

Dynamic Contracts with Moral Hazard and Adverse Selection*

Alex Gershkov[†] Motty Perry[‡]

25-02-10

Abstract

This paper studies a novel dynamic principle agent setting with moral hazard and adverse selection (persistent as well as repeated). In the model an expert whose skills is his private information, faces a finite sequence of tasks, one after the other. Each task's level of difficulty is an independent random variable revealed, upon arrival, to the expert only. On each task in turn the expert choose whether to pass or to work, and how much effort to exert. While the choice of work/pass is public, his effort is his private information.

The optimal contract-pair which takes advantage of the dynamic nature of the interaction is characterized. It is shown that as the length of the contract increases, the expected transfer per-period goes down and in the limit approaches the optimal payment when agent's skills are publicly known.

One example of such a dynamic interaction is the one occurs between a money manager who receives funds from investors, and then observes a sequence of investment opportunities. Another example that nicely fits this model is the design of optimal contracts to surgeons of different quality, to treat a flow of patients whose problems are the surgeon's private information.

*We wish to thank Michael Borns for his invaluable editorial work. Both authors gratefully acknowledge financial support from the Israel Science Foundation. Alex Gershkov is grateful to Falk Institute for financial support.

[†]Alex Gershkov: Department of Economics, Hebrew University of Jerusalem, alexg@huji.ac.il.

[‡]Motty Perry: Department of Economics, Hebrew University of Jerusalem & University of Warwick, motty@huji.ac.il.

1 Introduction

This paper studies a novel dynamic principal–agent setting with moral hazard and adverse selection (persistent as well as repeated). In the model an expert whose skills are his private information faces a finite sequence of tasks, one after the other. Each task’s level of difficulty is an independent random variable revealed, upon arrival, to the expert only. On each task in turn (and before the next task arrives), the expert chooses first whether to work on the task or to pass, and if his choice is to work, how much effort to exert. While the choice of work/pass is public, his effort is his private information and is the source of the moral hazard in the model.

One example of such a dynamic interaction is that which occurs between a venture capital manager who receives funds from investors who desire to invest their money but lack the knowledge to do so personally. Indeed, contracts between investors and their portfolio managers have become a recent focus of both academic research and public interest. For a money manager investment opportunities arrive sequentially; some are easy to assess and manage, others, more difficult. The probability of success of a given investment is a function of the manager’s skill, the complexity of the investment, and the effort exerted in first analyzing and then following and directing once the investment is made. The interesting question then is how to design an optimal compensation contract in the context of the moral hazard and adverse selection problems that arise from the unobservability of the efforts made by the manager as well as of his skill and of the types of investment opportunities available. In particular, how to create incentive schemes that distinguish between unskilled money managers who cannot do well in complex investments from skilled managers who can.

One can think of many other examples that fit this model, but in an attempt to give focus to the discussion, and to fix a language, we discuss this problem in the context of a health-care insurer or public official who employs surgeons, whose quality he does not observe, to treat a flow of patients, the severity of whose problems is also the surgeon’s private information. Our principal’s problem, then, is to design a system of contracts that guarantee that surgeries are performed, and effort is exerted, if and *only* if the surgeon’s quality matches the severity of the patient’s problem, and to do so at minimal costs. The recent controversy over the health-care report-card system illustrates the type of incentives problems that sometimes arise in the medical industry (see Dranov et al. (2003)). This system entails a public disclosure of patient-health outcomes at the level of the individual physician. Many private insurers use this information and supporters argue that

the system gives providers powerful incentives to improve quality. Skeptics counter that report cards may encourage providers to “game” the system by avoiding sick patients, seeking healthy patients, or both.

While the use of money is natural in many markets, one may be tempted to believe (as indeed Fong (2010) suggests) that when it comes to medical treatment money does not play an important role in providing incentives. Empirical studies, however, do not support this view.¹ For example, Gruber et al. (1999) empirically show that the frequencies of cesarean deliveries compared to normal child births react positively to fee differentials of health insurance programs. Along the same line, Hughes and Yule (1992) documented that the number of cervical cytology treatments is correlated with the fee for this treatment. Indeed, not only do surgeons react to financial incentives, but they may overreact in a way that is not in the patient’s best interest. Emons (1997) cites a Swiss study reporting that the average person’s probability of receiving one of seven major surgical interventions is one third above that of a physician or a member of a physician’s family, and Wolinsky (1993, 1995) refers to a study by the Federal Trade Commission that documents the tendency of optometrists to prescribe unnecessary treatment.

