correct* action, no matter what the agent believes, and has this principal type pay just the expected wage of the standard model. Such a contract, nevertheless, must give incentives to provide effort for any given principal type, which is feasible when the set of distribution over profits is linearly independent. A more intricate argument is necessary to rule out that any principal type obtains higher profits than in the standard case.

We do not prove that linear independence is also a necessary condition for the irrelevance of the principal's information, because lots of equilibria can be sustained in a signaling game due to the free choice of off-equilibrium path beliefs. But we provide this result in a special case of the model of independent interest (Theorem 2) in which we maintain the agent's risk neutrality, but allow for only three distributional vectors over profits. In that case, equilibrium conditions already impose so much structure on the payoffs of the equilibrium contracts that these can only be reached under linearly independent distribution vectors. In particular, we employ the fact that a triple of linearly dependent distribution vectors allows a representation of one vector as a convex combination of the other two (Lemma 1) which follows from their normalization to unity. We then turn to the problem of equilibrium existence under linear independence. Because in equilibrium each type obtains his profits from the standard model, it is convenient—especially for the construction of pooling equilibria—to identify the set of contracts that can be offered by each type in the standard model and provide him with these payoffs. Lemma 2 proves that this set is nonempty; again, this follows from the wide range of outcome vectors that can be reached under linear independence of the distributions. Contracts from this set serve as candidates for equilibrium contracts. The main problem of an equilibrium construction is to find an adequate assignment of beliefs for off-equilibrium path contracts such that potential deviations for the principal types are removed. A principal can only improve on his payoff from the standard model if the agent—errorenoulsy believing to face a different type of principal—responds to a given contract differently than in the standard case. This problem can be resolved by assuming that the agent always believes to face the principal type that benefits from this contract. Theorem 3 shows that such an assignment can be made and that an equilibrium exists. We then scrutinize the linear independence condition in relation to the number of profit states. While it is obviously never fulfilled with two profit levels, we argue that it is generically satisfied if there are three or more states (Proposition 1). Because private information of the principal causes distortions if and only if the distributions vectors are independent, the irrelevance of the principal's information can be directly linked to the number of attainable profit levels. To this end we demonstrate how costly signaling occurs in our model if there are just two profit states (Proposition 2) and discuss this phenomenon employed in some works of the literature.

Finally, we transfer our model to the case of a strictly risk averse agent. Now, linear independence of the set of distribution vectors over profits is no longer sufficient to effect the profits of the standard model, because complexity of contracts comes with a price. Contracts that implement a specific action—no matter what type of principal the agent expects—still exist, but may require an additional risk premium due to the greater uncertainty the agent faces. We provide an example to illustrate this point. More generally, we show that no principal type can obtain more than his standard profits—irrespective of the fact whether the set of distribution vectors over profits is linearly independent or not (Theorem 4).

Section 2.7 concludes; omitted proofs are relegated to Section 2.8.

2.2 Related Literature

We begin by reviewing some works in principal-agent theory that focus on irrelevance of private information. Thereafter, we discuss related papers analyzing informed principal models with moral hazard.

Irrelevance of information in a principal-agent setting is studied in Riordan and Sappington (1988). A principal wants to buy a quantity of a good from an agent, but does not know her cost function. There exists, however, an imperfect signal of the agent's cost, which can be observed after contracting and is verifiable so that the principal's contract can be made contingent on it. Both the signal and the type space are finite. When the conditional distributions of the agent's cost types given a signal are linearly independent, the principal can select a quantity-price menu that leaves the agent's information inconsequential. More precisely, the principal attains the same profit as if the agent's cost type was known to him. He can do so, because linear independence allows him to design menus that discriminate strongly across types and keep the agent's information rents at zero. Irrelevance of information under a linear independence assumption in an auction setting is analyzed in Crémer and McLean (1988). In our setup, we also find this distributional condition to render the principal's information immaterial (Theorem 1). But there are three major differences to our work: First, regarding the agent's information, we deal with a moral-hazard and not an adverse-selection problem. Second, we also introduce private information on behalf of the principal. Importantly, contract offers of the principal now transmit information about his type, which allows for signaling effects in our model. And third, in our model the principal's information becomes irrelevant, whereas in their models it is the agent's.

Maskin and Tirole (1990) consider optimal mechanisms in an informed principal setting. They show that generically the principal is strictly better off than when his information is known to the agent. Intuitively, under private information the principal has to satisfy his incentive and individual rationality constraint only in expectation instead of fulfilling them for each type individually. For instance, a violation of the incentive constraint for one type can be offset by a slack incentive constraint by another type and vice versa for the individual rationality constraint. This additional leeway allows the principal to improve on her full-information payoff. Maskin and Tirole show, however, that for quasi-linear objective functions of the players no gains from such constraint trade can be made. It follows that in this special case the principal obtains the same payoff as in the full-information situation. More recently, Mylovanov and Tröger (forthcoming) further advance this result. They identify in an informed principal model a general condition under which the principal's optimal mechanism coincides with the one if his information is known. In particular, this condition is satisfied in the special case of Maskin and Tirole. In comparison to our work, however, both papers share three noteworthy differences: First, the irrelevance conditions presented in their articles rest on the players' preferences; our focus is on a distributional assumption. Second, in both models private value environments are considered, i.e. the princpal's type does not directly enter the agent's utilitiy. Our framework, however, is not of private values, because the principal's type influences the probabilites of the agent's payoff. Thus, he affects her preferences directly. And third, in contrast to our setup, moral hazard is explicitly ruled out in their models.

More distantly, an early and very general discussion of the informed principal problem is Myerson (1983). He studies the principal's mechanism selection problem under virtually no restriction on the players' preferences and interactions, except a finite type and action space. Several solution concepts for the principal's problem are proposed and related with each other. While Myerson's framework includes moral-hazard problems and notably covers our model, his main focus lies on optimal mechanism design and he does not analyze under which circumstances the principal's information is relevant or not.

Inderst (2001) also does not study informational irrelevance, but has a principal-agent structure very similar to ours. The principal has private information regarding the stochastic relationship between effort and profits, but cannot observe the agent's effort choice. Moreover, both the agent and the principal are risk neutral. Yet, there are some important differences: the agent's task level can be adjusted continuously, only two outcome levels (success and failure) can occur, and the relationship between effort and the probability of success is linear. The high type of the principal is identified by a steeper probability function than the low type. As in our model, selling-the-firm contracts are not incentive compatible anymore; the low type, who can demand from the agent only a relatively low price for the firm, wants to mimic the high type and ask for a higher price. In order to separate from the low type principal, therefore, the high type increases the wage in the case of failure and decreases it in the case of success compared to his selling-the-firm contract. Since the low type has a higher probability for failure, this rearrangement of the wage schedule is more costly to him than to the high type.[1] Inderst shows that there exists a unique Perfect Bayesian Equilibrium satisfying the Intuitive Criterion in which the high type separates from the low type by providing a higher base wage and a lower bonus than in the standard moral hazard model. He finds these results in line with the observation that real world incentive schemes are quite low-powered compared to theoretical predictions of agency models. Beaudry (1994) uses an almost identical model structure to establish that the agent obtains a positive rent from the principal's signaling activity. In contrast, for finite task levels, our Theorems 1 and 2 show that distortive effects can only occur if the set of distribution vectors over profits is linearly dependent and is not a general feature of informed principal models with moral hazard.

Another branch of research focuses on the informed principal problem with moral hazard under the assumption of a strictly risk averse agent. As in the literature for the risk neutral case, there is no direct work on the irrelevance of the principal's information. Chade and Silvers (2002) study equilibrium outcomes in a two profit, two effort level setup that cannot arise in the standard model. In particular, there exist equilibria in which the type with the more informative technology obtains strictly less profits than the type with the less informative one, and equilibria that allow the agent to accrue a strictly positive rent. Silvers (2012) generalizes the analysis to a finite action and profit space. One finding in this paper is that private information of the principal is never beneficial to him and conditions are provided under which the principal is even strictly worse off. This result is consistent with Theorem 4 of our model, but does not generalize it, even though we impose a more restrictive action space for the risk averse case; Silvers assumes the distribution

[1] This wage structure also induces the agent to work less hard and further depresses the high type's profits. Note, however, that manipulating the wage schedule in the other direction is more harmful to the high than the low type and does not allow for a separation.

functions in his model to satisfy a monotone likelihood ratio property we can dispense with. In fact, we have no constraints at all on the distributions vectors in the risk averse case.

2.3 Model

We begin by introducing the timing of the game before we turn to the players' preferences.

First, nature draws the principal's type $t \in \mathcal{T} = \{t_1, \dots, t_T\}$ with probability λ_t, where $\lambda_t \in (0,1)$ for all $t \in \mathcal{T}$. After observing his type, the principal offers the agent a take-it-or-leave-it contract. The agent can turn down the offer in which case the principal obtains a payoff normalized to zero, while the agent gains the value of her outside option $\bar{U} \in \mathbb{R}$. If the agent accepts the contract, she selects an action $a \in \mathcal{A} = \{a_1, \dots, a_A\}$, which is unobservable to the principal. After that a profit level from the vector $\pi = (\pi_1, \dots, \pi_n) \in \mathbb{R}^n$ is realized ($n \geq 2$) according to some probability distribution $f(t,a) = (f_1(t,a), \dots, f_n(t,a)) \in \mathbb{R}^n$, where $f_i(t,a) > 0$ for all $i \in \{1, \dots, n\}, a \in \mathcal{A}$, and $t \in \mathcal{T}$.

The prinicipal designs a wage scheme $w = (w_1, \dots, w_n) \in \mathbb{R}^n$ to maximize his expected profits[2]

$$V(w,a) \;=\; \sum_{i=1}^{n} f_i(t,a)\,(\pi_i - w_i) \;=\; f(t,a) \cdot (\pi - w).$$

The agent's preferences over wage $\tilde{w} \in \mathbb{R}$ and action a can be represented by an additively separable von Neumann-Morgenstern utility function $\tilde{u}(\tilde{w}) - c(a)$, where $\tilde{u} : (\underline{\tilde{w}}, \infty) \to \mathbb{R}$ (with $\underline{\tilde{w}} \in \mathbb{R} \cup \{-\infty\}$) is a continuous and strictly increasing function and $c : \mathcal{A} \to \mathbb{R}$. Moreover, $\lim_{\tilde{w} \searrow \underline{\tilde{w}}} \tilde{u}(\tilde{w}) = -\infty$. Thus, overall preferences over incentive scheme w and action a can be summarized by

$$U(w,a) \;=\; \sum_{i=1}^{n} f_i(t,a)\tilde{u}(w_i) - c(a) \;=\; f(t,a) \cdot u(w) - c(a),$$

where $u(w) := (\tilde{u}(w_1), \dots, \tilde{u}(w_n)) \in \mathbb{R}^n$. Regarding attitudes towards risk, the agent can either be risk neutral ($\tilde{u}(w) = w$) or strictly risk averse ($\tilde{u}(w)$ strictly concave).

[2]Throughout the paper the dot symbol denotes the inner product of two vectors in the Euclidean space. The scalar product and the product of a scalar with a vector are both depicted without any algebraic symbol, but the meaning of a certain operation should be clear from the context.

Furthermore, we assume that for any $a \in \mathcal{A}$ there exists $\tilde{w}(a) \in \mathbb{R}$ such that $\tilde{u}(\tilde{w}(a)) = \bar{U} + c(a)$.[3] For each contract proposal w, the agent holds a belief $\mu(t|w)$ to face type $t \in \mathcal{T}$.

We look for pure strategy Perfect Bayesian Equilibria and make use of the Intuitive Criterion proposed by Cho and Kreps (1987) at times in order to give a sharper prediction for the outcome of the game.

2.4 Analysis: General Case, Risk Neutral Agent

We begin our analysis under the assumption of a risk neutral agent. In this section, we provide sufficient conditions for the stochastic relationship between effort and profits that cause the principal's private information to be irrelevant. More precisely, in any Perfect Bayesian Equilibrium the principal's paypoff is identical to the one in the standard principal-agent model of moral hazard. For that purpose, we briefly review some basic facts of this benchmark case before we return to our extended model.

Let's fix the principal's type at $t \in \mathcal{T}$ for a moment and allow for a concave utility function of the agent before we turn to the special case of risk neutrality. Following the approach of Grossman and Hart (1983), the principal can decompose his profit maximization problem in two subproblems. In the first step, he infers for each action $a \in \mathcal{A}$ the minimum implementation cost by solving

$$\min_{w \in \mathbb{R}^n} \quad f(t,a) \cdot w$$

$$\text{s. t.} \qquad f(t,a) \cdot u(w) \geq \bar{U} + c(a) \tag{IR}$$

$$f(t,a) \cdot u(w) - c(a) \geq f(t,\hat{a}) \cdot u(w) - c(\hat{a}), \qquad \forall \hat{a} \in \mathcal{A}. \tag{IC}$$

The individual rationality constraint (IR) induces the agent to accept the contract, whereas the incentive constraints (IC) lets action a be an optimal choice for her. Then, in the second step, the principal maximizes the difference between expected profits and implementation costs over actions. This procedure leads to a solution (their Proposition 1). Let the pair $\left(a^{SB}(t), w^{SB}(t)\right)$ be such a second-best solution, where incentive scheme $w^{SB}(t)$ implements action $a^{SB}(t)$, and denote with $a^{FB}(t)$ the first-best action implemented in

[3] Under this additional assumption, for a fixed type $t \in \mathcal{T}$, our setup satisfies all assumptions of the principal-agent model of Grossman and Hart (1983). Thus, our reference to some of their standard results further below is warranted.

the observable action case. Clearly, the principal's first-best profits always weakly exceed the second-best ones.

It is, however, a well-known result that risk neutrality of the agent allows the principal to attain first-best profits (their Proposition 3(2)). In fact, there exists an especially prominent incentive scheme for this purpose: the so-called selling-the-firm contract induces the agent to select $a^{FB}(t)$ and is defined by $w_i^{stf}(t) := \pi_i - \left[B(t) - \bar{U} - c(a^{FB}(t))\right]$. The agent becomes the residual claimant of the gross profits, reduced by the value of the firm $B(t) := f(t, a^{FB}(t)) \cdot \pi$ up to a compensation for the agent's outside option and her first-best effort cost. Since the agent has no concerns regarding risk, his optimal action choice is characterized by finding the task level that maximizes the difference between expected firm profits and effort costs. The solution to this problem is the efficient action, which coincides with $a^{FB}(t)$.[4] Thus, in the second-best outcome $a^{SB}(t) = a^{FB}(t)$, the principal obtains the first-best payoff $B(t) - \left[\bar{U} + c(a^{FB}(t))\right]$, and the agent accrues a net utility of $\bar{U}$.[5] Note, in particular, that the principal's profit is constant across each state of nature.

In our general model, nature selects the principal's type $t \in \mathcal{T}$ at the beginning of the game. But now it is not incentive compatible that each type offers his selling-the-firm contract. Unless the first-best profit $B(t) - \left[\bar{U} + c(a^{FB}(t))\right]$ is the same for all $t \in \mathcal{T}$, less profitable types will try to mimic more profitable ones, because these contracts provide a constant payment to the principal. Thus, there's no obvious contract that implements the outcome of the benchmark model; moreover, payoffs may vary across different equilibria.

We now establish sufficient conditions for the prinicpal's payoff to be first-best in any Perfect Bayesian Equilibrium. For an intuition, first note that a principal type cannot exceed first-best profits in a separating equilibrium, because his type is revealed. Hence, he must pool his wage offer with other types. But as the agent's individual rationality constraint must hold in expectation, there must exist a type with profits strictly less than first-best. Thus, in order to show our result, it suffices to restrict attention to the case of types with below first-best payoffs. Such a type may not deviate to a better contract, say

[4]Note that under observable effort $a^{FB}(t)$ is identical with the efficient action $a^*(t)$, which maximizes the difference between expected firm profits $f(t, a) \cdot \pi$ and effort cost $c(a)$. Under observable effort, the principal also selects this action, because there are no additional implementation costs involved: He could simply write a contract that compensates the agent with $\bar{U} + c(a^*(t))$ if $a^*(t)$ is observed and else imposes a penalty close to $\underline{\tilde{w}}$. A similar argument extends to the case of a risk-averse agent.

[5]The first- or second-best action need not be unique. In this case, the set of optimal actions in the first-best and second-best case coincide.

his selling-the-firm contract, because the agent holds a belief for this offer that induces her to select an action different from his first-best action or reject the contract altogether. But if there exists a contract that makes the agent choose the type's first-best action no matter what kind of principal she expects and gives the type his first-best profit, a profitable deviation exists for the principal. Thus, in equilibrium this type cannot obtain a payoff below first-best.

How can we guarantee that such a contract exists for every type? Roughly speaking, for a fixed type, the basic requirement for the contract is to generate a high utility for the agent for some distributions (when the first-best action is chosen) and a rather low one for others (when unfavorable actions are chosen). Because the first-best action to be implemented can be different from type to type, in order to find such contracts, a rich set of payoff vectors for the agent has to be reached with the set of effort-profit distributions. If this set is linearly dependent, only a subspace of all possible payoff vectors can be attained and one cannot guarantee the existence of such a contract. If, however, this set is linearly independent, for each payoff vector a corresponding contract exists. In particular, we can find one with the desired properties from above.

Theorem 1. *Assume the agent is risk neutral. If the set of vectors* $\{f(t,a)\}_{t\in\mathcal{T},a\in\mathcal{A}}$ *is linearly independent, in any Perfect Bayesian Equilibrium each type of principal obtains his first-best profits.*

PROOF: Assume there exists a Perfect Bayesian Equilibrium in which type $t_i \in \mathcal{T}$ does not obain his first-best profits. We will show that at least one type of principal has a profitable deviation.

First, assume that type $t_i \in \mathcal{T}$ obtains strictly lower profits and let $\epsilon > 0$ be the difference between first-best and equilibrium profits. For notational convenience, renumber the actions of the agent such that $a_i = a^{FB}(t_i)$. Consider contract $w \in \mathbb{R}^n$ satisfying $f(a,t) \cdot w = \bar{U} + c(a) + x$ for all $a \in \mathcal{A},\ t \in \mathcal{T}$ with $x = \frac{\epsilon}{2}$ if $a = a_i$ and $x = 0$ else, that is

$$\underbrace{\begin{pmatrix} f(t_1,a_1) \\ \vdots \\ f(t_1,a_i) \\ \vdots \\ f(t_1,a_A) \\ f(t_2,a_1) \\ \vdots \\ f(t_T,a_A) \end{pmatrix}}_{=:\,D} \cdot w \quad = \quad \underbrace{\begin{pmatrix} \bar{U} + c(a_1) \\ \vdots \\ \bar{U} + c(a_i) + \frac{\epsilon}{2} \\ \vdots \\ \bar{U} + c(a_A) \\ \bar{U} + c(a_1) \\ \vdots \\ \bar{U} + c(a_A) \end{pmatrix}}_{=:\,b}$$

A standard result in Linear Algebra shows that a solution to this system exists if and only if the vector b lies in the column space of matrix D (Fischer (2005, p. 130)). By assumption the set of vectors $\{f(t,a)\}_{t\in\mathcal{T},a\in\mathcal{A}}$ is linearly independent and the row rank of D is TA. Since the row rank of a matrix coincides with its column rank (Fischer (2005, p. 99)) the column space of D spans the whole $\mathbb{R}^{TA}$. In particular, b is in the column space of D, and we can conclude that there exists a $w \in \mathbb{R}^n$ solving the system.