In our model, agents sign contracts for T periods, and afterward in every period $t \in \{1, 2, \dots, T\}$, they encounter one task and decide whether to pass or to accept it and in the latter case, whether to exert a costly and unobservable effort. The probability of a successful accomplishment of a task in period t is monotonic in the agent’s effort at t , *but* it is also a function of the agent’s skills as well as the complexity of the task in question. While the agent’s skill is determined once and for all at $t = 0$, the type of the task is drawn independently each period, and both the agent’s skills and the types of the task are the agent’s private information.

For ease of exposition we confine our attention to a special case, probably the more interesting one, in which the principal’s preferences are lexicographic in the sense that his first goal is to align the complexity of the tasks with the quality of the agent while payment is only secondary. To be precise, depending on the agent’s skills and the complexity of the task, a different action is desired by the principal. If the task is simple, all types of agents (low- or high-skilled) should exert effort, but if the task is complex *only* the high-skilled agent should exert effort while the other one should pass on it.

¹On the other hand, the principal might find the use of money problematic for either moral or practical reasons.

As is always the case in solving for adverse selection, the principal presents the agent with a menu of contracts, one for each skill level, that promise financial rewards as a function of the observed history, i.e., whether a task was accepted and if so, whether it was successful. Conditional on the contracts providing the right incentives, the optimal contract is the one that minimizes payments. Thus, this is a dynamic model of repeated (and persistent) adverse selection and moral hazard problems. Moral hazard may arise since an agent's effort is unobservable, and adverse selection is a consequence of the superior information the agent has both about his own competence and the complexity of the task.

With this rather stylized model we characterize the optimal contract-pair which takes advantage of the dynamic nature of the interaction. It is shown that as the length of the contract increases, the expected transfer per period goes down and in the limit approaches the optimal payment when an agent's skills are publicly known. This can be done because the dynamic enables us to make the contract of the high-skilled agent unattractive to the low-skilled one, and more so as the contract gets longer. The intuition beyond this result is rather simple. Exerting effort on a task is a gamble whose probability of success is higher, the more skilled the agent is. A success in a sequence of tasks is exponentially less likely for a low-skilled agent. An optimal contract for a high-skilled agent takes advantage of this fact and stipulates high rewards conditional on a long history of successes. To construct these type of sequences and at the same time preserve incentives to exert effort, the optimal contract stipulates that in every period t , a success is rewarded *only* if it is followed by an uninterrupted sequence of successes until period T , the end of the contract.

The optimal contract for the high-skilled agent, in which he is compensated for a success in period t only if he keeps succeeding in every period thereafter till the end of the contract, looks rather extreme when T becomes large. Not only does it become very risky, it also has the unpleasant feature that no payment is guaranteed until the end of the contract at T . Of course, when the agent and the principal are risk-neutral and both discount the future at the same rate, as is assumed here, these features are irrelevant. Yet, we show that for T large enough, the set of optimal contracts is not a singleton and there are other contracts in the set in which these two features are relaxed dramatically. As we discuss in Section 6 once risk aversion is introduced, the optimal contract will take a less extreme form even for short-term contracts. Nevertheless, the essential phenomenon of using the dynamic to make the contract of the high-skilled agent more risky is preserved and is very much in line with common practice in, say, the medical

profession, where a high discontinuous increase in salary in the form of a promotion is promised only after a long history of successes. This practice helps weed out those surgeons whose private information suggests that they are low-quality and hence have a low probability of being promoted. They will opt out in order to avoid the inferior pay and work conditions that obtain before promotion. This result is also in line with the recent proposals to reform the fee structure for venture capital managers by delaying managers' bonus payments in an attempt to separate the skilled from the unskilled, but it goes further and explains why and how such a delay can be useful.²