Note that by construction, after receiving contract offer w, the agent implements action $a_i = a^{FB}(t_i)$ under any principal type $t \in \mathcal{T}$ and consequently for any equilibrium belief $(\mu(t|w))_{t\in\mathcal{T}}$ she may hold. Thus, if type t_i proposes contract w, he obtains first-best profits minus $\frac{\epsilon}{2}$. Thus, he has a profitable deviation.

Finally, assume that type $t_i \in \mathcal{T}$ obtains strictly higher equilibrium profits than first-best. Obviously, this cannot happen in a separating equilibrium, because the principal's type is revealed. So assume type t_i belongs to a subgroup $\hat{\mathcal{T}} \subset \mathcal{T}$ that offers a pooling contract $\hat{w}$ in equilibrium, responded to with action $\hat{a}$ by the agent. Because type t_i obtains profits strictly above first-best, it holds that

$$\begin{aligned} f(t_i,\hat{a}) \cdot (\pi - \hat{w}) \quad &> \quad f(t_i, a^{FB}(t_i)) \cdot \pi - \left[\bar{U} + c(a^{FB}(t_i))\right] \\ &\geq \quad f(t_i,\hat{a}) \cdot \pi - \left[\bar{U} + c(\hat{a})\right], \end{aligned}$$

which implies $f(t_i, \hat{a}) \cdot \hat{w} < \bar{U} + c(\hat{a})$. Since the agent's participation constraint for action $\hat{a}$ has to be satisfied in expectation over types in $\hat{\mathcal{T}}$, there exists a type $t_j \in \hat{\mathcal{T}}$ with $f(t_j, \hat{a}) \cdot \hat{w} > \bar{U} + c(\hat{a})$. Thus, we can bound his profits by

$$\begin{aligned} f(t_j, \hat{a}) \cdot (\pi - \hat{w}) &< f(t_j, \hat{a}) \cdot \pi - \left[\bar{U} + c(\hat{a})\right] \\ &\leq f(t_j, a^{FB}(t_j)) \cdot \pi - \left[\bar{U} + c(a^{FB}(t_j))\right], \end{aligned}$$

i.e. type t_j obtains strictly less than his first-best profits. By the above reasoning, this type has a profitable deviation. □

Theorem 1 parallels results from Riordan and Sappington (1988) as well as Crémer and McLean (1988). While they also show a linear independence assumption to render private information inconsequential for the contract proposing party, their works differ in at least two points: First, regarding the private information of the agent, they focus on an adverse selection problem, whereas in our framework it is one of hidden action. Second, they demonstrate how the agent's information can become irrelevant; in our setup it is the principal's information. In fact, the principal has no private information at all in their models. Importantly, in our model the principal's contract signals information about his type to the agent. To our knowledge irrelevance conditions in an informed principal problem with moral hazard have not yet been studied and are novel.

Signaling models are well known for harboring a large number of equilibria, since, without further refinements, no restrictions on off-equilibrium path beliefs are imposed. Thus, unlike the two papers mentioned above, we do not provide necessary conditions for the irrelevance of the principal's information on this level of generality. In Section 2.5 we do, however, establish a full characterization result for a special case of the model.

Obviously, a necessary condition for linear independence of the set of distribution vectors over profits is that the number of profit states is relatively large compared to the action and type space, i.e. $n \geq TA$. While crucial for our model, we postpone the discussion of this issue until the end of Section 2.5. By then we have finished our analysis of the risk neutral case and further links related to the linear independence assumption can be discussed.

Theorem 1 also relies on the agent's risk neutrality. Even though linear independence would be a sufficient condition for implementing a type's prefered action—in the proof of our theorem we can just replace the wage vector with a utility vector in the equation system, solve it, and retrieve from the solution the associated wage vector—it does not provide him in general with his payoffs from the standard model. A risk averse agent does not only care about the expected wage payment, but also about her exposure to risk. If the class of contracts that implements the type's favored action entails more risk than the contract in the standard model, this principal type cannot enforce this action without higher expenses breaking the logic of Theorem 1. We investigate these issues in Section 2.6.

2.5 Analysis: Special Case, Risk Neutral Agent

In this section, we analyze a special case of the model. We maintain the agent's risk neutrality, but allow for only two types of principals and two actions, a high and a low one. Moreover, there is no private information of the principal regarding the low effort choice of the agent, i.e. the distribution over profits given the low action is the same for both types of principals. We also assume that the high action is first-best. One interpretation of these assumptions is that the low action represents a rusty business-as-usual practice. The agent can adapt this policy from his predecessor and its stochastic impact on profits is known from the past. In contrast, the high action can be seen as an innovative and more profitable way to lead the company, but involves considerably more risk for the agent.

In this framework, we provide a full characterization of the irrelevance of the principal's information and show equilibrium existence. We then relate the number of profit levels to distortive effects of the principal's information. To illustrate this point we reproduce two of such distortions—the agent obtains a surplus beyond her reservation utility and the wage spread of her contract is reduced—for the case of two profit levels ($n = 2$). Both effects can be found in related models of the literature and we briefly compare them to our work.

The technical adjustments to our model are as follows: Let $\mathcal{T} = \{\mathbb{P}, \mathbb{Q}\}$ be the set of principal types and the high and low action in $\mathcal{A} = \{a^h, a^l\}$ satisfy $c(a^h) = c > 0$ and $c(a^l) = 0$. For the sake of convenience, we define $p^h := f(\mathbb{P}, a^h)$, $q^h := f(\mathbb{Q}, a^h)$ and $p^l := f(\mathbb{P}, a^l) = f(\mathbb{Q}, a^l)$; we assume, however, that the three vectors p^h, q^h, p^l are mutually distinct. Moreover, the high action a^h is first-best (or, equivalently, second-best) for both types of the principal, that is $a^{FB}(\mathbb{P}) = a^{FB}(\mathbb{Q}) = a^h$.

Before we look at the characterization theorem of this section, we establish a technical auxiliary result. It shows that linear dependence of three distribution vectors in $\mathbb{R}^n$ allows a representation of one vector as convex combination of the other two; this follows from their normalization to unity.

Lemma 1. *Let $p, q, r \in \mathbb{R}^n_+$ be mutually distinct probability distributions. Then the set of vectors $\{p, q, r\}$ is linearly dependent if and only if one vector is a convex combination of the other two with a unique weight $\alpha \in (0, 1)$.*

PROOF: The proof is relegated to the appendix. □

In the special case of our model, linear independence of the set of profit distributions is a necessary and sufficient condition for effecting the first-best outcome for the principal in any equilibrium. Sufficiency is less surprising, because Theorem 1 can almost immediately be applied. Necessity follows, since equilibrium conditions impose substantial diversity on the payoff structure of the equilibrium contracts, which is incompatible with linear dependence. For instance, consider a separating equilibrium in which both principal types obtain their first-best profits: The equilibrium contract of the $\mathbb{P}$-type exactly compensates the agent for his reservation utility and the effort cost under p^h, because the $\mathbb{P}$-type obtains first-best profits. Under q^h, however, it has to pay the agent a higher expected wage—otherwise the $\mathbb{Q}$-type has an incentive to mimic the $\mathbb{P}$-type's equilibrium offer. And under p^l the $\mathbb{P}$-type's contract has to pay much less than the reservation utility and effort cost to incentivize the agent to choose the high action, which is first-best. Now assume that the three distribution vectors are linearly dependent. By Lemma 1, this implies that one vector is a convex combination of the other two distributions. It follows that the expected wage payment of this vector must be a convex combination of the payments of the other two vectors. The above reasoning shows that p^h is necessarily a convex combination of q^h and p^l. The same logic, however, also applies to the equilibrium contract of the $\mathbb{Q}$-type implying that q^h has to be a convex combination of p^h and p^l. This yields a contradiction and the probability distributions have to be linearly independent.

Theorem 2. *Assume the agent is risk neutral. The set of vectors $\{p^h, q^h, p^l\}$ is linearly independent if and only if both types of the principal obtain their first-best profits in any Perfect Bayesian Equilibrium.*

Proof: For the only-if part, consider the proof of Theorem 1. Because $f(\mathbb{P}, a^l) = f(\mathbb{Q}, a^l)$ we can drop the condition $f(\mathbb{Q}, a^l)\cdot w = \bar{U}+c(a^l)$ from the equation system $D\cdot w = b$. Linear independence of $\{p^h, q^h, p^l\}$ is then sufficient for the existence of a deviation contract. The rest of the proof remains unaffected, so that its application is warranted here.

For the if part of the statement, we prove the case of pooling equilibria first before we turn to separating equilibria. Note that in any case action a^h is implemented in an equilibrium that provides both types of the principal first-best payoffs. If the agent chooses action a^l in response to an equilibrium offer of w, at least one type of principal, say the $\mathbb{P}$-type, has to provide the agent with a minimum expected remuneration of $\bar{U}$. This implies, however, that the $\mathbb{P}$-type's equilibrium payoff is strictly below first-best, because

$$p^l \cdot (\pi - w) \ \leq \ p^l \cdot \pi - \bar{U} \ < \ p^h \cdot \pi - \left[\bar{U} + c\right].$$

Thus, in any equilibrium with first-best payoff for the principal, action a^h is chosen by the agent and each type of principal pays $\bar{U} + c$ in expectation to her.

First, consider an equilibrium with pooling contract w^{pool} that implements first-best payoffs for each type of principal. Because action a^h is weakly prefered to a^l by the agent, contract w^{pool} satisfies

$$\begin{pmatrix} p^h \\ q^h \\ p^l \end{pmatrix} \cdot w^{pool} \quad = \quad \begin{pmatrix} \bar{U} + c \\ \bar{U} + c \\ \bar{U} - \epsilon \end{pmatrix}$$

for some $\epsilon \in \mathbb{R}_+$. Obviously, neither distribution can be written as a convex combination of the other two without violation of any equation. By Lemma 1, the set of vectors $\{p^h, q^h, p^l\}$ must be linearly independent.

Now assume there exists an equilibrium with separating contracts $(w_{\mathbb{P}}, w_{\mathbb{Q}})$ that implement first-best payoffs for the principal types. Compared to the pooling case, for each equilibrium contract there is the additional requirement that no type mimics the other type, i.e. $w_{\mathbb{P}}$ and $w_{\mathbb{Q}}$ have to satisfy

$$\begin{pmatrix} p^h \\ q^h \\ p^l \end{pmatrix} \cdot w_{\mathbb{P}} \quad = \quad \begin{pmatrix} \bar{U} + c \\ \bar{U} + c + \delta_{\mathbb{P}} \\ \bar{U} - \epsilon_{\mathbb{P}} \end{pmatrix}$$

and

$$\begin{pmatrix} p^h \\ q^h \\ p^l \end{pmatrix} \cdot w_{\mathbb{Q}} \quad = \quad \begin{pmatrix} \bar{U} + c + \delta_{\mathbb{Q}} \\ \bar{U} + c \\ \bar{U} - \epsilon_{\mathbb{Q}} \end{pmatrix}$$

for some $(\delta_{\mathbb{P}}, \epsilon_{\mathbb{P}}, \delta_{\mathbb{Q}}, \epsilon_{\mathbb{Q}}) \in \mathbb{R}^4_+$. Now assume that the set $\{p^h, q^h, p^l\}$ is linearly dependent. By Lemma 1, one distribution must be a convex combination of the other two, and the equations of contract $w_{\mathbb{P}}$ show that this must be p^h. But this violates the set of equations associated with $w_{\mathbb{Q}}$. Thus, the set $\{p^h, q^h, p^l\}$ is linearly independent. □

We now provide an equilibrium existence result under the assumption of linear independence of the distribution vectors. We know from Theorem 2 that our search for equilibrium contracts can be restricted to offers that provide at least one type of the principal with his first-best profits. A subset of this class is particularly convenient for an equilibrium construction; it consists of contracts that generate first-best profits for both types of the principal. Such offers are not only necessary for the existence of pooling equilibria, but also helpful in separating ones, because, by design, each type has no incentive to mimic the other's equilibrium contract. Let us define this set by

$$W_{\mathbb{P},\mathbb{Q}} \quad := \quad \{w \in \mathbb{R}^n \mid p^h \cdot w = \bar{U} + c, \ q^h \cdot w = \bar{U} + c, \ p^l \cdot w \leq \bar{U}\}.$$

Under our linear independence assumption this set is nonempty. Again, this is due to the large range of outcome vectors that can be attained under linear independence.

Lemma 2. *If the set of vectors $\{p^h, q^h, p^l\}$ is linearly independent, then $W_{\mathbb{P},\mathbb{Q}}$ is nonempty.*

Proof: We have to show that there exists a $w \in \mathbb{R}^n$ satisfying

$$\begin{pmatrix} p^h \\ q^h \\ p^l \end{pmatrix} \cdot w = \begin{pmatrix} \bar{U} + c \\ \bar{U} + c \\ \bar{U} - \epsilon \end{pmatrix}$$

for some $\epsilon \in \mathbb{R}_+$. By following the proof of Theorem 1 it is immediate that such a $w \in \mathbb{R}^n$ exists. □

Example 1. Let $p^h = (\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$, $q^h = (\frac{1}{6}, \frac{2}{3}, \frac{1}{6})$ and $p^l = (\frac{1}{6}, \frac{1}{6}, \frac{2}{3})$. Moreover, assume $\bar{U} = 0$, $c = 0.5$ and $\pi = (3, 2, 0)$. A brief calculation shows that implementing action a^h is first-best. Also, it is immediately clear from Lemma 1 that the vectors p^h, q^h and p^l are linearly independent, because each contains for a particular state the highest probability among the distributions. In this example, $W_{\mathbb{P},\mathbb{Q}}$ has the parametrized form $W_{\mathbb{P},\mathbb{Q}} = (\frac{2}{3}, \frac{2}{3}, -\frac{1}{3}) + \epsilon\ (\frac{1}{3}, \frac{1}{3}, -\frac{5}{3})$ where $\epsilon \geq 0$.

We will use contracts in $W_{\mathbb{P},\mathbb{Q}}$ as candidates for an equilibrium construction. Thus, on the equilibrium path both types obtain their first-best profits. What remains to be done, is to find an adequate choice of off-equilibrium path beliefs that rules out any deviation opportunities. Note that a principal type can only improve on his first-best profits if the agent accepts a contract and selects a specific action, because she believes to face the other type. Fortunately, this is impossible whenever the agent chooses the low action a^l; since there's no private information regarding the effort-profit relationship, each type that offers such a contract compensates the agent at least with her reservation utility and therefore cannot improve on his first-best profits. This implies that profitable deviations only occur if the agent selects the high action a^h. Thus, for each type a deviation set is identified by contracts that have him pay less than $\bar{U} + c$, but the other type more than $\bar{U} + c$, if the agent chooses a^h. A simple way to avoid the agent's choice of a^h for these contracts, is to assign the belief that the agent indeed faces the principal type paying her less than $\bar{U} + c$. Hence, she will either reject the contract or select a^l, and no type of principal has an incentive to switch to these contracts. Consequently, this constraint on off-equilibrium path beliefs establishes equilibrium existence.

Theorem 3. *Assume the agent is risk neutral. If the set of vectors $\{p^h, q^h, p^l\}$ is linearly independent, then Perfect Bayesian Equilibria in pooling and separating strategies exist. Moreover, these equilibria satisfy the Intuitive Criterion.*

Proof: For each type of principal define a deviation set

$$\begin{aligned} W_{\mathbb{P}}^{dev} &:= \left\{ w \in \mathbb{R}^n \mid p^h \cdot w < \bar{U} + c,\ q^h \cdot w \geq \bar{U} + c \right\}, \\ W_{\mathbb{Q}}^{dev} &:= \left\{ w \in \mathbb{R}^n \mid q^h \cdot w < \bar{U} + c,\ p^h \cdot w \geq \bar{U} + c \right\}. \end{aligned}$$

It is easy to check that

$$W_{\mathbb{P}}^{dev} \cap W_{\mathbb{Q}}^{dev} = \emptyset, \qquad W_{\mathbb{P}}^{dev} \cap W_{\mathbb{P},\mathbb{Q}} = \emptyset, \qquad W_{\mathbb{Q}}^{dev} \cap W_{\mathbb{P},\mathbb{Q}} = \emptyset.$$

First, consider the case of separating equilibria. Assume that the $\mathbb{P}$- and $\mathbb{Q}$-type offer contracts $w_{\mathbb{P}} \in W_{\mathbb{P},\mathbb{Q}}$ and $w_{\mathbb{Q}} \in W_{\mathbb{P},\mathbb{Q}}$ respectively, where $w_{\mathbb{P}} \neq w_{\mathbb{Q}}$. Such contracts exist by Lemma 2.[6] The agent is assumed to accept each offer and implement action a^h. Off the equilibrium path the agent chooses for each contract the option that generates the highest expected utility given her belief for this contract: (1) select a^h, (2) select a^l, or (3) reject the offer. In case of any indifference she is assumed to select the option with the lowest number.

The belief system is structured as follows: Since the principal's type is revealed on the equilibrium path $\mu(\mathbb{P}|w_{\mathbb{P}}) = 1$ and $\mu(\mathbb{Q}|w_{\mathbb{Q}}) = 1$. Off the equilibrium path beliefs are specified as

$$\begin{aligned} &\mu(\mathbb{P}|w) = 1 && \text{for all } w \in W_{\mathbb{P}}^{dev}, \\ &\mu(\mathbb{Q}|w) = 1 && \text{for all } w \in W_{\mathbb{Q}}^{dev}, \\ &\mu(\mathbb{P}|w) = x(w) && \text{for all } w \in \mathbb{R}^n \setminus \left(W_{\mathbb{P}}^{dev} \cup W_{\mathbb{Q}}^{dev} \cup \{w_{\mathbb{P}}, w_{\mathbb{Q}}\} \right), \end{aligned}$$

where $x(w) \in [0,1]$. Note that such a belief assignment is valid by the above disjointness properties.

We now verify that these strategies and beliefs form a Perfect Bayesian Equilibrium. It is readily checked that the beliefs are updated according to Bayes' Rule whenever possible.

[6] Because the choice of $\epsilon \in \mathbb{R}_+$ was arbitrary in the proof of Lemma 2, the set $W_{\mathbb{P},\mathbb{Q}}$ is uncountably infinite and we can find at least two distinct contracts.

In equilibrium both types of the principal obtain their first-best payoff and the agent receives her outside option value of $\bar{U}$. The agent has no profitable deviation: Because $w_{\mathbb{P}}, w_{\mathbb{Q}} \in W_{\mathbb{P},\mathbb{Q}}$ the agent cannot improve by choosing a^l or rejecting either contract. By construction the agent's behavior is also optimal off the equilibrium path. Regarding the principal, we consider deviations for each action of the agent separately. We start with the $\mathbb{P}$-type of the principal. First, note that he does not deviate to contracts $w \in \mathbb{R}^n$ that induce the agent to choose a^l. Because the distribution over profits is the same for both types under a^l, he has to provide the agent with an expected wage of at least $\bar{U}$. His profits are then strictly less than first-best because

$$p^l \cdot (\pi - w) \;\; \leq \;\; p^l \cdot \pi - \bar{U} \;\; < \;\; p^h \cdot \pi - \left[\bar{U} + c\right].$$

Now consider deviation contracts if the agent responds with action a^h. To attain higher profits than first-best, the $\mathbb{P}$-type has to pay strictly less than $\bar{U}+c$ to the agent. Certainly, action a^h is not implemented for $w \in \{\tilde{w} \in \mathbb{R}^n \mid p^h \cdot \tilde{w} < \bar{U} + c,\ q^h \cdot \tilde{w} < \bar{U} + c,\}$, because for any belief the agent would rather select a^l or reject the contract. But action a^h is also not chosen for $w \in \{\tilde{w} \in \mathbb{R}^n \mid p^h \cdot \tilde{w} < \bar{U} + c,\ q^h \cdot \tilde{w} \geq \bar{U} + c,\} = W_{\mathbb{P}}^{dev}$, since she believes to face the $\mathbb{P}$-type given our belief assignment. It follows that the $\mathbb{P}$-type has no profitable deviation among contracts responded with a^h. Hence, equilibrium contract $w_{\mathbb{P}}$ is optimal for him. Due to the symmetry of the equilibrium construction, the $\mathbb{Q}$-type does not have a profitable deviation either. Thus, we have found a strategy profile in which players behave sequentially rational, and a belief system with beliefs updated by Bayes' Rule whenever possible, i.e. we have constructed a Perfect Bayesian Equilibrium.

The existence proof for pooling equilibria is completely analogous to the case of separating equilibria with the adjustments $w_{pool} = w_{\mathbb{P}} = w_{\mathbb{Q}}$ and equilibrium beliefs $\mu(\mathbb{P}|w_{pool}) = \lambda_{\mathbb{P}}$ and $\mu(\mathbb{Q}|w_{pool}) = \lambda_{\mathbb{Q}} = 1 - \lambda_{\mathbb{P}}$. We show in the appendix that our equilibrium constructions satisfy the Intuitive Criterion. □

We now examine more closely the assumption of linear independence of the distribution vectors. A necessary condition for this property to occur is that the number of profit states is large compared to the size of the action and type space ($n \geq 3$ in the simplified model, $n \geq TA$ in the general case). In the following, we want to elaborate on two points in that respect.

First, if there are many profit states relative to actions and types, we argue that linear independence of the distributions is not a strong assumption—a triple of distributions $p, q, r \in \mathbb{R}^n$ is almost always linearly independent. For an intuition imagine two vectors in $\mathbb{R}^2$ as minute- and hour-hand of a clock; generically, their position does not form a straight line. As formal criterion, we use the genericity condition of Fudenberg and Tirole (2002, p. 479), which requires the set of linearly independent distribution triples to be open and dense in the set of all distribution triples.[7] That said, this rationale implicitly presumes that a priori there are no reasons, why an occurence of a linearly dependent triple is particularly likely. If this is the case, we obtain

Proposition 1. *Assume $n \geq 3$. Generically, any triple of distributions $p, q, r \in \mathbb{R}^n$ is linearly independent.*

PROOF: The proof is relegated to the appendix. □

Second, we highlight the importance of the number of profit states for the distortive nature of the principal's information. By Proposition 1 and Theorem 2, if the number of profit states is large compared to actions and types, the distribution vectors are generically linearly independent, and the principal's information is generically irrelevant. In contrast, for a low number of profit levels, the distribution vectors are necessarily linearly dependent, and there must be a distortive effect from the information asymmetry. More precisely, in at least one Perfect Bayesian Equilibrium the principal's first-best payoff is not realized.[8] Thus, there exists a relationship between the number of available signal states and the relevance of the principal's information; to our knowledge this connection has not yet been studied. For an intuition, observe that the principal is unable to adequately fine tune the agent's monetary incentives if the profit signal is not sufficiently pronounced. This prevents the principal from attaining first-best profits in two dimensions. First, the set of profitable deviations is reduced. For instance, as in the proof of Theorem 1, there do not exist contracts that induce the agent to choose the principal's type first-

[7] Denseness of a set without openness does not fully comprise the essence of *almost all*. For instance, the set of rational numbers $\mathbb{Q}$ is not open, but dense in $\mathbb{R}$, because any irrational number can be approximated by a sequence in $\mathbb{Q}$. But the complement of $\mathbb{Q}$ in $\mathbb{R}$, the irrational numbers, is much larger (uncountably infinite) than $\mathbb{Q}$ (countably infinite).

[8] This conclusion cannot be drawn solely from Theorem 1. While a relatively low number of profit levels implies linear dependence, this need not lead to a distortive effect. The principal's types could obtain first-best profits in any equilibrium despite linearly dependent distribution vectors.

best action independent of her belief and still provide this type with first-best profits. Therefore, the principal can wind up in equilibria with strictly higher or lower payoffs for him, because profitable deviation contracts can be made unattractive by imposing an appropriate belief system.[9] More importantly, second, the set of equilibrium contracts for first-best implementations is smaller. For instance, if $n = 2$, no pooling equilibrium exists, since the only contract that compensates the agent with $\bar{U} + c$ under p^h and q^h is the contract that pays $\bar{U} + c$ in both profit states—a full insurance contract. Clearly, such an offer does not provide any incentives to the agent to select the first-best action a^h.

In the following we demonstrate how a distortive effect arises in our model if the number of profit states is low ($n = 2$). We focus on equilibria satisfying the Intuitive Criterion to obtain a sharper prediction. Thereafter, we relate our results to the papers of Beaudry (1994) and Inderst (2001).

So assume that there are only two profit levels $\pi_1 > \pi_2$ and $p_1^h > q_1^h > p_1^l$, that is, the $\mathbb{P}$-type generates larger gross profits with action a^h than the $\mathbb{Q}$-type. The basic features of this setup can be explained best by considering Figure 2.1.[10] First, note that the vector π is situated below the 45°-line since $\pi_1 > \pi_2$. Note also that the respective selling-the-firm contracts $w_{\mathbb{P}}^{stf}$ and $w_{\mathbb{Q}}^{stf}$ can be found on a parallel to the 45°-line that runs through π, because these contracts transfer the complete profit realizations to the agent minus a state-independent payment to the principal. Since the $\mathbb{P}$-type generates a higher gross profit, he can charge a larger prize for the firm and his selling-the-firm contract specifies a lower absolute wage payment in both states of the world than the one of the $\mathbb{Q}$-type. Under both contracts the agent is just compensated for his effort and his outside option. Thus, two isowage curve can be drawn through $w_{\mathbb{P}}^{stf}$ and $w_{\mathbb{Q}}^{stf}$, which characterize all pairs of wage levels that give an expected wage payment of $\bar{U} + c$ to the agent under the distributions p^h and q^h.

There is also a third isowage line: it depicts all contracts that provide the agent with an expected payment of $\bar{U}$ if he chooses action a^l. This curve crosses the parallel of the 45°-

[9] Note that a principal type can only obtain higher profits than first-best in a pooling contract at the expense of another type. Thus, if not all principal types obtain their first-best payoff in equilibrium, at least one type gets strictly less than first-best and has an incentive to deviate to a contract as described above.

[10] All values that refer to the first entry of a vector are depicted on the horizontal axis whereas those of the second entry lie on the vertical axis.

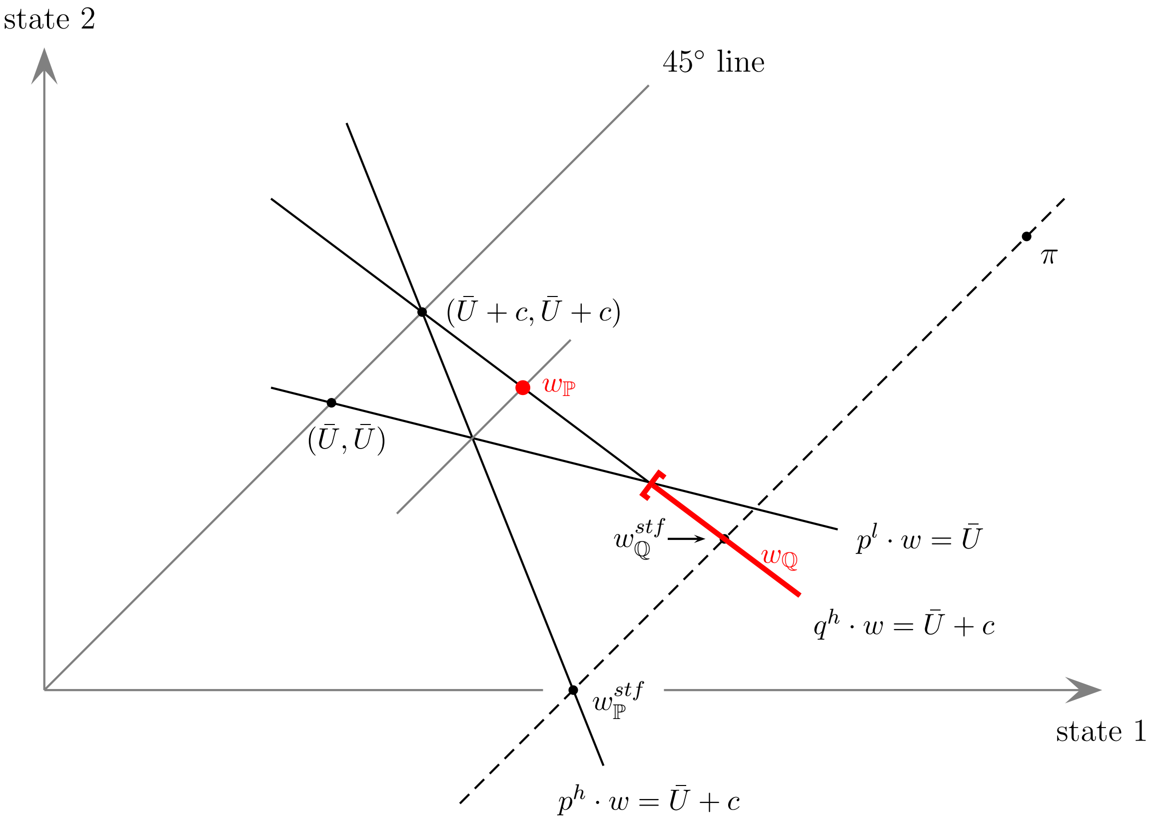

Figure 2.1: Graphical representation of the simplified model for $n = 2$ in which contract vectors can be represented by a point in $\mathbb{R}^2$. Full insurance contracts lie on the 45°-line, whereas the selling-the-firm contracts $w_{\mathbb{P}}^{stf}$ and $w_{\mathbb{Q}}^{stf}$, respectively, are situated on a parallel running through the gross profit vector π. Three further lines are depicted; for each distribution p^h, q^h and p^l isowage curves are drawn, yielding the agent a respective payment of $\bar{U}+c$ and $\bar{U}$. Equilibrium contracts satisfying the Intuitive Criterion are displayed in red color. While there is only a unique contract $w_{\mathbb{P}}$ for the $\mathbb{P}$-type, all contracts on the red line can serve as equilibrium offers $w_{\mathbb{Q}}$ for the $\mathbb{Q}$-type.

line closer to π than both selling-the-firm contracts; otherwise the benefit of implementing a^h does not outweigh its costs and action a^h is not first-best for both types of principals.

Certainly, it is not an equilibrium if both types offer their selling-the-firm contract and the agent selects action a^h for both offers. The $\mathbb{Q}$-type would have an incentive to deviate to the selling the firm contract of the $\mathbb{P}$-type, because this contract lies on a lower isowage curve. In order to separate himself from the $\mathbb{Q}$-type, the $\mathbb{P}$-type, starting at his selling-the-firm contract, moves upwards his isowage curve towards the 45°-line. By decreasing the payment in the high profit state and increasing it for the low one, the $\mathbb{P}$-type keeps his expected payment constant, but reduces the wage spread. The $\mathbb{Q}$-type, however, feels increasingly dissatisfied with these contracts. Since he has a lower probability for high profit states (and a higher one for low states), his average payment to the agent is increasing.[11] Yet this change in the contract structure is not sufficient for a separation of types. Lower wage spreads provide less incentives for the agent to exert costly effort, and eventually, her incentive constraint becomes binding (intersection of the isowage curves of p^h and p^l). By a movement parallel to the 45°-line, the $\mathbb{P}$-type advances separation, but also leaves the agent's individual rationality slack. Not before the $\mathbb{P}$-type's contract offer has reached the isowage-line of the $\mathbb{Q}$-type, the least-cost separating contract, the $\mathbb{Q}$-type has no incentives to mimic the $\mathbb{P}$-types's offer anymore, and an equilibrium can be sustained. Equilibrium contracts $(w_{\mathbb{P}}, w_{\mathbb{Q}})$ that satisfy the Intuitive Criterion are marked in red color. Unlike $w_{\mathbb{P}}$ there are infinitely many candidates for $w_{\mathbb{Q}}$ that can be supported as equilibrium contracts. We summarize this result in

Proposition 2. *Let $n = 2$. In any Perfect Bayesian Equilibrium that satisfies the Intuitive Criterion, the agent implements action a^h and the equilibrium strategies $(w_{\mathbb{P}}, w_{\mathbb{Q}})$ satisfy*

$$\begin{aligned} p^h \cdot w_{\mathbb{P}} &> \bar{U} + c, & (p^h - p^l) \cdot w_{\mathbb{P}} &= c, \\ q^h \cdot w_{\mathbb{Q}} &= \bar{U} + c, & (q^h - p^l) \cdot w_{\mathbb{Q}} &\geq c, \\ q^h \cdot w_{\mathbb{P}} &= \bar{U} + c. \end{aligned}$$

In particular, no pooling equilibrium exists.

Proof: The proof is relegated to the appendix. □

[11] Note that a movement along the isowage curve away from the 45°-line does not allow the $\mathbb{P}$-type to generate this effect.

Note that the $\mathbb{P}$-type's equilibrium contract $w_{\mathbb{P}}$ exhibits a lower wage spread than $w_{\mathbb{P}}^{stf}$, because parallels to the 45°-line denote contracts that have the same wage spread. In addition, his signaling activity provides the agent with a surplus beyond her outside option. Thus, the $\mathbb{P}$-type is unable to attain first-best profits and the principal's information acts distortive. These effects also arise in two related informed principal models with moral hazard, which basically share all features of our binary framework except continuous effort. In Inderst (2001) the low wage spread of the $\mathbb{P}$-type's contract is put forward as an explanation for the prevalence of low-powered incentive schemes in reality. In contrast, Beaudry (1994) emphasizes the agent's surplus and sees it as a rationale for efficiency wages. By that means, he counters criticism against agency theory that claims that it provides no explanation why an employer would leave ex ante rents to an employee.

Even though we can accomodate these results in our model, they heavily rely on the restriction to two profit states. In particular, if $n \geq 3$, both principal types obtain their first-best payoff and the agent does not enjoy a net surplus beyond her outside option value. Also, equilibrium contracts need not feature a low wage spread.[12] To see this, consider Example 1 and the set $W_{\mathbb{P},\mathbb{Q}}$ parametrized by $\epsilon \geq 0$. When ϵ increases the associated wage grows to infinity in the first and second profit state, and approaches minus infinity in the third one, i.e. the incentives of the contract become very accentuated. But the proof of Theorem 3 shows that any contract in $W_{\mathbb{P},\mathbb{Q}}$ can serve as an equilibrium offer. Thus, equilibrium contracts of both principal types can be of arbitrarily strong power. Note that this feature is not restricted to our specific example and is generally true for any $n \geq 3$. We omit the general result.

All results of this section were derived for a special case of the model. In particular, we assumed that both types of the principal agreed on the first-best action a^h, and the agent had perfect knowledge about the relationship between effort and profits for the low action a^l. Both assumptions are not innocuous and are employed in Theorem 2 and 3. Nevertheless, we do not feel too concerned about the restrictiveness of these results; we

[12] There is a caveat regarding the reduced wage spread if $n = 2$. We have implicitly assumed that the $\mathbb{P}$-type would select $w_{\mathbb{P}}^{stf}$ in a second-best situation without private information of the principal. In fact, an infinite number of contracts could serve as second-best contracts, and in particular contracts with a low wage spread (c.f. Figure 2.1). We argue, however, that $w_{\mathbb{P}}^{stf}$ is an especially salient contract and may well serve as a reasonable benchmark. This feature does not arise in the Inderst (2001) model, because he assumes effort to be continuous.

rather see them as first steps towards *completing* the analysis of our general model. A more detailed discussion of the issues is presented in Section 2.7.

2.6 Analysis: Special Case, Risk Averse Agent

In this section, we consider our special case model when the agent is strictly risk averse. We show that the principal's private information can never be advantageous for him and provide an example in which he is strictly worse off. For our general result we strengthen, however, the solution concept by focusing on Perfect Bayesian Equilibria satisfying the Intuitive Criterion. But first, we continue our discussion of the standard moral hazard model from Section 2.4 under the assumption of a strictly risk averse agent.

For our purpose, there are two important differences that arise in the standard model when moving from a risk neutral to a strictly risk averse agent. First, unless the principal implements the least-cost action, he obtains strictly lower profits in a second-best than in a first-best situation. (Grossman and Hart (1983), Proposition 3(6)).[13] Intuitively, inducing the agent to select a specific action requires the principal to incentivize the agent with a nonconstant wage scheme. Due to the agent's risk aversion, however, her exposure to risk has to be compensated with an expected payment exceeding the one in the first-best situation, where a flat wage is feasible. Thus, a real tradeoff between providing incentives and insuring the agent exists. Second, strict risk aversion of the agent implies that the contract that implements the second-best action is unique (their Remark 2).

Returning to the informed principal model, we maintain our simplifying assumption that action a^h is second-best for both principal types, i.e. $a^{SB}(\mathbb{P}) = a^{SB}(\mathbb{Q}) = a^h$, and assume it is implemented by the unique second-best contracts $w_{\mathbb{P}}^{SB}$ and $w_{\mathbb{Q}}^{SB}$ respectively. Moreover, as before, there is no private information of the principal regarding the low action a^l, i.e. $p^l = f(\mathbb{P}, a^l) = f(\mathbb{Q}, a^l)$ and we denote with $w^l := (\tilde{u}^{-1}(\bar{U}), \ldots, \tilde{u}^{-1}(\bar{U}))$ the flat-wage contract for implementing action a^l at least cost.[14]

[13] To implement the least-cost action, the principal simply offers a flat wage scheme as in the first-best situation, which compensates the agent in any state of the world for her reservation utility and effort cost. Because this contract does not provide any incentives except participation, the agent selects the least-cost action, and, since this contract fully insures the agent, there are no additional implementation costs for the principal.

[14] Thus, the distributions p^h, q^h and p^l satisfy $p^h \cdot (\pi - w_{\mathbb{P}}^{SB}) > p^l \cdot (\pi - w^l)$ and $q^h \cdot (\pi - w_{\mathbb{Q}}^{SB}) > p^l \cdot (\pi - w^l)$.