Our result might also provide some support to the proponents of HMOs (Health Maintenance Organizations) who often emphasize their ability to contain costs by implementation of a payment scheme that aligns physicians' incentives with those of the health-care plan. This argument, however, has never been made precise and there remains considerable debate over the sources of cost savings generated by HMOs. The result presented in this paper indicates that long-term contracts for physicians at HMOs are considerably less costly than the payment resulting from the traditional Fee-for-Service system, which can be translated to our setting as a one-period contract. For example, *The New York Times* reported that, for two triple coronary bypass surgeries performed only months apart, George Washington University Hospital received \$28,113 from a traditional insurer, but received only \$10,987 from the HMO Kaiser Permanente (Milt Freudenheim, 1994). The recent call from the Institute of Medicine for government payers to increase payments to health-care providers who deliver high-quality care is one of several signs that, contrary to what is often assumed in the literature, medical practitioners share a strong feeling that moral hazard and adverse selection are important problems in the industry and unless incentives to exert effort are provided, one should expect under-investment even in the provision of medical treatment.

The remainder of the paper is organized as follows. Section 2 is a brief survey of the literature. We present the basic setup in Section 3. In Section 4 we define the notion of an admissible contract as one that provides incentives for agents to exert effort only if their skills match the complexity of the task. The optimal contract-pair is characterized in Section 5 and we conclude in Section 6. Most of the proofs are relegated to the Appendix.

²The promotion system in academia is another case where, instead of using a linear compensation system, promotion is guaranteed only after a series of successful publications.

2 Related Literature

The literature on dynamic agency can be roughly divided into two groups according to how time is treated in the model: continuous- versus discrete-time models. Although studying similar economics problems, the two groups have little in common when it comes to the theory employed. While our paper belongs to the second group to which most of this section is devoted to, we shall start with a model by Fong (2009) which, although in continuous time, is close to ours in spirit and motivation. In Fong's model the principal's objectives are the same as those of the principal in the model studied here, but Fong does not allow for the use of money as an instrument in the contracts. It follows that the only available tool for providing incentives is the flow rate of patients and Fong's first result is that there is no need to consider complicated contracts because an optimal policy takes the form of a stopping rule that specifies if and when to permanently fire a surgeon. The main result is a characterization of the optimal contract-pair that takes the form of scoring rules in which the surgeon's past performance is summarized by a single score and the surgeon is fired if his score falls below a threshold, and is tenured if his score climbs above some other threshold. Contracts for surgeons of different quality levels are different in their sensitivity to successes and failures. Fong's work can be viewed as a rationalization (and refinement) of the report card system. Another continuous-time paper on dynamic agency with moral hazard and adverse selection is Sannikov (2007). In the paper a principal employs an agent of unknown skills where the principal observes no information during the contract period and needs to condition his compensation only on the reports of the agent. In this environment, to prevent manipulation by the agent, the optimal contract requires very specific conditioning of the compensation on the reported information. More precisely, the agent gets a credit line and he is compensated, only if during the whole contract period the balance of the line was above the prespecified cutoff. Also in continuous time, Biais et. al. (2010) analyzed the problem of dynamic moral hazard where an agent with limited liability suppose to exert unobservable effort to prevent the high scale losses to the firm. It is shown that in the optimal contract, the agent is compensated by the fixed per-period flow of payments if and only if no losses occurs for a significantly long period of time. Moreover, sometimes in the optimal contract the principal may downsize the firm in order to make the shirking less favorable action for the agent.

Discrete-time models evolved gradually from dynamic models of moral hazard only to models in which moral hazard as well as adverse selection

problems are present, and from models in which only short-term contracts are offered to those in which the principal can commit to a long-term contract. In an attempt to describe the development along these lines of research we list below only a small sample of these papers, and no attempt is made to provide an exhaustive survey of a very productive field.

One of the first papers on dynamic agency is Rubinstein and Yaari (1983) who considered an infinitely repeated moral hazard problem and demonstrated the existence of a strategy for the principal that yields the first best in an environment in which the principal cannot commit to a strategy that governs the relation. Note however that the infinitely repeated aspect of their problem is crucial in deriving their result, which indeed falls within the realm of the theory of repeated games. In a pioneering paper on career concern and reputation, Holstrom (1982) studied the provision of incentives to exert effort when the agent's ability is unobserved in finitely repeated interactions without output-contingent multi-period contracts.