Under strict risk aversion linear independence of the set $\{p^h, q^h, p^l\}$ is no longer sufficient to guarantee the principal his second-best profits. As in Theorem 1 the agent can still be incentivized to choose a specific action—independent of her beliefs about the principal's type—but this may not be possible at second-best costs. The agent not only values her expected payment, but also her exposure to risk, and linear independence does not account for this. This problem also extends to the construction of equilibria. In fact, we now apply Example 1, which employs linearly independent distribution vectors, to the risk averse case and show that in any equilibrium at least one type of the principal does not obtain his second-best profits. Thus, the principal's information acts distortive despite linear independence.

Example 2. As in Example 1 let $p^h = (\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$, $q^h = (\frac{1}{6}, \frac{2}{3}, \frac{1}{6})$, and $p^l = (\frac{1}{6}, \frac{1}{6}, \frac{2}{3})$. Again, assume $\bar{U} = 0$, $c = 0.5$ and $\pi = (3, 2, 0)$, but now the agent's utility function is given by $\tilde{u}(w) = \ln(w)$. We use numerical methods to obtain the second-best contracts $w_{\mathbb{P}}^{SB} = (1.970, 1.913, 0.718)$ and $w_{\mathbb{Q}}^{SB} = (1.740, 1.992, 0.733)$ respectively.[15]

We now show that it is impossible that both types obtain their second-best profits. First, consider separating equilibria. Because second-best contracts are unique and both types are revealed in a separating equilibrium, the contracts $w_{\mathbb{P}}^{SB}$ and $w_{\mathbb{Q}}^{SB}$ must be offered in equilibrium and action a^h is chosen by the agent. Otherwise both types do not obtain their second-best profits. But since $p^h \cdot w_{\mathbb{Q}}^{SB} = 1.656 < 1.742 = p^h \cdot w_{\mathbb{P}}^{SB}$, the $\mathbb{P}$-type has a profitable deviation: He has an incentive to mimic the $\mathbb{Q}$-type's contract, because by offering $w_{\mathbb{Q}}$ the $\mathbb{P}$-type implements a^h at much lower costs than under his own equilibrium contract.

For the case of pooling equilibria, note that uniqueness of the second-best contracts is not sufficient to rule out second-best profits for both types, because the individual rationality and incentive constraint have to be satisfied only in expectation. So, for instance, a violation of the (IR) by the $\mathbb{P}$-type could be compensated by a slack (IR) of the $\mathbb{Q}$-type and vice versa for the (IC). Therefore, we consider all contracts that provide both types with second-best profits, i.e. we look at contracts of the set $\tilde{W}_{\mathbb{P},\mathbb{Q}} := \{w \in \mathbb{R}^3 \mid p^h \cdot w = p^h \cdot w_{\mathbb{P}}^{SB},\ q^h \cdot w = q^h \cdot w_{\mathbb{Q}}^{SB}\} = (1.939, 1.932, 0.774) - z(\frac{1}{5}, \frac{1}{5}, -1)$ with $z \in \mathbb{R}$. Let us choose the prior $\lambda_{\mathbb{P}} = \lambda_{\mathbb{Q}} = 0.5$. We scrutinize (IR) and (IC) for all contracts in $\tilde{W}_{\mathbb{P},\mathbb{Q}}$ to see if one

[15] It is readily verified that for both types of the principal, given π, implementing action a^h with $w_{\mathbb{P}}^{SB}$ and $w_{\mathbb{Q}}^{SB}$, respectively, generates higher profits than implementing a^l with $w^l = (1, 1, 1)$.

of these contracts jointly satisfies both constraints. Numerical methods show that (IC) is satisfied only if $z \leq -0.059$, but (IR) is violated for any $z \leq -0.050$. Thus, whenever the incentive constraint is satisfied the individual rationality constraint is violated. It follows that any of the pooling equilibrium candidates does not implement a^h, and second-best profits for both types are not attained.

Both cases together shows that there does not exist a Perfect Bayesian Equilibrium in which both types obtain their second-best profits.

We do not attempt to identify the distributional assumptions that fully characterize the irrelevance of the principal's private information under risk aversion. Instead, by imposing harsher restrictions on off-equilibrium path beliefs with the help of the Intuitive Criterion, we establish upper and lower bounds of the principal's payoff in any Perfect Bayesian Equilibrium. In particular, neither type can improve on his second-best profits, but cannot fall short of his payoff for implementing the least-cost action a^l either. The latter result is quite intuitive, because the agent will always respond a flat-wage offer with action a^l. Regarding the upper-bound the analysis is more involved. Clearly, neither principal type is able to obtain more than his second-best profits in a separating equilibrium. In a pooling equilibrium a violation of the upper-bound allows for a profitable deviation: Assume that one type, say the $\mathbb{P}$-type, obtains strictly more than his second-best profits under the pooling contract. Necessarily, a^h has to be implemented and he violates either his individual rationality or his incenctive or both constraints, which is compensated by a slack constraint of the $\mathbb{Q}$-type. Because there is no private information of the principal regarding a^l one can show that in each case even both constraints of the $\mathbb{Q}$ type have to be slack under a pooling contract. Thus, for contracts close to the pooling contract, the agent implements a^h if he expects to face the $\mathbb{Q}$-type principal. Now, a second aspect comes into play. Since each type is identified with a different distribution over profits, there exist contracts arbitrarily close to the pooling contract that do not allow the $\mathbb{P}$-type to improve on his equilibrium payoff no matter how the agent responds, but do so for the $\mathbb{Q}$-type if the agent implements a^h. By the Intuitive Criterion the agent should believe to face the $\mathbb{Q}$-type for these contracts. But then, due to the slack constraints of the $\mathbb{Q}$-type close to the pooling contract, the agent indeed implements a^h in this contract region, which gives the $\mathbb{Q}$-type a profitable deviation. This yields a contradiction.

Theorem 4. *Assume the agent is strictly risk averse. In any Perfect Bayesian Equilibrium that satisfies the Intuitive Criterion with equilibrium contracts* $(w_{\mathbb{P}}, w_{\mathbb{Q}})$ *and equilibrium actions* $(a(w_{\mathbb{P}}), a(w_{\mathbb{Q}}))$ *it holds*

$$\begin{aligned} p^l \cdot (\pi - w^l) \quad &\leq \quad f(\mathbb{P}, a(w_{\mathbb{P}})) \cdot (\pi - w_{\mathbb{P}}) \quad \leq \quad p^h \cdot (\pi - w_{\mathbb{P}}^{SB}), \\ p^l \cdot (\pi - w^l) \quad &\leq \quad f(\mathbb{Q}, a(w_{\mathbb{Q}})) \cdot (\pi - w_{\mathbb{Q}}) \quad \leq \quad q^h \cdot (\pi - w_{\mathbb{Q}}^{SB}). \end{aligned}$$

Proof: Assume there exists a Perfect Bayesian Equilibrium in which one type obtains equilibrium profits strictly above or below the specified bounds. Without loss of generality let this be the $\mathbb{P}$-type. We will show that at least one type of principal then has a profitable deviation.

First, assume that the $\mathbb{P}$-type receives a payoff strictly below $p^l \cdot (\pi - w^l)$ and let $\epsilon > 0$ be the difference of this lower bound and equilibrium profits. Consider contract $\hat{w}^l := \left(\tilde{u}^{-1}(\bar{U}) + \frac{\epsilon}{2}, \ldots, \tilde{u}^{-1}(\bar{U}) + \frac{\epsilon}{2}\right)$. Because this contract fully insures the agent and gives her a surplus beyond $\bar{U}$, she accepts it and implements action a^l. Thus, if the $\mathbb{P}$-type proposes contract $\hat{w}^l$, he obtains the lower-bound profits minus $\frac{\epsilon}{2}$. Thus, he has a profitable deviation.

Regarding the upper bound, we analyze the cases of separating and pooling equilibria separately. So assume that the $\mathbb{P}$-type receives an equilibrium payoff strictly above his second-best profits. Obviously, this cannot happen in a separating equilibrium, since the principal's type is revealed in equilibrium.

The proof for pooling equilibria is a bit more intricate and we will proceed in several minor steps. We will establish first that the $\mathbb{Q}$-type's individual rationality and incentive constraint are both slack for contracts close to the pooling contract w^{pool}. We then show that we can find a set of contracts arbitrarily close to the pooling contract w^{pool} that give higher-than-equilibrium payoffs to the $\mathbb{Q}$-type if a^h can be implemented, but lower-than-equilibrim payoffs to the $\mathbb{P}$-type independent of the action choice of the agent. Hence, by the Intuitive Criterion, the agent should believe with probability one to face the $\mathbb{Q}$-type when offered such a contract. Then, given this belief and by choosing a contract sufficiently close to w^{pool} from this set, the agent strictly prefers to implement action a^h, because the agent's individual rationality and incentive constraint become slack. But this contract must be a profitable deviation for the $\mathbb{Q}$-type, which concludes the proof.

Step 1: We first prove that the $\mathbb{Q}$-type's individual rationality and incentive constraint are both slack under the pooling contract w^{pool}. Note that necessarily a^h is implemented under w^{pool}, since otherwise the $\mathbb{P}$-type cannot improve on his upper bound. Thus, we want to show that the $\mathbb{Q}$-type's constraints satisfy $q^h \cdot u(w^{pool}) > \bar{U} + c$ and $(q^h - p^l) \cdot u(w^{pool}) > c$.

To see this, realize that (IR) and (IC) have to hold in expectation, that is

$$\begin{aligned} \lambda_{\mathbb{P}}\ [p^h \cdot u(w^{pool})]\ +\ [1-\lambda_{\mathbb{P}}]\ [q^h \cdot u(w^{pool})] &\geq \bar{U} + c, \\ \lambda_{\mathbb{P}}\ [(p^h - p^l) \cdot u(w^{pool})]\ +\ [1-\lambda_{\mathbb{P}}]\ [(q^h - p^l) \cdot u(w^{pool})] &\geq c. \end{aligned}$$

The $\mathbb{P}$-type must either violate his individual rationality or his incentive or even both constraints under w^{pool}—otherwise $w^{SB}_{\mathbb{P}}$ does not generate second-best profits, which are exceeded by the $\mathbb{P}$-type. Clearly, if he violates both constraints, the $\mathbb{Q}$-type's constraints have to be slack in order to satisfy the constraints in expectation. Thus, it remains to be shown that the $\mathbb{Q}$-type's constraints are even slack if the $\mathbb{P}$-type violates just a single constraint.

Assume that the $\mathbb{P}$-type violates (IR), but not (IC). Evidently, the (IR) of the $\mathbb{Q}$-type must be slack, but it also implies that $p^l \cdot u(w^{pool}) < \bar{U}$. Combining both conclusions shows that the $\mathbb{Q}$-type's (IC) is slack, too.

Now assume that the $\mathbb{P}$-type violates (IC), but not (IR). It follows immediately that the $\mathbb{Q}$-type's (IC) is slack. Using the (IC) of each type together yields

$$p^h \cdot u(w^{pool})\ <\ p^l \cdot u(w^{pool}) + c\ <\ q^h \cdot u(w^{pool}).$$

If $q^h \cdot u(w^{pool}) \leq \bar{U} + c$, the inequality chain above implies that $p^h \cdot u(w^{pool}) < \bar{U} + c$, violating the (IR) for the $\mathbb{P}$-type. This is a contradiction. Hence, it holds $q^h \cdot u(w^{pool}) > U + c$, i.e. the (IR) is slack for the $\mathbb{Q}$-type.

Step 2: We now construct contracts where the $\mathbb{Q}$-type is better and the $\mathbb{P}$-type worse off than in equilibrium if action a^h is chosen by the agent. Since $p^h \neq q^h$ one can find $i, j \in \{1, \ldots, n\}$ such that $p^h_i > q^h_i$ and $p^h_j < q^h_j$. Now take $\delta \in (1, \frac{p^h_i}{q^h_i})$, $\epsilon > 0$ and define the contract w^ϵ as follows:

$$\begin{aligned} w^\epsilon_i &:= w^{pool}_i + \frac{\delta\epsilon}{p^h_i}, \\ w^\epsilon_j &:= w^{pool}_j - \frac{\epsilon}{p^h_j}, \\ w^\epsilon_k &:= w^{pool}_k, \qquad \text{for all } k \neq i, j. \end{aligned}$$

By construction we have

$$
\begin{aligned}
p^h \cdot w^\epsilon &= p^h \cdot w^{pool} + \delta\epsilon - \epsilon \\
&> p^h \cdot w^{pool},
\end{aligned}
$$

and

$$
\begin{aligned}
q^h \cdot w^\epsilon &= q^h \cdot w^{pool} + \delta\epsilon \frac{q_i^h}{p_i^h} - \epsilon \frac{q_j^h}{p_j^h} \\
&< q^h \cdot w^{pool} + \epsilon - \epsilon \quad = \quad q^h \cdot w^{pool}.
\end{aligned}
$$

Hence, our contract w^ϵ fulfills the desired property. In addition $w^\epsilon \to w^{pool}$ as $\epsilon \to 0$, so that we can find such a contract arbitrarily close to w^{pool}.

Step 3: We now show that the equilibrium fails the Intuitive Criterion. Since by Step 1 both constraints are always slack for the $\mathbb{Q}$-type under w^{pool} and the agent's utility function is continuous, we can construct an open ball around w^{pool} such that the constraints remain slack within this ball. By choosing ϵ small enough, we can find an off-equilibrium contract $\hat{w}^\epsilon$ within this ball with the property described in Step 2. Observe that the $\mathbb{P}$-type is strictly better off with his equilibrium strategy than with offering $\hat{w}^\epsilon$ independent of the agent's best response. This is obvious if the agent finds it optimal to reject the contract and an immediate consequence of Step 2 if the agent responds with a^h. For the agent's choice of a^l note that the agent is compensated with a least $\bar{U}$ under $\hat{w}^\epsilon$—otherwise she would prefer rejecting the contract and a^l was no best response—and therefore the $\mathbb{P}$-type's profits are strictly below second-best and in particular below his equilibrium payoff. In contrast, the $\mathbb{Q}$-type can improve upon his equilibrium payoff due to the property of $\hat{w}^\epsilon$ if the agent finds it optimal to select a^h. Beause $\hat{w}^\epsilon$ lies in the open ball, a^h is indeed a best response for the agent under the belief to face the $\mathbb{Q}$-type with probability one. Thus, by the Intuitive Criterion, only the $\mathbb{P}$-type is discarded from offering $\hat{w}^\epsilon$ and the agent believes unambiguously to face the $\mathbb{Q}$-type when she observes this contract. Under this revised belief, however, action a^h is the unique best response to $\hat{w}^\epsilon$ for the agent. Thus, the $\mathbb{Q}$-type obtains with $\hat{w}^\epsilon$ strictly higher payoffs than in equilibrium under any choice from the revised set of best responses of the agent. Hence, the equilibrium fails the Intuitive Criterion. □

Theorem 4 is in line with the conclusion in Silvers (2012) showing that the principal's private information can never be beneficial for the principal. Even though his model covers a more general setup regarding actions, he restricts attention to distributions satisfying a monotone likelihood ratio property. While we do not impose this assumption, our result hinges on the simplification that the principal has no private information with respect to the low action. Thus, both approaches are complementary.

Note from Example 2 that the upper bounds in Theorem 4 are not always attained. Since our example does not employ the Intuitive Criterion, a principal type may, however, even obtain profits above this threshold. Observe also that the distributions in the example do not satisfy the monotone likelihood ratio property of Silvers (2012).[16] Hence, the example is not merely a special case of his model.

2.7 Discussion

We consider a discrete informed principal model with moral hazard in which the principal has private information regarding the stochastic relationship between effort and profits. We study distributional conditions that let the principal's information become irrelevant. Under a risk neutral agent linear independence of the set of distribution vectors over profits is sufficient for effecting the first-best payoff for the principal. Because a multitude of equilibria can be sustained in signaling models, we do not show linear independence to be also a necessary condition for informational irrelevance in the general setup, but derive this result in a special case. This full characterization allows us to tie the impact of the principal's private information to the number of available profit states. We illustrate this relationship by reproducing distortive effects of the principal's information that are encountered in the literature. Finally, an equilibrium existence result is presented.

We then turn to the case of a strictly risk averse agent. Because the agent not only cares about her expected payment, but also her exposure to risk, linear independence of the distribution vectors is no longer a sufficient condition to render the principal's informa-

[16] If the monotone likelihood ratio condition imposed in Silvers (2012) is applied to our example, the distributions p^h, q^h, p^l have to satisfy

$$\frac{p_i^h}{p_i^l} \geq \frac{p_j^h}{p_j^l} \qquad \text{and} \qquad \frac{q_i^h}{p_i^l} \geq \frac{q_j^h}{p_j^l}$$

for all $i < j$ with $i, j \in \{1, 2, 3\}$. This is violated for distributions q^h and p^l in our example.

tion inconsequential and we provide an illustrating example. Without any distributional assumptions, but an equilibrium refinement, we show that the principal's payoffs can be bounded above by the second-best profits from the standard principal-agent model. Thus, the principal's information can never be beneficial to him.

We want to discuss some of our results and provide further avenues for research. First, consider the case of a risk neutral agent. It seems difficult to extend Theorem 2 (full characterization of informational irrelevance) and Theorem 3 (equilibrium existence) to the general model, because our simplifying assumptions are active drivers of these results.

For instance, equilibrium existence is hard to establish whenever the principal has also private information about the low action distribution. In that case profitable deviations of the principal under the low action have also to be taken into account. This can be a source of major complications: For instance, fix an equilibrium outcome and imagine an off-equilibrium contract that provides a profitable deviation for the $\mathbb{P}$-type if the agent implements a^h, and the agent indeed does so, when she believes to face the $\mathbb{Q}$-type. Assume the same for the low action a^l with the role of the principal types reversed. It is not at all clear that an appropriate belief can be assigned to this contract such that an equilibrium can be sustained. Further distributional assumptions probably need to be imposed to rule out the existence of such constellations.

On the other hand, the proof of the characterization theorem greatly simplifies with a^h being the first-best action for both types. For example, consider the case when both types agree on a^l as first-best. A pooling contract $\hat{w}$ that provides both types with their first-best profits could satisfy $(p^h, q^h, p^l) \cdot \hat{w} = (\bar{U}+c-\epsilon_1, \bar{U}+c-\epsilon_2, \bar{U})$ where $(\epsilon_1, \epsilon_2) \in \mathbb{R}^2_+$. But this equilibrium condition alone does not rule out linear dependence of the distribution vectors by Lemma 1; to obtain further structure for the distribution vectors, one has to incorporate the equilibrium condition that neither principal type has an incentive to deviate to an off-equilibrium contract. This, in turn, is hard, because the choice of off-equilibrium path beliefs is arbitrary and allows for a lot of leeway in sustaining equilibrium strategies. It suggests, however, that equilibrium refinements will help to further progress in this direction; we have not yet taken this approach.

In the case of risk aversion, the analysis is even more intricate, because of the agent's preferences over risk and technical challenges remain. In particular, providing sufficient conditions for an information irrelevance result as in the proof of Theorem 1 is difficult.

They have to guarantee that a nonlinear equation system is solvable and a solution to this system lies on a hyperplane defined by a principal type's second-best profits and the distribution vector of his second-best action. We leave this issue open for further research.

2.8 Appendix

PROOF OF LEMMA 1: The if part is obvious. For the only-if part note, first, that we can find a $k \in \{1, \ldots, n\}$ such that $p_k > r_k$ because $p \neq r$. Then, possibly after relabeling our vectors, it holds that $p_k > r_k$ and $q_k \in [r_k, p_k]$. Second, since $p \neq r$ and $\sum p_i = \sum r_i = 1$ the vectors p and r are linearly independent. Thus, if p, q, r are linearly dependent, the vector q can be written as $q = \alpha p + \beta r$, which implies

$$1 \;=\; \sum_{i=1}^{n} q_i \;=\; \alpha \sum_{i=1}^{n} p_i + \beta \sum_{i=1}^{n} r_i \;=\; \alpha + \beta,$$

or $\beta = 1 - \alpha$. Hence, the vector q is a convex combination of p and r if and only if $\alpha \in (0, 1)$.

Obviously, $0 \neq \alpha \neq 1$, because p, q, r are mutually distinct. So let $\alpha < 0$ and $\beta > 1$. Then

$$q_k \;=\; \alpha p_k + \beta r_k \;<\; \alpha r_k + \beta r_k \;=\; r_k \;\leq\; q_k.$$

This yields a contradiction. So assume $\alpha > 1$ and $\beta < 0$. This implies

$$q_k \;=\; \alpha p_k + \beta r_k \;>\; \alpha p_k + \beta p_k \;=\; p_k \;\geq\; q_k.$$

Again, a contradiction. Hence, $\alpha \in (0, 1)$.

For uniqueness assume there exists a $\gamma \in (0, 1)$ with $\gamma \neq \alpha$ and $q = \gamma p + (1 - \gamma) r$. But then

$$\begin{aligned} q_k &= \gamma p_k + (1-\gamma) r_k \\ &= (\alpha + \gamma - \alpha) p_k + (1 - \alpha - [\gamma - \alpha]) r_k \\ &= \underbrace{\alpha p_k + (1-\alpha) r_k}_{= q_k} + \underbrace{[\gamma - \alpha](p_k - r_k)}_{\neq 0} \;\neq\; q_k. \end{aligned}$$

Contradiction, so $\gamma = \alpha$. □

REMAINING PROOF OF THEOREM 3: We show that the Perfect Bayesian Equilibria constructed in the proof satisfy the Intuitive Criterion. Recall that the Intuitive Criterion refines equilibria by restricting beliefs associated with off-equilibrium path contract offers. The underlying idea is that the agent rationalizes which principal type could have offered a certain off-equilibrium contract. In particular, the agent should expect offers only from types who can improve on their equilibrium outcome. To make things more precise, fix an equilibrium outcome and consider a specific off-equilibrium contract. The Intuitive Criterion proceeds in two steps: First, it discards any principal type whose equilibrium payoff is strictly above the gain under the off-equilibrium contract independent of the agent's choice from the set of best responses. Note that the agent's best response set contains all actions that are optimal for some distribution over the principal types. Second, it checks if among the remaining types there exists one whose equilibrium payoff is strictly below the gain under the contract independent of the agent's choice from a restricted set of best responses. This restricted set now only contains optimal actions for distributions over the remaining types. Thereafter, the procedure stops. If there exists an off-equilibrium contract with a principal type surviving the first and satisfying the second step, we say that the equilibrium fails the Intuitive Criterion. If for any off-equilibrium contract no such type exists, we say that the equilibrium satisfies the Intuitive Criterion.

Now consider off-equilibrium contracts of our equilibrium construction. We begin with contracts $w \in \mathbb{R}^n \backslash \left(W_{\mathbb{P}}^{dev} \cup W_{\mathbb{Q}}^{dev} \cup \{w_{\mathbb{P}}, w_{\mathbb{Q}}\} \right)$. Note that neither principal type has a profitable deviation to these contracts under any possible belief, because the choice of these beliefs was arbitrary in the equilibrium construction. Thus, neither type has an incentive to deviate to any of these contracts under a restricted belief structure. This is a sufficient condition that the equilibria do not fail the Intuitive Criterion for these contracts. Now focus on off-equilibrium contracts $w \in W_{\mathbb{P}}^{dev}$ and $w \in W_{\mathbb{Q}}^{dev}$ with

$$\begin{aligned} \mu(\mathbb{P}|w) &= 1 \qquad \text{for all } w \in W_{\mathbb{P}}^{dev}, \\ \mu(\mathbb{Q}|w) &= 1 \qquad \text{for all } w \in W_{\mathbb{Q}}^{dev}. \end{aligned}$$

We only examine contracts $w \in W_{\mathbb{P}}^{dev}$—the case of $w \in W_{\mathbb{Q}}^{dev}$ is entirely analogous. Note that by construction $w \in W_{\mathbb{P}}^{dev}$ strictly improves upon the $\mathbb{P}$-type's equilibrium payoff if the agent responds with a^h, whereas it does not for the $\mathbb{Q}$-type. Moreover, both types obtain the same payoff if the contract is responded by a^l or rejected. Thus, whenever the $\mathbb{P}$-type is discarded by the first step of the Intuitive Criterion, this is also true for the $\mathbb{Q}$-type. So for any first round type elimination, the agent can justify to face the $\mathbb{P}$-type with probability one. This implies that a^l or rejecting the contract belongs to the restricted set of best responses of the agent. Since both principal types do not improve upon their equilibrium payoff if the agent chooses a^l or rejects, the second step of the Intuitive Criterion doesn't apply to either type.[17] Hence, our equilibria do not fail the Intuitive Criterion for contracts $w \in W_{\mathbb{P}}^{dev}$ (and $w \in W_{\mathbb{Q}}^{dev}$) either. Therefore our equilibria satisfy the Intuitive Criterion. $\square$

Proof of Proposition 1: Let $n \geq 3$ and consider the following notation:

$$\begin{aligned} X &:= \{ (p,q,r) \in \mathbb{R}^{n\times 3} \mid p,q,r \text{ probability distributions } \}, \\ A &:= \{ (p,q,r) \in X \mid p,q,r \text{ linearly independent } \}, \\ A^c &:= \{ (p,q,r) \in X \mid p,q,r \text{ linearly dependent } \}, \\ A_{\text{limit}} &:= \{ (p,q,r) \in X \mid (p,q,r) \text{ limit point of } A \}, \\ \text{cl } A &:= A \cup A_{\text{limit}}. \end{aligned}$$

We show A to be dense and open in X in two separate steps.

Claim 1. The set A is dense in X, i.e. cl $A = X$.

Proof: Obviously, it holds that $A \cup A^c = X$ and cl $A \subset X$. Hence, we just have to show $A^c \subset A_{\text{limit}}$, because this implies $X = A \cup A^c \subset A \cup A_{\text{limit}} = \text{cl } A$. So let $(p,q,r) \in A^c$. We establish the result for three different cases.

Case 1: (p,q,r) are mutually distinct. By Lemma 1 (and possibly after relabeling) we can write $q = \alpha p + (1-\alpha) r$ for some $\alpha \in (0,1)$. Since $p \neq r$ there exist $k,l \in \{1,\dots,n\}$ with $p_k > r_k$ and $p_l < r_l$. Define the following sequence $(p^m,q^m,r^m)_{m\in\mathbb{N}}$ with $p^m = p,\ r^m = r$ and

$$\begin{aligned} q_k^m &= q_k - \frac{\epsilon}{m}, \\ q_l^m &= q_l, \\ q_i^m &= q_i + \frac{\epsilon}{m(n-2)}, \quad \text{for all } i \in \{1,\dots,n\}\backslash\{k,l\} \end{aligned}$$

for all $m \geq 1$ and $\epsilon \in (0, q_k - r_k)$. By construction, p^m, q^m, r^m are probability distributions and $p_k^m > q_k^m > r_k^m$ for all $m \geq 1$. Furthermore $\lim_{m\to\infty}(p^m,q^m,r^m) = (p,q,r)$. It remains to be shown that p^m,q^m,r^m are linearly independent to establish $(p,q,r) \in A_{\text{limit}}$.

[17] If a^l belongs to the restricted set of best responses, the agent has to obtain at least his reservation utility $\bar{U}$ under contract w. This guarantees that neither principal type gets more than in equilibrium under a^l.

So assume by contradiction that p^m, q^m, r^m are linearly dependent. Since p^m, q^m, r^m are mutually distinct (consider the k^{th} entry), by Lemma 1, one vector is a convex combination of the other two. As $p_k^m > q_k^m > r_k^m$ this vector must be q, that is

$$q^m = \alpha^m p^m + (1-\alpha^m) r^m = \alpha^m p + (1-\alpha^m) r$$

for some $\alpha^m \in (0,1)$. But note that from $q_k^m < q_k$ and $p_k > r_k$ we have $\alpha^m < \alpha$ for all $m \in \mathbb{N}$. On the other hand, because $p_l < r_l$

$$q_l^m = \alpha^m p_l + (1-\alpha^m) r_l > \alpha p_l + (1-\alpha) r_l = q_l$$

contradicting $q_l^m = q_l$. Hence, p^m, q^m, r^m are linearly independent.

Case 2: exactly two vectors of p, q, r *are equal.* Without loss of generality assume $p = q \neq r$. As before there exist $k, l \in \{1, \dots, n\}$ with $p_k > r_k$ and $p_l < r_l$. Define a sequence $(p^m, q^m, r^m)_{m \in \mathbb{N}}$ with $p^m = p, r^m = r$ and

$$\begin{aligned} q_k^m &= p_k - \frac{\epsilon}{m}, \\ q_l^m &= q_l \ (= p_l), \\ q_i^m &= q_i + \frac{\epsilon}{m(n-2)}, \quad \text{for all } i \in \{1, \dots, n\} \backslash \{k, l\} \end{aligned}$$

for all $m \geq 1$ and $\epsilon \in (0, p_k - r_k)$. Again, p^m, q^m, r^m are probability distributions, $p_k^m > q_k^m > r_k^m$ for all $m \geq 1$ and $\lim_{m \to \infty}(p^m, q^m, r^m) = (p, q, r)$. Hence, $(p, q, r) \in A_{\text{limit}}$ if p^m, q^m, r^m are linearly independent for all $m \in \mathbb{N}$.

Assume by contradiction that p^m, q^m, r^m are linearly dependent. Because our vectors are mutually distinct (consider the k^{th} entry), by Lemma 1

$$q^m = \alpha^m p^m + (1-\alpha^m) r^m = \alpha^m p + (1-\alpha^m) r$$

for some $\alpha^m \in (0,1)$. But this implies

$$q_l^m = \alpha^m p_l + (1-\alpha^m) r_l > p_l = q_l,$$

violating $q_l^m = q_l$. Thus, p^m, q^m, r^m are linearly independent.

Case 3: $p = q = r$. There exists a $k \in \{1, \ldots, n\}$ such that $0 < p_k \leq 1$. Define $p^m = p$ and

$$\begin{aligned} q_k^m &= p_k - \frac{\epsilon}{2m}, \\ q_l^m &= p_l + \frac{\epsilon}{2m}, \\ q_i^m &= p_i, \qquad \text{for all } i \in \{1, \ldots, n\} \backslash \{k, l\}, \end{aligned}$$

and

$$\begin{aligned} r_k^m &= p_k - \frac{\epsilon}{m}, \\ r_j^m &= p_j + \frac{\epsilon}{m}, \\ r_i^m &= p_i, \qquad \text{for all } i \in \{1, \ldots, n\} \backslash \{k, j\}, \end{aligned}$$

where k, l, j are mutually distinct (which is feasible if $n \geq 3$) and $\epsilon \in (0, p_k)$. Again, p^m, q^m, r^m are probability distributions and $\lim_{m \to \infty}(p^m, q^m, r^m) = (p, q, r)$. Hence, $(p, q, r) \in A_{\text{limit}}$, if p^m, q^m, r^m are linearly independent for all $m \in \mathbb{N}$.

Assume that p^m, q^m, r^m are linearly dependent. Since $p_k^m > q_k^m > r_k^m$ our vectors are mutually distinct and we can make use of Lemma 1. It follows that

$$q^m = \alpha^m p^m + (1 - \alpha^m) r^m$$

for some $\alpha^m \in (0, 1)$. This gives us

$$q_l^m = \alpha^m p_l^m + (1 - \alpha^m) r_l^m = p_l.$$

But by construction $q_l^m = p_l + \epsilon/2m > p_l$, contradiction. Thus, p^m, q^m, r^m are linearly independent.

Cases 1-3 yield $A^c \subset A_{\text{limit}}$. □

Claim 2. The set A is open in X.

PROOF: Assume A is not open in X, i.e. there exists $(p, q, r) \in A$ such that any open ball centered at (p, q, r) contains a triple of distributions that is linearly dependent. Thus, we can construct a sequence $(p^m, q^m, r^m)_{m \in \mathbb{N}}$ with $\lim_{m \to \infty} (p^m, q^m, r^m) = (p, q, r)$ and

(p^m, q^m, r^m) linearly dependent for all $m \in \mathbb{N}$. It follows that for any $m \in \mathbb{N}$ there exists $\tilde{\lambda}^m \in \mathbb{R}^3 \backslash \{0\}$ such that

$$\tilde{\lambda}_1^m p^m + \tilde{\lambda}_2^m q^m + \tilde{\lambda}_3^m r^m = 0.$$

Define $\lambda_i^m := \tilde{\lambda}_i^m / ||\tilde{\lambda}^m||$ for all $i \in \{1, \ldots, n\}$ and note that $||\lambda^m|| = 1$ for all $m \in \mathbb{N}$, that is, the sequence $(\lambda^m)_{m \in \mathbb{N}}$ is bounded. Thus, by the Bolzano-Weierstrass Theorem there exists a subsequence $(\lambda^{m_k})_{k \in \mathbb{N}}$ converging to some $\lambda \in \mathbb{R}^3$. Since the norm is continuous $||\lambda|| = \lim_{k \to \infty} ||\lambda^{m_k}|| = 1$, implying that $\lambda \neq 0$. But then

$$\begin{aligned} \lambda_1 p + \lambda_2 q + \lambda_3 r &= \lim_{k \to \infty} (\lambda_1^{m_k} p^{m_k} + \lambda_2^{m_k} q^{m_k} + \lambda_3^{m_k} r^{m_k}) \\ &= \lim_{k \to \infty} 1/||\tilde{\lambda}^{m_k}|| \left(\tilde{\lambda}_1^{m_k} p^{m_k} + \tilde{\lambda}_2^{m_k} q^{m_k} + \tilde{\lambda}_3^{m_k} r^{m_k} \right) \\ &= 0 \end{aligned}$$

that is, the triple (p, q, r) is linearly dependent. But this violates $(p, q, r) \in A$. Thus, A is open in X. □

Claim 1 and 2 establish Proposition 1. □

Proof of Proposition 2*:* We develop the result in several steps.

Step 1: In any Perfect Bayesian Equilibrium both types of principals implement action a^h and the $\mathbb{Q}$-type pays at most $\bar{U} + c$ to the agent. Assume one type does not implement a^h. Then this type gets an equilibrium payoff of at most $p^l \cdot \pi - \bar{U}$. Since a^h is first-best, we have $p^h \cdot \pi - [\bar{U} + c] > q^h \cdot \pi - [\bar{U} + c] > p^l \cdot \pi - \bar{U}$. Let $\epsilon := q^h \cdot \pi - [\bar{U} + c] - [p^l \cdot \pi - \bar{U}] > 0$. Define $\hat{w}_i := w_{\mathbb{Q},i}^{stf} + \frac{\epsilon}{2}$ for all $i \in \{1, \ldots, n\}$. If the principal type offers this contract, the agent accepts and implements action a^h independent of her belief, because under the $\mathbb{P}$-type's distribution even higher gross profits are generated than under the $\mathbb{Q}$-type's one. But our type then obtains profits of $q^h \cdot \pi - [\bar{U} + c] - \frac{\epsilon}{2} > p^l \cdot \pi - \bar{U}$ and has a profitable deviation. This also shows that the $\mathbb{Q}$-type cannot fall short of his first-best profits and therefore pays at most $\bar{U} + c$ to the agent.

Step 2: In any pooling ($w_{\mathbb{P}} = w_{\mathbb{Q}}$) or separating equilibrium ($w_{\mathbb{P}} \neq w_{\mathbb{Q}}$) it holds $w_{\mathbb{F},1} > w_{\mathbb{F},2}$ and $p^h \cdot w_{\mathbb{F}} > q^h \cdot w_{\mathbb{F}}$ for all $\mathbb{F} \in \{\mathbb{P}, \mathbb{Q}\}$. The first inequality follows, because a^h is implemented in equilibrium (Step 1), and the second is due to $p_1^h > q_1^h$.

Step 3: No pooling equilibrium exists satisfying the Intuitive Criterion. Assume there exists a pooling equilibrium with contract w^{pool}, which, by Step 1, implements action a^h.

(i) *(IR) and (IC) of the $\mathbb{P}$-type are both slack.* By Step 1, the $\mathbb{Q}$-type never pays more than $\bar{U} + c$ to the agent. Thus, to meet (IR) in expectation the $\mathbb{P}$-type has to pay more than that if the $\mathbb{Q}$-type violates (IR) individually. But he also pays more if the $\mathbb{Q}$-type just satisfies it, because of Step 2. Thus, the $\mathbb{P}$-type's (IR) is slack. Regarding (IC), assume it is not slack for the $\mathbb{P}$-type, i.e. $(p^h - p^l) \cdot w^{pool} \leq c$. But then, by Step 2, $(q^h - p^l) \cdot w^{pool} < c$ and (IC) cannot be satisfied in expectation. Thus, the $\mathbb{P}$-type's (IC) is also slack.

(ii) *There exist contracts arbitrarily close to w^{pool} for which the $\mathbb{P}$-type obtains strictly higher payoffs than in equilibrium and the $\mathbb{Q}$-type strictly lower ones if the agent chooses a^h.* Consider contract w^ϵ with

$$w_1^\epsilon := w_1^{pool} - \frac{\delta\epsilon}{p_1^h}, \qquad w_2^\epsilon := w_2^{pool} + \frac{\epsilon}{p_2^h},$$

and $\epsilon > 0$ and $\delta \in (1, p_1^h/q_1^h)$. Then

$$\begin{aligned} p^h \cdot w^\epsilon &= p^h \cdot w^{pool} - \delta\epsilon + \epsilon \\ &< p^h \cdot w^{pool} \end{aligned}$$

and

$$\begin{aligned} q^h \cdot w^\epsilon &= q^h \cdot w^{pool} - \epsilon\, \overbrace{\delta \frac{q_1^h}{p_1^h}}^{<1} + \epsilon\, \overbrace{\frac{q_2^h}{p_2^h}}^{>1} \\ &> q^h \cdot w^{pool}. \end{aligned}$$

Hence, if the agent chooses a^h, contract w^ϵ gives the $\mathbb{P}$-type a strictly higher payment than in equilibrium and a strictly lower one to the $\mathbb{Q}$-type. Moreover, $w^\epsilon \to w^{pool}$ as $\epsilon \to 0$.