An important contribution is Holmstrom and Milgrom (1987), who studied a finitely (as well as a continuous-time) repeated moral hazard problem, but, unlike the Rubinstein-Yaari model, and along the lines we are pursuing in our paper, the principal in their model can commit to a long-term strategy that governs the relations in all periods. That is, the principal pays the agent at the end of the last period based on the entire observable history. It is shown that the optimal compensation scheme is a simple linear function of observable events. Similarly, Malcomson and Spinnewyn (1988), Rey and Salanie (1990), and Fudenberg, Holmstrom, and Milgrom (1990) studied the question of when the long-term optimal contract can be replicated by a sequence of short-term (spot) contracts.

Laffont and Tirole (1988) explored a dynamic two-period model of moral hazard *and* adverse selection and identified the ratchet effect that occurs whenever the principal is constrained to offer a short-term contract. That is, the equilibrium is characterized by much pooling in the first period as agents internalize the cost involved in revealing their type. Baron and Besanko's (1984) model of moral hazard and adverse selection is one in which the principal can commit to a long-term strategy but the moral hazard problem is not dynamic. In particular, they study the case of a regulated monopoly that first invests in R&D and then, in future periods, observes privately its marginal cost, which depends stochastically on the level of investment in R&D in period zero. Thus, their model is a one-shot moral hazard problem followed by a multi-period incentive scheme under adverse selection.

Our model incorporates all the incentives problems mentioned above. First, it is a dynamic moral-hazard problem, as the surgeon's choice of effort

in any given period is unobservable. Furthermore, there are also two types of adverse selection problems to overcome: a persistence adverse selection problem due to the unobservability of the surgeon’s quality as determined in period zero, and a dynamic adverse selection since the type of patient, which is different in every period, is observable only by the surgeon.

3 The Model

In what follows and for the sake of clarity only, we refer to the agent as surgeons, and discuss the abstract problem in the context of health economics, where the principal is a care insurer or public official who employs surgeons, whose quality he does not observe, to treat a flow of patients, the severity of whose problems are also the surgeon’s private information.

Basic set-up

Consider a surgeon who is employed by a principal for T periods. In every period $t \in \{1, 2, \dots, T\}$, the surgeon sees one patient and has to decide whether to perform a surgery and if so whether to exert a costly effort $C \in \{0, c\}$. While the probability of a successful operation in period t is positive only if $C = c$, it is also a function of the surgeon’s quality denoted by s , as well as the severity of the problem of the patient who shows up at t , which is denoted by p_t and is referred to as the patient’s “type” at t .

Surgeons are of two quality levels: high and low, denoted by $s \in \{h, l\}$, respectively. Conditional on exerting effort c , a surgeon of type h has a higher probability of a successful operation on a given patient. Similarly, the arriving patient in period t has either a minor or major problem ($p_t \in \{e, d\}$, respectively), and conditional on the surgeon’s quality, the chances of a successful operation are higher when the patient’s problem is minor. We assume that for all $t \in \{1, 2, \dots, T\}$, the patient’s type $p_t \in \{e, d\}$ is independently drawn and the probability of an arrival of type d is q and type e is $(1 - q)$. Finally, the quality of the surgeon is his private information and the patient’s type p_t is revealed only at t and only to the surgeon.

Technology

The probability $\Pi : \{0, c\} \times \{l, h\} \times \{e, d\} \rightarrow [0, 1]$ that a surgery will be successful is given by

$$\Pi(C, s, p_t) = \begin{cases} 0 & \text{if } (C, s, p_t) = (0, s, p_t) \\ \pi_{(h,e)} & \text{if } (C, s, p_t) = (c, h, e) \\ \pi_{(h,d)} & \text{if } (C, s, p_t) = (c, h, d) \\ \pi_{(l,e)} & \text{if } (C, s, p_t) = (c, l, e) \\ \pi_{(l,d)} & \text{if } (C, s, p_t) = (c, l, d) \end{cases}$$

where for $s \in \{h, l\}$ and $p_t \in \{e, d\}$ we have

$$\begin{aligned} (i) \quad & 0 < \pi_{(s, p_t)} < 1 \\ (ii) \quad & \pi_{(h, e)} > \pi_{(h, d)} \text{ and } \pi_{(l, e)} > \pi_{(l, d)} \\ (iii) \quad & \pi_{(h, e)} > \pi_{(l, e)} \text{ and } \pi_{(h, d)} > \pi_{(l, d)}. \end{aligned}$$