(iii) *No pooling equilibrium exists.* Since by (i) both constraints are always slack for the $\mathbb{P}$-type under w^{pool} and the agent's utility function is continuous, we can construct an open ball around w^{pool}, such that the constraints remain slack within this ball. By choosing ϵ small enough, we can find an off-equilibrium contract $\hat{w}^\epsilon$ within this ball with the feature described in (ii). Observe that the $\mathbb{Q}$-type is strictly better off with his equilibrium strategy than with offering $\hat{w}^\epsilon$ independent of the agent's best response. This is obvious if the agent finds it optimal to reject the contract and a direct consequence of (ii) if the agent responds with a^h. For the agent's choice of a^l note that the agent is compensated with a least $\bar{U}$ under $\hat{w}^\epsilon$—otherwise she would prefer rejecting the contract and a^l was no best response—and therefore the $\mathbb{Q}$-type's profits are strictly below first-best and in particular below his equilibrium payoff (Step 1). In contrast, the $\mathbb{P}$-type can improve upon his equilibrium payoff due to the property of $\hat{w}^\epsilon$ if the agent finds it optimal to select a^h. Beause $\hat{w}^\epsilon$ lies in the open ball, a^h is indeed a best response for the agent under the belief to face the $\mathbb{P}$-type with probability one. Thus, by the Intuitive Criterion, only the $\mathbb{Q}$-type is discarded from offering $\hat{w}^\epsilon$ and the agent believes unambiguously to face the $\mathbb{P}$-type when she observes this contract. Under this revised belief, however, action a^h is the unique best response to $\hat{w}^\epsilon$ for the agent. Thus, the $\mathbb{P}$-type obtains with $\hat{w}^\epsilon$ strictly higher payoffs than in equilibrium under any choice from the revised set of best responses of the agent. Hence, the equilibrium fails the Intuitive Criterion.

Step 4: $q^h \cdot w_{\mathbb{Q}} = \bar{U} + c$ *and* $p^h \cdot w_{\mathbb{P}} > \bar{U} + c$. Since in a separating equilibrium the principal's type is revealed, the $\mathbb{Q}$-type pays at least $\bar{U} + c$ to the agent. But, by Step 1, he also does not pay more than that, showing the first equality. Moreover,

$$p^h \cdot w_{\mathbb{P}} \;>\; q^h \cdot w_{\mathbb{P}} \;\geq\; q^h \cdot w_{\mathbb{Q}} \;=\; \bar{U} + c,$$

where the first inequality is a consequence from Step 2 and the second follows, because in equilibrium no type has an incentive to deviate to the contract of the other type.

Step 5: $(q^h - p^l) \cdot w_{\mathbb{Q}} \geq c$ *and* $(p^h - p^l) \cdot w_{\mathbb{P}} = c$. The first equality simply results from the $\mathbb{Q}$-type's (IC) and type revelation in separating equilibria. Regarding the $\mathbb{P}$-type's (IC), note that a slack (IC) had both his constraints slack by Step 4. In such a case, with the help of the Intuitive Criterion, one can find a profitable deviation for the $\mathbb{P}$-type close to $w_{\mathbb{P}}$. The procedure is completely analogous to Step 3, but also uses the equilibrium condition that the $\mathbb{Q}$-type's payoff under $w_{\mathbb{Q}}$ is weakly higher than under $w_{\mathbb{P}}$. Thus, the $\mathbb{P}$-type's (IC) necessarily holds with equality.

Step 6: $q^h \cdot w_{\mathbb{P}} = \bar{U} + c$. We know already from Step 4 that $q^h \cdot w_{\mathbb{P}} \geq \bar{U} + c$. So assume $q^h \cdot w_{\mathbb{P}} > \bar{U} + c$. We show that the $\mathbb{P}$-type has a profitable deviation. Consider contracts that are identical to $w_{\mathbb{P}}$, but pay less to the agent in the low profit state, i.e. state 2. For a sufficiently small decrease, we can find a contract $\hat{w}$ still satisfying $q^h \cdot \hat{w} > \bar{U} + c$ and $p^h \cdot \hat{w} > \bar{U} + c$, because the $\mathbb{P}$-type's (IR) is also slack at $w_{\mathbb{P}}$ (Step 4). Note, by Step 5 and $p_2^h < p_2^l$, that the $\mathbb{P}$-type's (IC) is slack at $\hat{w}$, too. By making use of the Intuitive Criterion one can show entirely analogous to Step 3 (iii) that $\hat{w}$ is a profitable deviation for the $\mathbb{P}$-type. Thus, it must hold $q^h \cdot w_{\mathbb{P}} = \bar{U} + c$.

Steps 1-6 establish Proposition 2. □

Chapter 3

Shopping Clubs, Clearance Sales, and Distribution Channel Choice

In the recent past, shopping clubs have established themselves as popular online distribution channels for clearance sales. We study their emergence in a durable-goods model in which they have a cost advantage compared to traditional store outlets. Shopping clubs are therefore prefered distribution channels for producers at the end of the season. We show, however, that under relatively small cost differences, producers favor traditional shops in early sales periods, because the club fails to internalize how prices in early periods influence future producer profits. This explains the observation that shopping clubs are used for clearance, but not for regular sales.

3.1 Introduction

Shopping clubs are online distribution channels that specialize in clearance sales for brand producers. Their product categories encompass mainly clothing, shoes, and accessories, but some have also focused on high-end electronics or young-family goods. Products are offered with substantial discounts to club members in sales deals, which last just a few days and are announced to customers on a weekly basis. While the first clubs were founded quite recently in 2001, the industry has already grown to considerable size. *Vente priveé*, for instance, a French shopping club active in Western Europe, had almost 1500 employees and attained revenues just above one billion euros in 2011; in comparison, the German online market for apparel and shoes amounted to roughly three billion euros at that time.

Despite their widespread success, the business model of shopping clubs seems to be less transparent. Darius Zamani-Achtiani, the Head of Sourcing of *Brands4Friends*, a German club, says that shopping clubs assist their brand partners who have "issues with remaining overstock and end-of-season products." Producers are very intent on not damaging the image of the brand and avoid clearance sales on rummage tables. Melanie Bauer, a sales manager at *Michalsky*, a German fashion label, misses an "appropriate distribution channel to sell our overstock production without depreciation of our brand image." Shopping clubs advertise themselves for closing this gap by selling their goods to a selected clientele in an exclusive environment. To this end different measures are taken: First, access to membership is restricted and entry criteria vary from club to club.[1] Second, online platforms of clubs are protected from search engine activity. Thus, it is impossible for nonmembers to obtain information about the club's current offers and prices in this way. And third, shopping clubs claim to accept only well-known brands for their assortment. This does not only dispel customers' doubts about quality, but also reassures producers that their goods are sold in an appropriate environment. Thus, as Stephan Zoll, CEO of *Brands4Friends*, puts it: "The advantages for the manufacturers of partnering with *Brands4Friends* are obvious. [...] The concept of an exclusive shopping club with a high fashion and brand competence supports the brand image of the manufacturer and sometimes even helps to evolve it."

This self-attributed exclusiveness, however, does not match with the business conduct of shopping clubs. In fact, the most successful clubs have already such a large and still growing member base that it is hard to speak of a selected community.[2] This condition mainly owes to a very loose entry policy to membership.[3] The lack of exclusiveness also extends to the clubs' offers and prices: Although search engines are barred from this data, web pages have emerged that collect weekly offers from all clubs in one place and make them visible to the public. And finally, many clubs have succumbed to the temptation of

[1] Some clubs allow existing members to send a limited number of invitations to their friends; others require to hand in an online application and grant access at their own discretion.

[2] For instance, in 2011, *Brands4Friends* had 4 million members in Germany with a monthly inflow of 150.000 new customers.

[3] Many clubs try to boost membership rates by giving vouchers to existing customers if they convince others to join the club. In addition, old registration links from past promotion campaigns are often still active and can easily be spotted with a web search. Finally, given the large user base, access to membership is often a formality, because one can usually find a club member in his social circle and ask for an invitation.

diluting their portfolio of high-end producers—some platforms distribute the goods of far more than 1000 brands, many of which are hardly known to consumers.

Thus, there seems to be little evidence that supports the alleged business idea of shopping clubs. And yet, they have been thriving over the last years and are still growing impressively.[4] In this paper we present a new rationale that coherently explains the use of shopping clubs for clearance, but not for regular sales. We argue that the club has a relative cost advantage compared to traditional store outlets. Hence, the crucial question in our analysis is: Why do producers use shopping clubs just for end-of-season sales and not also in earlier periods?

In Section 3.2 we discuss related literature before we introduce our model in Section 3.3. The firm, a monopolistic brand manufacturer, wants to sell a durable good over two periods. She has two distribution channels available—her own store and the shopping club.[5] Both channels reach the same group of strategic consumers, which have preferences over prices, but not over the two distribution channels.[6] Distribution through her own store allows the firm to set the consumer price of the period, but is more costly than for the club: shop space has to be rented nationwide and employees have to be hired to operate these stores. In contrast, there are low maintenance costs for the club's web page and it benefits from cheap shipping fares due to its high turnover. When the firm decides to sell the good via the club, she can, however, only specify the capacity of goods available to the club in this period and the transfer price for this capacity. Even though the firm has all the bargaining power, the club chooses independently its optimal consumer price of the period. There is also a commitment problem for the firm: each period she has to choose the distribution channel and the period consumer price or the contract with the club anew. Finally, we assume demand to be stochastic at the beginning of the game and

[4] Just recently, two global e-commerce firms secured their access to the shopping club industry in Europe: In 2010, *Ebay* bought *Brands4Friends* for estimated 150 million euros, whereas *Amazon* acquired *BuyVIP*, a Spanish club, for purported 70 million euro.

[5] Of course, brand producers have often more options for distributing their goods, e.g. large department stores. We abstract from this issue in the remainder of this paper. One justification could be that these channels serve entirely different customer groups—for instance, if brand and department store are located at opposite positions on a Salop circle and transportation costs are relatively high. In that case, both stores serve disjoint intervals of consumers and do not influence the other's pricing decision.

[6] Even though it seems that an online distributor can reach a larger customer base, we assume that the firm's store has the same access to consumers. This assumption is particularly plausible for specific goods, e.g. high-end fashion brands, where potential customers and stores are mainly located in urban areas.

the state of nature is revealed between the firm's decision of the distribution channel and the price announcement in the first period.[7]

Section 3.4 contains all results of the paper. We solve the game by backward induction and start in period 2. Because the club is more cost efficient and the firm can accrue all of its profits due to her bargaining power, she prefers to distribute the good through the club (Lemma 1). Turning to the first period, an optimal selling scheme for the firm, i.e. selling under vertical integration of firm and club, combines two components: Obviously, it is profit maximizing to use the more cost efficient technology of the club in the first period. On the other hand, because consumer behavior and second-period profits are influenced by the first-period consumer price, optimal pricing strikes a balance between profits in both periods. Each distribution channel, however, can satisfy only one of these two aspects.

Since the firm has direct control over shop prices, the shop's first-period price correctly internalizes the intertemporal tradeoff across period profits. Nevertheless, inefficiently high distribution costs lead to suboptimal profits if the shop is employed in the first period (Lemma 2). This picture is reversed for the club: While distribution proceeds at low cost in the first period, the consumer-price choice of the club does not maximize total firm profits. Intuitively, second-period profits are increasing in the first-period price, because customers respond by postponing consumption to the second period. This implies that optimal pricing involves a price above the one that maximizes first-period profits. But since the club anticipates that all its second-period profits will be absorbed, its first-period price just maximizes first-period profits. Therefore, if the club is not capacity constrained in its pricing decision, it chooses a first-period consumer price that is too low from the firm's perspective. The firm can ameliorate this problem by reducing the club's first-period capacity and thus enforce a higher price. Because capacity is determined before the demand state is realized, this constraint leads to a distorted price in at least one demand state. Thus, even though the firm manages to extract all profits from the club in the first period (Lemma 3), the profits remain suboptimal if the club is used (Lemma 4).

[7]The underlying idea is that the shop and the club can quickly experiment with consumer prices to learn the demand state before they fix this price until the end of the first period. If this time of experimentation is relatively small compared to the remainder of the period, this assumption is well-justified. We further discuss this issue in Section 3.5.

Therefore, in period 1, both distribution channels only partially incorporate the optimal selling policy.[8] As the distribution cost of the shop decreases, the store's advantage of optimal consumer pricing outweighs the benefits from low distribution costs under the club. Hence, if the distribution cost of the shop is sufficiently close to the club's, the firm always prefers selling through her own shop in the first period (Theorem 1). This result is in line with the observation that firms use shopping clubs for clearance sales, but not for regular selling in earlier periods.

Section 3.5 contains a discussion of the model; all proofs are relegated to the appendix in Section 3.6.

3.2 Related Literature

Highly related to our paper is Karp and Perloff (1996). In their model a durable-good monopolist has two cost technologies available: an inferior (high increasing marginal cost) and a superior one (low increasing marginal cost). The monopolist lacks commitment power and each period he reconsiders his technology choice. Karp and Perloff show that under any strictly positive technology switching cost the monopolist uses the inferior technology at any point in time if the slope of the superior marginal cost function is sufficiently low. The intuition for this result builds on an insight from Kahn (1986): A durable-good monopolist with strictly increasing marginal cost can secure strictly positive profits in comparison to a monopolist with constant marginal cost.[9] Returning to the Karp-Perloff model, this implies that the monopolist would like to commit to the inferior technology, because profits under the superior technology converge to zero as the marginal cost function becomes more and more horizontal. But this commitment is also credible; although some consumers can still be profitably served with the superior technology after market saturation, the monopolist resists the temptation to use it—and destroy his promise to

[8] One may object that the firm could open her own online platform to benefit from small distribution costs. We return to this issue in Section 3.5.

[9] A durable-good monopolist with strictly increasing marginal cost is reluctant to flood the market with his goods too extensively after having served consumers in earlier periods. Because he can credibly commit to keep prices high to some extent, consumers are less eager to wait for lower offers and the monopolist obtains strictly positive profits. Thus, he can partially evade the famous Coase Conjecture: a durable-good monopolist competes so intensely with his future self such that he winds up selling the efficient quantity on the market right away and obtains zero profits. In fact, Stokey (1981) formally proves this result under constant marginal cost.

keep prices relatively high over time—since there are strictly positive technology switching costs.[10] While our model is also concerned with the question why a superior technology (the more cost efficient club) is not used, the respective drawbacks of usage are quite different. In their model, the cost benefit of the superior technology backfires by giving the monopolist a lower commitment value to keep prices high in future periods. In contrast, the monopolistic firm in our model refrains to use the club in the first period, because it fails to internalize the firm's second-period profits and does not keep the price high in the current period.

The distribution channel conflict between a firm and its retailer has received broad attention in the marketing literature. A particularly relevant contribution with regard to our work is Desai et al. (2004). They employ a two-period model in which a durable-good monopoly uses each period a two-part tariff to sell its good to a retailer.[11] The retailer has no additional skills compared to the upstream firm that would make him as an intermediary especially attractive for the firm. They find that the optimal two-part tariff attains the same profits as a vertically integrated distribution channel. This result holds true whether the monopoly has commitment power or not. Intuitively, a convenient wholesale-price choice induces the retailer to reproduce the pricing pattern under vertical integration. In addition, the fixed fee guarantees a full profit extraction from the retailer. Desai et al. informally argue that establishing a commitment with a single retailer is often easier than with numerous individual consumers. Hence, they provide a justification for intermediaries in a durable-good setting, even though there are no direct retailer benefits, like privileged access to consumers etc. In comparison to our model, however, they do not try to explain why a retailer should be prefered to direct sale in some periods but not in others.

Lehmann and Weinberg (2000) study a monopolist who has two versions of a good, each of which can be distributed through a specific distribution channel. This is, for instance, the case for movies in the theaters and home videos. The channels are interrelated with each other so that opening the second channel while the first channel is still active leads to cannibalization effects. Moreover, attainable profits from each distribution channel—

[10] Kutsoati and Zábojník (2005) also show that it can be optimal to suppress a superior technology. Their superior technology, however, raises consumers' willingness-to-pay after an initial period without effect.

[11] For distribution channel choice in a nondurable-good context, McGuire and Staelin (1983) and Chiang et al. (2003) are very instructive contributions.

open or not—shrink over time as consumers become less and less interested in the good. Lehmann and Weinberg derive the optimal instant when to switch from one channel to the other. A common feature with our model is the optimal dynamic channel choice, but there are at least two diverging aspects: First, the monopolist can commit to his timing choice at the beginning of the game and, second, consumers do not behave strategically.[12]

More remotely, there is a part of the literature that scrutinizes clearance sales in the context of durable goods. They all differ, however, from our model by excluding multiple distribution channels. In Conlisk et al. (1984) new consumers—consisting of high and low willingness-to-pay types—enter each period. They find that a monopolistic seller who cannot commit to a specific pricing policy follows a cyclic pattern in equilibrium: He charges most of the time relatively high prices that bar low-type customers from buying and allows for a large profit extraction from the high types. Nevertheless, after some periods the mass of low types is so large that he drastically reduces the price and all consumers on the market buy the good. Then the cycle starts again. Sobel (1984) extends this result to an oligopoly setting. Considering the monopolistic seller again, Sobel (1991) identifies further equilibria and also analyzes the game under full commitment power. In this case, the seller chooses a constant price over time that either caters exclusively to high-type consumers or serves both consumer types. This result is in line with the fact that uniform pricing is the optimal dynamic pricing strategy in a broad class of durable-good models with commitment.

Nocke and Peitz (2007) challenge this view and provide a rationale for clearance sales despite commitment power. In a two-period model, a monopolistic firm chooses her total capacity and the price for each period at the beginning of the game. They show that clearance sales (high price in the first period and a low one in the second) can be an optimal policy, because it allows to price discriminate between high and low valuation consumers. High-valuation types are detered from waiting for the low price, since the good is (randomly) rationed in the second period and there is a positive probability of not obtaining the good at all.

[12]In fact, consumer behavior is not explicitly modeled but summarized in a parametrized revenue function.

3.3 Model

In our model, there are two periods and the set of players consists of the firm, the club, and an infinite number of consumers. The firm produces a durable good at a marginal cost of zero. Each consumer wants to purchase exactly one unit of the good, but can freely choose in which of the two periods to buy (if at all). Consumers' willingness-to-pay is uniformly distributed on $[0, V]$ with a total consumer mass of M. Demand is stochastic: With probability $\lambda \in (0,1)$ demand is low, $(V, M) = (V_l, M_l)$, and with probability $1-\lambda$ demand is high, $(V, M) = (V_h, M_h)$, where $V_l, V_h > 0$ and $0 < M_l < M_h$.

The firm has two distribution channels available to reach consumers: She can either sell the good directly to consumers in her shop with marginal distribution cost $c \in \left(0, \bar{V}\right)$, where $\bar{V} := \min(V_l, V_h)$, or sell a specific capacity at a transfer price to the club for distribution.[13] We assume that simultaneous sales via both distribution channels are not feasible.[14] Moreover, the firm cannot commit to a certain second-period policy in any dimension of her choice variables in the first period. If the club obtains the good for distribution, it can sell the good at marginal distribution cost of zero up to its acquired capacity. The club has no possibility to store the good, but can dispose it free of charge.[15]

The precise timing of the game in period 1 is as follows:

1. Nature chooses the demand state, which is unobservable to all players and which remains fixed until the end of the game.

2. The firm decides whether to sell in her shop or sell quantity $q_1 \in \mathbb{R}_+$ at transfer price $T_1 \in \mathbb{R}_+$ to the club. If the club obtains an offer, it can accept it or not. In case of a refusal, the firm and the club obtain zero profits in this period.

3. All players observe the demand state.[16]

[13] Selling a whole load of goods is a common practice in industries, where production has already occured and producers try to minimize the risk of overstock at the end of the season. Thus, besides capacity, we rule out contractual agreements that influence the club's pricing decision, e.g. through wholesale prices or two-part tariffs.

[14] This restriction is innocuous but simplifies the analysis substantially. We show in Section 3.5 that our results also go through if we allow for parallel sales of shop and club.

[15] In fact, shopping clubs collect customer orders in sales deals and procure goods directly from the brand producers' storage facilities for further packaging and delivery.

[16] Section 3.5 and Footnote 7 provide a motivation for the timing how risk is resolved.

4. The player distributing the good sets a consumer price $p_1 \in \mathbb{R}_+$. If the club sells the good, it is restricted to consumer prices that do not violate its capacity constraint (no rationing).

5. Each consumer decides whether to buy or not. If a consumer buys the good, she disappears from the market until the end of the game.

In the second period, steps 2, 4, and 5 are repeated with the adjustment to prices $T_2, p_2 \in \mathbb{R}_+$ and capacity $q_2 \in \mathbb{R}_+$. After period 2, the game ends.

The firm and the club maximize their expected profits and each consumer her surplus. All players have an outside option of value zero and discount their respective second-period profit and surplus with a common discount factor $\delta \in (0, 1)$.

We characterize pure-strategy Subgame Perfect Equilibria of this game.

3.4 Analysis

We solve the game by backward induction. It is convenient, however, to start with subgames that follow first-period consumer-price announcements. This allows us to derive the firm's optimal second-period behavior and establish a unique second-period demand and consumer price.

So let H_1 be a first-period history consisting of the demand state (V, M), the firm's channel decision, and the first-period consumer price p_1 and consider the subgame following H_1. It is readily verified that in any equilibrium of this subgame there exists a cutoff consumer $v(H_1) \in [0, V]$ with the property that consumers with $v > v(H_1)$ buy in period 1 and consumers with $v < v(H_1)$ do not. This follows from consumers' optimal behavior in equilibrium. Note that $v(H_1)$ is always strictly larger than zero unless the first-period price in H_1 is zero. Obviously, independent of the choice of distribution channel, a first-period price of zero gives the firm a profit over both periods of at most zero, which is never optimal.[17] Thus, for the remainder of the paper, we restrict attention to subgames following histories H_1 with $p_1 > 0$. This implies that $v(H_1) > 0$ and second-period demand consists of an interval of low-valuation consumers that can still be served. Because the

[17] The firm could sell the good in the first period at the shop at a price of $p_1 > V$ such that all customers postpone consumption until period 2. It is then optimal for her to delegate the second-period distribution to the club and extract all club profits. Therefore, from a first-period perspective the firm can obtain the strictly positive one-period monopoly profits discounted by δ.

club is more cost-efficient than the shop and the firm can extract all of the club's profits due to her bargaining power, she always chooses the club for distributing the good in period 2.

Lemma 1. *In any equilibrium of the subgame following first-period history H_1, the club serves consumers in period 2. Moreover, the club obtains zero profits and the firm attains the club's one-period monopoly profits of second-period consumers.*

PROOF: The proof is relegated to the appendix. □

Lemma 1 allows us to pin down the second-period consumer price in an equilibrium following first-period history H_1. If $v(H_1)$ is the equilibrium cutoff consumer, second-period demand is given by $D_2(p_2|H_1) := \frac{M}{V}(v(H_1)-p_2)$. The club chooses the second-period price $p_2 = \frac{v(H_1)}{2} = \arg\max_p \; pD_2(p|H_1) - T_2$ to maximize profits.[18] This observation has two useful implications: First, the cutoff consumer $v(H_1)$ and second-period consumer price p_2 are *unique* across equilibria following first-period history H_1. This follows from a simple intuition: Let there be two equilibria after history H_1 with different cutoff consumers. The higher cutoff consumer requires a lower second-period price in that equilibrium, because less consumers buy in period 1. But this is incompatible with subgame perfect behavior of the club—a higher cutoff consumer is associated with a higher second-period price.[19, 20]

Second, due to the uniqueness of $(v(H_1), p_2)$ after history H_1, we can neglect the dependence of the cutoff consumer on the choice of the distribution channel in period 1. Because consumers only care about prices it holds that $v(H_1) = v(\tilde{H}_1)$ for all first-period histories

[18]Note also that the quantity constraint of the club in the maximization problem is not binding, because the firm extracts the maximum profits from consumers by Lemma 1.

[19]A rigorous argument is as follows: Let there be two equilibria after history H_1 with cutoff consumer and second-period price $(v(H_1), p_2)$ and $(\hat{v}(H_1), \hat{p}_2)$ respectively, where $(v(H_1), p_2) \neq (\hat{v}(H_1), \hat{p}_2)$. Note that $v(H_1) = \hat{v}(H_1)$ implies $p_2 = \hat{p}_2$. Hence, without loss of generality, let $0 < v(H_1) < \hat{v}(H_1) \leq V$, which involves $p_2 < \hat{p}_2$. Consider consumer $v \in (v(H_1), \hat{v}(H_1))$ and let p_1 be the first-period price of history H_1. For this consumer it must hold that $\delta(v - p_2) \leq v - p_1 \leq \delta(v - \hat{p}_2)$, but this yields the contradiction $p_2 \geq \hat{p}_2$.

[20]Note that this rationale does not hold true if there also exist equilibria following H_1 in which the shop is used in period 2. In that case a shop equilibrium (with a low cutoff consumer and a high price) can coexist with a club equilibrium (with a high cutoff consumer and a low price). The low price in the latter equilibrium follows from the club's cost advantage.

H_1 and $\tilde{H}_1$ with the same demand state and first-period price p_1.[21] For convenience, in slight abuse of notation, we will also suppress the dependence on the demand state. Thus, in the following we will write $v(p_1)$ for the cutoff consumer with the implicit understanding that this value depends on the realization of the random variables (V, M).

We now derive the cutoff consumer, the second-period consumer price, and demand in both periods as a function of p_1. By definition of the cutoff consumer, for all $v(p_1) \in (0, V)$ this consumer is indifferent between consumption in period 1 and 2. Thus, it holds that $v(p_1) - p_1 = \delta(v(p_1) - p_2) = \delta\left(v(p_1) - \frac{v(p_1)}{2}\right)$, and the cutoff consumer is given by

$$v(p_1) \quad = \quad \begin{cases} \frac{2}{2-\delta}p_1 & \text{if } p_1 \in \left[0, \frac{2-\delta}{2}V\right], \\ V & \text{if } p_1 > \frac{2-\delta}{2}V. \end{cases}$$

It is easy to verify that the interior second-period consumer price is

$$p_2(p_1) \quad := \quad \frac{v(p_1)}{2} \quad = \quad \frac{1}{2-\delta}p_1.$$

First- and second-period demand then amount to

$$D_2(p_1) \quad := \quad \frac{M}{V}\,(v(p_1) - p_2(p_1)) \quad = \quad \frac{M}{V}\,\frac{1}{2-\delta}p_1$$

and

$$D_1(p_1) \quad := \quad \frac{M}{V}\,(V - v(p_1)) \quad = \quad \frac{M}{V}\,\frac{V(2-\delta) - 2p_1}{2-\delta}$$

for such interior prices.[22]

We now scrutinize the firm's optimal first-period behavior. Before we do so, let us briefly consider the benchmark case of vertical integration. Obviously, profit maximization

[21] Without uniqueness of $(v(H_1), p_2)$, equilibria with different second-period demand and profits after history H_1 can exist. Thus, the first-period channel choice cannot be ignored, since a different channel can potentially induce a different second-period behavior.

[22] We do not explicitly write out the boundary cases for the second-period price and demand in both periods. They become relevant in the proofs of our results, but are not essential for an intuitive understanding in our exposition.

requires that the good is also sold via the club in period 1.[23] The optimal first-period consumer price trades off first-period and discounted second-period profits to maximize

$$\pi(p_1) \quad := \quad p_1 D_1(p_1) + \delta p_2(p_1) D_2(p_1).$$

A short calculation shows that the unique optimum is attained at $p_1^* = \frac{V(2-\delta)^2}{8-6\delta}$ where distributed quantity equals $x_1^* = \frac{M(2-2\delta)}{4-3\delta}$. Reintroducing the demand state again, optimal profits $\pi^*(V, M) := \pi(p_1^*)$ are a function of (V, M). Thus, expected profits before the demand state is revealed amount to

$$\pi^* = \lambda \pi^*(V_l, M_l) + (1-\lambda)\pi^*(V_h, M_h).$$

Observe that optimal distribution under vertical integration combines the cost advantage of the club with an optimal consumer-price choice, which maximizes the total profits of both periods. We now show that each distribution channel, shop and club, can achieve only one aspect of optimal distribution.

If the firm decides to sell her good via the shop, she bears strictly positive marginal distribution costs. On the other hand, she can freely choose the first-period price p_1 and obtains total profits of

$$\pi^S(p_1) \quad := \quad (p_1 - c) D_1(p_1) + \delta p_2(p_1) D_2(p_1).$$

This expression is uniquely maximized at $p_1^S = \frac{V(2-\delta)^2 + 2c(2-\delta)}{8-6\delta}$ if $c < (1-\delta)V$. When δ is large, the cost of not catering consumers in period 1 is relatively small and the firm prefers exclusive selling via the club at low distribution costs in period 2. Thus, for $c \geq (1-\delta)V$, she charges a price $p_1^S \geq \frac{2-\delta}{2} V$. Total profits under selling via the shop are $\hat{\pi}^S(V, M, c) := \pi^S(p_1^S)$. Hence, when the firm decides to sell via the shop she has expected profits of

$$\hat{\pi}^S(c) \quad = \quad \lambda \hat{\pi}^S(V_l, M_l, c) + (1-\lambda)\hat{\pi}^S(V_h, M_h, c).$$

[23] To see this, note that the club can always choose the same first-period price as the shop, but does not have to bear the strictly positive distribution cost.

While the shop channel involves optimal price setting, the good is distributed under strictly positive marginal costs. Therefore, when the firm has to make her channel decision, she expects profits of the shop to be strictly below the ones under vertical integration. As the marginal distribution cost of the shop decreases, however, the former converge to the latter. Not surprisingly, expected shop profits can only benefit from lower cost when the shop is actually involved in distribution, i.e. whenever $c < (1-\delta)V$.

Lemma 2. *In period 1, expected shop profits are strictly below expected profits under vertical integration, i.e.* $\hat{\pi}^S(c) < \pi^*$ *for any* $c \in (0, \bar{V})$. *Moreover,* $\hat{\pi}^S(c)$ *is continuous in* c *with* $\lim_{c \searrow 0} \hat{\pi}^S(c) = \pi^*$. *Finally,* $\hat{\pi}^S(c)$ *is strictly decreasing in* c *for* $c \in (0, (1-\delta)\bar{V})$.

Proof: The proof is relegated to the appendix. □

We now consider selling via the club. Recall from Lemma 1 that the firm extracts all second-period profits from the club. This result also extends to period 1. Intuitively, if the club expects strictly positive first-period profits after accepting the firm's contract, the firm can charge it a slightly higher transfer price T_1 for the same capacity of goods. Because this fixed fee is a sunk cost for the club, it does not alter its first-period consumer price in the respective demand states such that second-period profits of the firm remain unchanged. Thus, the firm has a profitable deviation.

Lemma 3. *In any Subgame Perfect Equilibrium in which the firm distributes the good via the club in period 1, she extracts all expected first-period profits from the club.*

Proof: The proof is relegated to the appendix. □

If the firm distributes the good over the club channel in period 1, Lemma 3 allows us to restrict attention to contracts (q_1, T_1), where T_1 extracts the whole surplus from the club. Thus, if the club chooses the first-period price $p^C_{1,l}$ and $p^C_{1,h}$ in the respective demand states, the firm obtains a total profit of

$$T_1 + \delta\left(\lambda T_2(V_l, M_l) + (1-\lambda)T_2(V_h, M_h)\right) \quad = \quad \lambda\pi(p^C_{1,l}, V_l, M_l) + (1-\lambda)\pi(p^C_{1,h}, V_h, M_h),$$

because the good is sold by the club in both periods.

Suppose that the club is not capacity constrained in its choice of consumer prices. Fixing the demand state for a second, note that it is not in the club's interest to select the first-period consumer price p_1^*, which maximizes $\pi(p_1)$. Since the club anticipates that its future second-period profits will be absorbed by the firm, it optimizes just its first-period profits. Hence, it chooses the one-period monopoly price $p_1^m := \frac{V(2-\delta)}{4} = \arg\max_p \; pD_1(p)$ and sells quantity $x_1^m := \frac{M}{2}$. Because the club does not internalize second-period profits, which are increasing in the first-period price, the club's consumer price is too low ($p_1^m < p_1^*$) and it sells too much ($x_1^m > x_1^*$) from the firm's perspective.

Let us now reincorporate the fact that the firm's capacity choice can influence the club's pricing behavior: By restricting capacity provided to the club, she can induce it to set a higher price than p_1^m. More precisely, if $D_1^{-1}(x_1)$ is the first-period inverse demand function and q_1 the firm's capacity choice, the club selects a first-period consumer price

$$p_1^C(q_1) := \begin{cases} p_1^m & \text{if } q_1 \geq x_1^m, \\ D_1^{-1}(q_1) & \text{if } q_1 \in [0, x_1^m). \end{cases}$$

Hence, including the dependence of $p_1^C(q_1)$ on the demand state again, if the firm distributes via the club under capacity choice q_1, she obtains profits of

$$\pi^C(q_1, V, M) \quad := \quad \pi(p_1^C(q_1)).$$

Observe, however, that capacity is contracted before the realization of the demand state. Therefore, when the firm has to decide on her first-period channel her expected profits from using the club amount to

$$\hat{\pi}^C \quad := \quad \max_{q_1} \; \lambda \pi^C(q_1, V_l, M_l) + (1-\lambda)\pi^C(q_1, V_h, M_h)$$

We now show that the firm's ability to influence the club's pricing through her capacity choice is limited: In at least one demand state, the club does not select the vertical-integration consumer price p_1^*. To see this, note that the vertical integration quantity x_1^* is strictly increasing in M and independent of V. Thus, x_1^* is larger in the high demand state than in the low one. Hence, a sufficiently small capacity leads to the vertical integration price in the low demand state, but an exceedingly high price under high demand. If the capacity constraint is large enough to implement the vertical integration price in the high demand state, it will be slack, however, under low demand. Therefore, in this demand state the club chooses a price that is too low from the firm's perspective. For any other capacity choices there will be distortions in both demand states. It follows from this discussion that at the time of channel selection, the firm's expected profits from using the club are strictly below the vertical integration profits. Even though the firm manages to distribute the good under efficient cost, she cannot provide appropriate incentives for the club to select the vertical-integration consumer prices. We formally prove this result in

Lemma 4. *In period 1, expected club profits are strictly below expected profits under vertical integration, i.e.* $\hat{\pi}^C < \pi^*$.

PROOF: The proof is relegated to the appendix. □

We now turn to our main result. Lemma 2 and 4 show that both distribution channels cannot reach the vertical integration profits π^*. While the club does not internalize the tradeoff between first-period and second-period profits, the shop sells the good under inefficient distribution costs. The shop's disadvantage, however, increasingly vanishes as the distribution cost falls and the firm's profits converge to π^* (Lemma 2). Thus, for sufficiently small distribution cost c, selling the good through the shop channel is more profitable. It follows that in any Subgame Perfect Equilibrium, the firm will sell the good via the shop in period 1 if the distribution cost is sufficiently small.

Theorem 1. *There exists a* $c^* > 0$ *such that for all* $c \in (0, c^*)$ *it holds: In any Subgame Perfect Equilibrium the firm uses the shop in period 1 and the club in period 2.*

PROOF: The proof is relegated to the appendix. □

Theorem 1 is in line with our observation that shopping clubs are used for clearance, but not for regular sales. If the club's cost advantage is not too big, the firm prefers her own store as a distribution channel, because the value of consumer-price control in early periods is more profitable than low distribution costs.

3.5 Discussion

We finish the analysis by discussing some assumptions of the model. First, it is important in our setup that the firm produces a good with at least some degree of durability. This is clearly the case for clothes and accessories, which are predominantly sold by shopping clubs. For simplicity we chose a fully durable good for our model, but our results also go through if just a fraction of consumers leaves the market after a purchase in period 1. Of course, if goods are entirely perishable, all consumers who buy in period 1 reappear in period 2 and we have two sales periods that are not intertemporally linked with each other. But then optimal profits are attained by using the club in both periods. This is inconsistent with the observation that shopping clubs are employed for clearance sales only.

Second, there has to be a commitment problem for the firm. This concerns not only her ability to set future consumer prices in her shop, but also to fix in advance a specific capacity to be sold to the club. It is well-known from the standard durable-goods problem that under commitment profits are maximized by charging in both periods the one-period monopoly price. Consumers do not expect prices to fall in the future and each consumer with a valuation above the first-period price purchases the good in period 1. This outcome is also achieved in our model. If one of the above types of commitment is possible, it is easily verified that the firm attains the same profits as under vertical integration. It involves, however, the club being active in period 1, which we do not observe in reality.

We believe that a lack of producer commitment power is a reasonable assumption for the goods that shopping clubs distribute. While the degree of commitment power may vary from industry to industry, the fashion sector is notorious for engaging frequently in clearance sales. Thus, in this industry it seems particularly hard for sellers to keep prices high over time.

Third, demand uncertainty is important for having a genuine conflict of interest between the firm and the club: In at least one demand state the club chooses a consumer price that is suboptimal from the firm's perspective. Under demand certainty, the firm can fine tune the sold capacity in such a way that the club selects an optimal first-period price for the firm. Again, this implies that the club sells in both periods.

Although demand uncertainty is a very realistic assumption, the timing of how it resolves is essential, too. If the demand state is revealed before the firm contracts with the club, one obtains the same result as under deterministic demand. As we have argued in Footnote 7, the seller's ability to learn instantaneously the demand state when engaging with consumers can be a good approximation. When the time of consumer-price experimentation to assess the demand state is relatively short compared to total period length, large parts of the period profits are generated from the seller's final price choice. Thus, the seller approximately behaves as if he knows the demand state.

Fourth, in contrast to the previous points, our results do not change if we allow for simultaneous sales of shop and club. To see this, assume that under parallel sales shop and club engage in Bertrand price competition and consumers buy from the provider with the lowest price. The club, of course, can only serve consumers up to its capacity constraint. Note that in the second period nothing changes: the club is used as sole distribution channel and all of its profits are absorbed by the firm. Regarding period 1, the analysis is a bit more intricate. We show that total firm profits under simultaneous sales are strictly below the ones under vertical integration when the shop's distribution cost becomes sufficiently low. Unlike pure-shop profits, profits from parallel sales do not converge to the vertical-integration profits. Hence, Theorem 1 remains valid and selling exclusively via the shop in period 1 is the dominant choice of the firm for small distribution cost.