Thus, conditional on exerting effort c , the surgeon's probability of success is higher if he is of high quality, for any type of patient; and is higher when the patient's problem is minor, for any type of surgeon. The analysis reveals that the nature of the optimal contract depends on whether $\pi_{(l, e)} > \pi_{(h, d)}$ or $\pi_{(h, d)} > \pi_{(l, e)}$. The bulk of the paper is devoted to the more interesting case $\pi_{(l, e)} \geq \pi_{(h, d)}$, while the treatment of the other case, being very similar, is provided in the Appendix.

Preferences

The surgeon's VNM utility is a function of efforts and payments only. In particular, the utility of a surgeon who exerts effort in k periods and receives a total payment of m is $m - ck$. Thus, the surgeon is assumed to be risk-neutral and to maximize expected payment minus costs. The outside option generates a stream of utilities, which, for simplicity, are normalized to zero per period. Consequently, due to limited liability, negative payments are ruled out.

The principal

The surgeon here is employed by a principal, say, the health authority. If the surgeon's quality is low, i.e., $s = l$, the principal would like him to operate and exert effort only if the patient's problem is minor, $p_t = e$, and otherwise not to operate on him. If, however, the surgeon is a high-quality one, $s = h$, then the principal would like him to exert effort on all types of patients. One possible scenario, leading to these preferences, is the existence of an alternative treatment whose probability of success is higher only when the patient's problem is major, and the surgeon's quality is low.

Conditional on the surgeon providing the right treatment, the principal's objective is to minimize expected payment. Thus, the principal's preferences are lexicographic. First and foremost, he is interested in providing incentives to the surgeon to perform an operation and to exert effort *only* when desirable. As there are many mechanisms that lead to these incentives, the principal is interested in the one that minimizes expected payment.³

³We assume for simplicity that both the principal and the surgeon do not discount the future. The result are qualitatively the same if we assume that they discount future payment at the same rate. To keep the model tractable, we do not assume different time preferences.

The principal can fully commit at time $t = 0$ to any observable history-dependent contract, governing the surgeon's payments. Because the effort C , the surgeon's quality s , and the types of the patients p_t , for $t \in \{1, \dots, T\}$, are not observable by the principal, the only information available to the principal at t is a specification, for every $t' \leq t$, as to whether an operation was conducted, and if so whether it was successful or not.

4 Contracts

Recall that in our setup the principal, in every period t , observes one of three possible outcomes: (i) successful operation, (ii) no operation, and (iii) failed operation, which we denote by $\{1, 0, -1\}$ respectively. A contract thus, specifies for every $t \in \{1, \dots, T\}$ the payment to the surgeon as a function of the observable history up to (and including) t which is a sequence of t elements from $\Psi = \{1, 0, -1\}$ and is denoted by ω_t where Ω_t denotes the set of all possible histories from time zero to t .⁴ Without loss of generality we can assume that all payments are postponed to the last period, T , and define a contract as follows.

Definition 1 *A T -periods contract is a mapping $\tau_T : \Omega_T \rightarrow \mathbb{R}^+$ specifying the payment to the surgeon as a function of the observed history $\omega_T \in \Omega_T$.*

As is typically the case in solving problems of adverse selection, the principal offers a menu of contracts, from which the surgeon chooses the one that is best for him given his quality. Without loss of generality, we can restrict our attention to a mechanism where only two contracts are offered by the principal: τ_T^h to the high-quality surgeon and τ_T^l to the low-quality one.

A two-period contract for a surgeon of type $s \in \{h, l\}$ is depicted below. Note that for every history $\omega_t \in \Omega_t$ we associate a subgame sub_{ω_t} that contains all possible observable histories following ω_t .⁵

Definition 2 *Admissible Contract-pair:* *A pair of contracts (τ_T^h, τ_T^l) is called admissible if it satisfies incentive compatibility (IC), individual rationality (IR), and efficiency (EF) where:*

⁴Following convention, we let \emptyset denotes the history in period zero.