Note that under simultaneous sales in period 1, the firm chooses the shop's price to maximize total profits given her expectations about the club's price choice. An important observation is that the firm's second-period profits are continuous in the first-period price. Therefore, in the first period, the firm has an incentive to undercut the club's price as long as it is above her distribution cost.

Assume for a second that the club is not capacity constrained. Then, the first-period behavior of shop and club entirely resembles price competition of homogeneous-good duopolists with different marginal cost. Thus, if the shop's distribution cost is sufficiently small, shop and club set a first-period equilibrium price equal to the shop's distribution cost and all consumers buy from the club.[24] The shop's distribution cost is sufficiently small if it is below the club's first-period monopoly price.[25]

[24]There exist further equilibria with prices strictly below the shop's distribution cost. First, following the reasoning further below, in these equilibria total profits fall short of vertical-integration profits even more. Second, if the club has binding capacity constraints, such equilibria will not occur, because the firm sells some units at a price below distribution cost.

[25]In contrast, if distribution cost differences are drastic, the club can behave virtually as a monopolist and charge the monopoly price. As we are only interested in the case when cost differences become small, we leave this scenario aside.

Note that these equilibrium prices are also obtained if the club is capacity constrained.[26] This follows because the shop can still serve the whole market and undercuts the club's offer as before.[27] To sum up, if distribution cost differences become small, the first-period price under simultaneous sales equals the shop's distribution cost. This implies that the first-period price is smaller or equal the first-period monopoly price, which, in turn, lies strictly below the one under vertical-integration. Hence, total firm profits under parallel sales fall short of vertical-integration profits and even further decrease as the shop's distribution cost gets smaller. Thus, Theorem 1 generalizes to the case when we allow for simultaneous sales of shop and club.

On a broader perspective, one may wonder why the firm does not open an own online channel and profit from low distribution costs. We suggest two arguments against this objection:

First, single firms may not generate enough turnover to benefit from distributing their goods online. Products have to be prepared for online sales and a logistic infrastructure needs to be maintained. Shopping clubs have many different labels in their brand portfolios and can take advantage of economies of scale: All goods are processed in a central department and clubs can negotiate very low fares with shipping companies. Thus, even though a firm can add an online branch to her existing store facilities, she may not enjoy the same cost advantage as shopping clubs do.

Second, if firms try to establish their brand in a new market region, there has to be a store, where consumers can virtually experience the good and get an impression of its quality. This is not possible in the anonymity of the internet. Thus, physical store presence preceeds online distribution. Given the high setup costs for launching an online store with its logistic system, firms may prefer to test the market environment before adding an own internet channel.[28] Shopping clubs can serve as handy distributors in this experimentation phase. This argument by itself implies that shopping clubs vanish over time as brand producers introduce their own online stores. This is in contrast to the

[26] What does change is the number of consumers the club is selling the good when its capacity constraint is binding. Therefore, it is optimal to impose no binding capacity constraints on the club to avoid the shop inefficiently distributing a positive amount of goods.

[27] For drastic cost differences, the club's price lies somewhere between its first-period monopoly price and the shop's distribution cost depending on the capacity constraint.

[28] For instance, *Zara*, a Spanish fashion label, has stores in many countries, but in just few of them also operates an online store.

impressive growth of the shopping club industry in recent years. In combination with our preceding cost argument, however, it explains how shopping clubs can emerge in this testing period and also persist on the market in the long run.

Finally, physical outlet stores are also distribution channels with very low distribution costs, because they are located in cheap rural and suburban areas. While brand producers can directly influence prices in such stores, they face worse demand conditions than in traditional shops: Consumers have to bear large search costs to visit these remote locations and just a fraction of customers will buy there. Thus, despite being an attractive distribution alternative in terms of costs, a brand producer may favor to incorporate shopping clubs in his sales strategy.

3.6 Appendix

Proof of Lemma 1: In each equilibrium following history H_1 there exists a cutoff consumer $v(H_1) > 0$. Hence, for a specific equilibrium, second-period demand is given by $D_2(p_2|H_1) := \frac{M}{V}(v(H_1) - p_2)$. If the club distributes the good to consumers, it charges the one-period monopoly price $p_2^m = \frac{v(H_1)}{2} = \arg\max_p \; pD_2(p|H_1) - T_2$ and sells quantity $x_2^m = D_2(p_2^m|H_1)$.

Now assume there exists an equilibrium in which the firm distributes the good via the shop and obtains profits π_2. Because the club is more cost efficient and $v(H_1) > 0$, we must have $\pi_2 < p_2^m x_2^m$. Define $\hat{\epsilon} := p_2^m x_2^m - \pi_2 > 0$ and consider the firm decision to sell via the club under contract $(\hat{q}_2, \hat{T}_2) := (x_2^m, p_2^m x_2^m - \frac{\hat{\epsilon}}{2})$. Observe that the club accepts this offer, because by setting its optimal price p_2^m it obtains a positive profit of $\frac{\hat{\epsilon}}{2}$. Thus, the firm receives a profit of $\hat{T}_2 > \pi_2$ and has a profitable deviation.

Finally, assume there exists an equilibrium in which the club distributes the good, but the firm does not obtain the club's one-period monopoly profits of the second-period consumers. (Note that this is particularly the case when the club obtains positive profits.) So let (q_2, T_2) be the equilibrium contract of the firm with the club, where $T_2 < p_2^m x_2^m$. Define $\tilde{\epsilon} := p_2^m x_2^m - T_2 > 0$ and consider contract $(\tilde{q}_2, \tilde{T}_2) := (x_2^m, T_2 + \frac{\tilde{\epsilon}}{2})$. A similar reasoning as above shows that this contract is a profitable deviation for the firm. □

Proof of Lemma 2: It is sufficient to show that the results hold for each demand state $(V, M) \in \{(V_l, M_l), (V_h, M_h)\}$ individually. Inserting p_1^* into $\pi(p_1)$ yields

$$\pi^*(V, M) \quad = \quad \frac{MV(2-\delta)^2}{4(4-3\delta)},$$

and $\pi^S(p_1)$ evaluated at p_1^S gives

$$\hat{\pi}^S(V, M, c) \quad = \quad \begin{cases} \frac{M(V^2(2-\delta)^2+4c^2-8cV(1-\delta))}{4V(4-3\delta)} & \text{if } c < (1-\delta)V, \\ \frac{MV\delta}{4} & \text{if } c \geq (1-\delta)V. \end{cases}$$

Observe that $\hat{\pi}^S(V, M, c)$ is continuous at $c = (1-\delta)V$. Moreover, each branch of the function is continuous in c, so that $\hat{\pi}^S(V, M)$ is continuous in c. It is also immediately verified that $\lim_{c \searrow 0} \; \hat{\pi}^S(V, M, c) = \pi^*(V, M)$.

To prove $\hat{\pi}^S(V,M,c) < \pi^*(V,M)$, consider the lower branch of $\hat{\pi}^S(V,M,c)$ and note that

$$\begin{aligned}
MV\,\frac{\delta}{4} &= \frac{MV\delta(4-3\delta)}{4(4-3\delta)} \\
&= \frac{MV(2-\delta)^2}{4(4-3\delta)} - \frac{MV((2-\delta)^2-\delta(4-3\delta))}{4(4-3\delta)} \\
&= \pi^*(V,M) - \frac{MV(4-4\delta+\delta^2-4\delta+3\delta^2)}{4(4-3\delta)} \\
&= \pi^*(V,M) - \frac{MV(4-8\delta+4\delta^2)}{4(4-3\delta)} \\
&= \pi^*(V,M) - \frac{MV(2-2\delta)^2}{4(4-3\delta)} \\
&< \pi^*(V,M).
\end{aligned}$$

For the upper branch of $\hat{\pi}^S(V,M,c)$ it holds that

$$\begin{aligned}
\frac{M(V^2(2-\delta)^2+4c^2-8cV(1-\delta))}{4V(4-3\delta)} &= \pi^*(V,M) + \frac{M(4c^2-8cV(1-\delta))}{4V(4-3\delta)} \\
&\overset{c\leq V(1-\delta)}{\leq} \pi^*(V,M) - \frac{M4c^2}{4V(4-3\delta)} \\
&< \pi^*(V,M).
\end{aligned}$$

Finally, we show that $\hat{\pi}^S(V,M,c)$ is strictly decreasing in c for $c \in (0,(1-\delta)\bar{V})$. Note that $\bar{V} \leq V$, such that the upper branch of $\hat{\pi}^S(V,M,c)$ is relevant for this cost range. By considering the first derivative of this branch, we get

$$\frac{\partial}{\partial c}\,\hat{\pi}^S(V,M,c) = \frac{M(8c-8V(1-\delta))}{4V(4-3\delta)} < 0.$$

□

PROOF OF LEMMA 3: Assume there exists a Subgame Perfect Equilibrium in which the firm distributes the good via the club in period 1 and the club obtains positive expected first-period profits $\epsilon > 0$. Let (q_1, T_1) be the firm's first-period contract offered to the club and let $p^C_{1,l}$ and $p^C_{1,h}$ the club's first-period prices in the low and high demand state respectively. Thus, because the firm extracts all second-period profits of the club by Lemma 1, her total expected equilibrium profits are

$$\pi \;:=\; T_1 \;+\; \delta\left[\lambda\, p_2(p_{1,l}^C)D_2(p_{1,l}^C) + (1-\lambda)\, p_2(p_{1,h}^C)D_2(p_{1,h}^C)\right].$$

Consider contract $(\hat{q}_1, \hat{T}_1) := (q_1, T_1 + \frac{\epsilon}{2})$. The club still accepts this offer and—since no change to the available capacity has been made—chooses again the first-period prices $p_{1,l}^C$ and $p_{1,h}^C$ in the respective demand states. Thus, the firm attains the same second-period profits as under (q_1, T_1), but higher first-period profits. Hence, she has a profitable deviation. □

Proof of Lemma 4: It is sufficient to show that for all $q_1 \in \mathbb{R}_+$ there exists a demand state $(V, M) \in \{(V_l, M_l), (V_h, M_h)\}$ such that $\pi^C(q_1, V, M) < \pi^*(V, M)$. Let $x_{1,l}^*$ and $x_{1,h}^*$ be the vertical integration quantities in the low and high demand state, respectively, and $p_{1,l}^*$ and $p_{1,h}^*$ the associated prices. Similarly, define $x_{1,l}^m$ and $x_{1,h}^m$ as the one-period monopoly prices in the low and high demand state. Observe that $x_{1,l}^* < x_{1,h}^*$ and $x_{1,i}^* < x_{1,i}^m$ for $i \in \{l, h\}$.

First, consider $q_1 \leq x_{1,l}^*$. Since $x_{1,h}^* > x_{1,l}^*$ the club chooses in the high demand state a first-period price $p_1 > p_{1,h}^*$. Thus, because $p_{1,h}^*$ uniquely maximizes $\pi(p_1)$ in the demand state (V_h, M_h), we have $\pi^C(q_1, V_h, M_h) < \pi^*(V_h, M_h)$ for all $q_1 \leq x_{1,l}^*$. Now consider $q_1 > x_{1,l}^*$. If the low-demand state realizes, the club sells quantity $x_1 = \min(q_1, x_{1,l}^m) > x_{1,l}^*$, such that $p_1 < p_{1,l}^*$. Because $p_{1,l}^*$ uniquely maximizes $\pi(p_1)$ in the demand state (V_l, M_l), we have $\pi^C(q_1, V_l, M_l) < \pi^*(V_l, M_l)$ for all $q_1 > x_{1,l}^*$. □

Proof of Theorem 1: The firm's club choice in period 2 was already shown in Lemma 1. Regarding the first period, we have to show that there exists a $c^* > 0$ such that for all $c \in (0, c^*)$ it holds that $\hat{\pi}^S(c) > \hat{\pi}^C$. Without loss of generality we restrict attention to $c \in (0, \hat{V})$, where $\hat{V} := (1-\delta)\bar{V}$. Note that, $\hat{\pi}^S(c)$ is a strictly decreasing function of c on that domain by Lemma 2. If $\hat{\pi}^S(c) > \hat{\pi}^C$ for all $c \in (0, \hat{V})$, set $c^* = \hat{V}$ and we are already done. Now assume there exists $\tilde{c} \in (0, \hat{V})$ such that $\hat{\pi}^S(\tilde{c}) \leq \hat{\pi}^C$. Because $\lim_{c \searrow 0} \hat{\pi}^S(c) = \pi^* > \pi_C^*$ and $\hat{\pi}^S(c)$ is continuous in c (Lemma 2), by the Intermediate Value Theorem, we can find a $\hat{c} \in (0, \tilde{c}]$ such that $\hat{\pi}^S(\hat{c}) = \hat{\pi}^C$. By the strict monotonicity of $\hat{\pi}^S(c)$ on $(0, \hat{V})$, we have $\hat{\pi}^S(c) > \hat{\pi}^C$ for all $c \in (0, \hat{c})$. Thus, we can set $c^* = \hat{c}$ and the proof is complete. □

Bibliography

BAKER, G., M. GIBBS, AND B. HOLMSTROM (1994): "The Wage Policy of a Firm," *The Quarterly Journal of Economics*, 109, 921–955.

BEAUDRY, P. (1994): "Why an Informed Principal May Leave Rents to an Agent," *International Economic Review*, 35, 821–832.

CHADE, H. AND R. SILVERS (2002): "Informed Principal, Moral Hazard, and the Value of a More Informative Technology," *Economic Letters*, 74, 291–300.

CHIANG, W. K., D. CHHAJED, AND J. D. HESS (2003): "Direct Marketing, Indirect Profits: A Strategic Analysis of Dual-Channel Supply-Chain Design," *Management Science*, 49, 1–20.

CHO, I.-K. AND D. M. KREPS (1987): "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, 102, 179–221.

CONLISK, J., E. GERSTNER, AND J. SOBEL (1984): "Cyclic Pricing by a Durable Goods Monopolist," *The Quarterly Journal of Economics*, 99, 489–505.

CRÉMER, J. AND R. P. MCLEAN (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247–1257.

DESAI, P., O. KOENIGSBERG, AND D. PUROHIT (2004): "Strategic Decentralization and Channel Coordination," *Quantitative Marketing and Economics*, 2, 5–22.

EPSTEIN, L. G. AND M. PETERS (1999): "A Revelation Principle for Competing Mechanisms," *Journal of Economic Theory*, 88, 119–160.

FISCHER, G. (2005): *Lineare Algebra*, Wiesbaden: Vieweg & Sohn Verlag.

Fudenberg, D. and J. Tirole (2002): *Game Theory*, Cambridge, Massachusetts: MIT Press.

Gibbons, R. and L. F. Katz (1991): "Layoffs and Lemons," *Journal of Labor Economics*, 9, 351–380.

Greenwald, B. C. (1986): "Adverse Selection in the Labour Market," *Review of Economic Studies*, 53, 325–347.

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

Inderst, R. (2001): "Incentive Schemes as a Signaling Device," *Journal of Economic Behavior & Organization*, 44, 455–465.

Kahn, C. (1986): "The Durable Goods Monopolist and Consistency with Increasing Costs," *Econometrica*, 54, 275–294.

Karp, L. S. and J. M. Perloff (1996): "The Optimal Suppression of a Low-Cost Technology by a Durable-Good Monopoly," *The RAND Journal of Economics*, 27, 346–364.

Kutsoati, E. and J. Zábojník (2005): "The effects of learning-by-doing on product innovation by a durable good monopolist," *International Journal of Industrial Organization*, 23, 83–108.

Lazear, E. P. (1986): "Raids and Offer Matching," in *Research in Labor Economics*, ed. by R. Ehrenberg, vol. 8, Part A, 141–165.

——— (2000): "Performance Pay and Productivity," *The American Economic Review*, 90, 1346–1361.

Lazear, E. P. and M. Gibbs (2009): *Personnel Economics in Practice*, Hoboken: John Wiley & Sons.

Lebow, D., L. Sheiner, L. Slifman, and M. Starr-McCluer (1999): "Recent Trends in Compensation Practices," Federal Reserve Board Working Paper.

Lehmann, D. R. and C. B. Weinberg (2000): "Sales through Sequential Distribution Channels: An Application to Movies and Videos," *Journal of Marketing*, 64, 18–33.

Maskin, E. and J. Tirole (1990): "The Principal-Agent Relationship with an Informed Principal: The Case of Private Values," *Econometrica*, 58, 379–409.

McGuire, T. W. and R. Staelin (1983): "An Industry Equilibrium Analysis of Downstream Vertical Integration," *Marketing Science*, 2, 161–191.

Mussa, M. and S. Rosen (1978): "Monopoly and Product Quality," *Journal of Economic Theory*, 18, 301–317.

Myerson, R. B. (1983): "Mechanism Design by an Informed Principal," *Econometrica*, 51, 1767–1797.

Mylovanov, T. and T. Tröger (forthcoming): "Informed-Principal Problems in Environments with Generalized Private Values," *Theoretical Economics*.

Nocke, V. and M. Peitz (2007): "A Theory of Clearance Sales," *The Economic Journal*, 117, 964–990.

Oyer, P. and S. Schaefer (2005): "Why Do Some Firms Give Stock Options to All Employees?: An Empirical Examination of Alternative Theories," *Journal of Financial Economics*, 76, 99–133.

Perri, T. J. (1995): "Is There a Winner's Curse in the Labor Market?" *Journal of Economic Behavior & Organization*, 28, 79–89.

Postel-Vinay, F. and J.-M. Robin (2004): "To Match or not to Match? Optimal Wage Policy with Endogenous Worker Search Intensity," *Review of Economic Dynamics*, 7, 297–330.

Riordan, M. H. and D. E. M. Sappington (1988): "Optimal Contracts with Public Ex Post Information," *Journal of Economic Theory*, 45, 189–199.

Salop, J. and S. Salop (1976): "Self-Selection and Turnover in the Labor Market," *The Quarterly Journal of Economics*, 90, 619–627.

Silvers, R. (2012): "The Value of Information in a Principal-Agent Model with Moral Hazard: The Ex Post Contracting Case," *Games and Economic Behavior*, 74, 352–365.

Sobel, J. (1984): "The Timing of Sales," *The Review of Economic Studies*, 51, 353–368.

——— (1991): "Durable Goods Monopoly with Entry of New Consumers," *Econometrica*, 59, 1455–1485.

Stokey, N. L. (1981): "Rational Expectations and Durable Goods Pricing," *The Bell Journal of Economics*, 12, 112–128.

Stole, L. A. (2007): "Chapter 34 - Price Discrimination and Competition," in *Handbook of Industrial Organization*, ed. by M. Armstrong and R. Porter, Elsevier, vol. 3, 2221–2299.

Waldman, M. (1984): "Job Assignments, Signalling, and Efficiency," *The RAND Journal of Economics*, 15, 255–267.

Weiss, A. (1990): *Efficiency wages*, Princeton: Princeton University Press.