⁵The notion of subgame here, although obvious, is not exactly the one used in game theory.

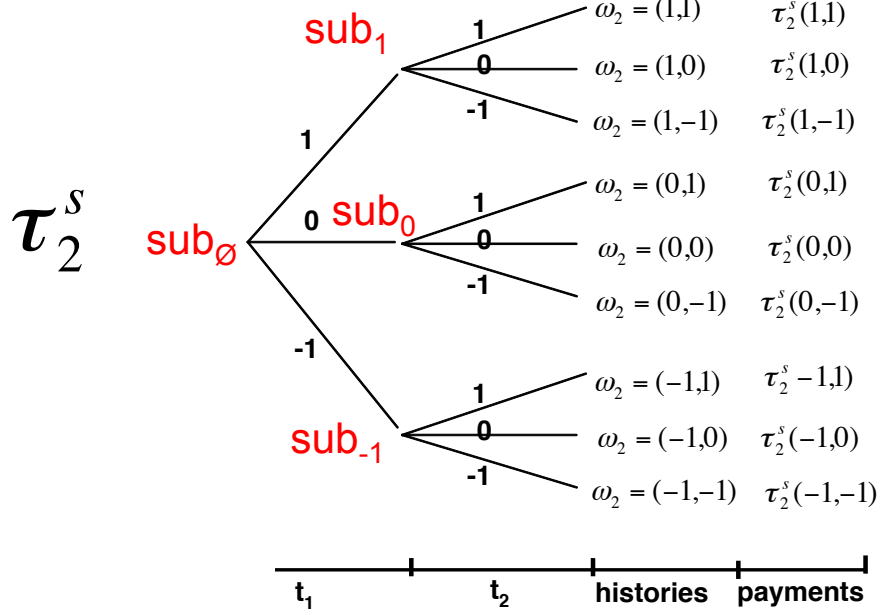


Figure 1: Two-period Contract

IC – a surgeon of quality h prefers the contract τ_T^h to τ_T^l , while the opposite holds for a surgeon of quality l .

IR – the contract τ_T^s yields a non-negative expected payoff to a surgeon of quality $s \in \{h, l\}$ starting after every history $\omega_t \in \Omega_t$ and for all $t \in \{1, 2, \dots, T\}$.

EF – for all $t \in \{1, 2, \dots, T\}$, a surgeon of quality h prefers to conduct an operation and exert effort on all types of patients, while a surgeon of quality l prefers to conduct an operation and exert effort only if at t the patient is of type e .

Remark 1 If (τ_T^h, τ_T^l) is admissible then τ_T^h must entail conducting an operation in every period along the equilibrium path. It follows that if $\tau_T^h(\omega'_T) > 0$ for some ω'_T containing an outcome of zero (no operation), then there exists another contract $\tilde{\tau}_T^h$ in which $\tilde{\tau}_T^h(\omega'_T) = 0$ and $\tilde{\tau}_T^h(\omega_T) = \tau_T^h(\omega_T)$ for all $\omega_T \neq \omega'_T$ such that the new pair $(\tilde{\tau}_T^h, \tau_T^l)$ is admissible and yields, in equilibrium, the same expected payment to the principal. Thus, without loss of generality, we hereafter restrict our attention to contracts for the high-quality surgeon that pay zero whenever the history contains an outcome of zero. That is, if (τ_T^h, τ_T^l) is admissible, then $\tau_T^h(\omega'_T) = 0$ whenever $\{0\} \in \omega'_T$.

Remark 2 Note that if at some t and ω_t the contract provides the surgeon with incentives to exert effort on a given patient's type, then the surgeon will exert effort whenever the arriving type has a higher success probability. It follows that a contract pair (τ_T^h, τ_T^l) satisfies EF if τ_T^h provides the high-quality surgeon with adequate incentives to exert effort whenever a patient with a major problem arrives (i.e., $p_t = d$), while τ_T^l provides the low-quality surgeon with incentives to exert effort only if the patient's problem is minor.

Remark 3 Since the surgeon can always refrain from operating, and since all payments are non-negative, all contracts satisfy IR.

Of all admissible contract-pairs, we are interested in the one that minimizes expected payment. So denote by $m^s(\tau_T^s, \text{sub}_{\omega_t})$ the ex-ante (before observing the patient's type in period $t + 1$) expected payment of τ_T^s to a surgeon of quality s conditional on ω_t and conditional on playing optimally thereafter and let $u^s(\tau_T^s, \text{sub}_{\omega_t})$ denote the ex-ante expected utility of τ_T^s to a surgeon of quality s conditional on ω_t and conditional on playing optimally thereafter. Note that $m^s(\tau_T^s, \text{sub}_{\emptyset})$ and $u^s(\tau_T^s, \text{sub}_{\emptyset})$ are monotonically related in all contracts τ_T^s satisfying EF. This is so because expected costs to a surgeon of quality s are the same in all contracts satisfying EF. In particular, given a T -period contract-pair (τ_T^h, τ_T^l) satisfying EF, it is straight forward to verify that

$$u^h(\tau_T^h, \text{sub}_{\omega_t}) = m^h(\tau_T^h, \text{sub}_{\omega_t}) - c(T - t)$$

and

$$u^l(\tau_T^l, \text{sub}_{\omega_t}) = m^l(\tau_T^l, \text{sub}_{\omega_t}) - c(1 - q)(T - t)$$

where, as defined above, $(1 - q)$ is the probability that the patient's problem is minor, i.e., $p_t = e$.

We are now in a position to define an optimal contract-pair.

Definition 3 An Optimal Contract-pair. A pair of contracts $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is called optimal if it is admissible, and if for every admissible contract-pair (τ_T^h, τ_T^l) we have

$$m^h(\tau_T^h, \text{sub}_{\emptyset}) \geq m^h(\hat{\tau}_T^h, \text{sub}_{\emptyset}) \text{ and } m^l(\tau_T^l, \text{sub}_{\emptyset}) \geq m^l(\hat{\tau}_T^l, \text{sub}_{\emptyset}).$$

Finally, denote by p^s the ex-ante probability of a successful operation by a quality s surgeon when effort is exerted. That is,

$$p^h = q\pi_{(h,d)} + (1 - q)\pi_{(h,e)}$$

and

$$p^l = q\pi_{(l,d)} + (1 - q)\pi_{(l,e)}.$$

5 The Optimal Contract-pair

In this section we maintain the assumption that $\pi_{(l,e)} > \pi_{(h,d)}$ and show that the optimal contract-pair is a separating pair, in the sense that surgeons of different quality sign different contracts. When this assumption does not hold (i.e., $\pi_{(l,e)} < \pi_{(h,d)}$) the unique optimal contract-pair is pooling. Since the analysis of the pooling case is very similar to that of the separating case, it is postponed to Appendix B.

We start by characterizing the set of optimal contracts when the surgeon is known to be a high-quality surgeon, and denote this set by Γ_T^h . We then show that when the surgeon's quality is *unobservable*, the contract offered to the high-quality surgeon belongs to Γ_T^h . Thus, when quality is unobservable, the contract assigned to the high-quality surgeon is the second-best contract and the binding constraint is the incentive constraint on the low-quality surgeon, whose purpose is to ensure that he will prefer the contract assigned to him to the one assigned to the high-quality surgeon.

While the high-quality surgeon is indifferent between all contracts in Γ_T^h (see point 2 below), this is not the case for the low-quality surgeon. The main theorem of this section establishes that the optimal contract for the high-quality surgeon is the contract in Γ_T^h that would minimize the payoff of the low-quality surgeon if he pretended to be a high-quality one and adopted it. In this contract a success in period t is rewarded *only if* it is followed by a success in *every* period following t . This contract, in a way, is the riskiest contract in Γ_T^h , however, and this is crucial, it is exponentially more risky to the low-quality surgeon than it is to the high-quality one. In contrast, the optimal contract to the low-quality surgeon is the contract that pays a fixed amount per successful operation and, makes the low-quality indifferent between the two contracts. It is shown that as T gets larger, the per-success expected payment in the optimal contract-pair approaches the expected amount paid when quality is observable.

5.1 Surgeon's Quality is Known to be High

We now characterize the set of optimal contracts to a surgeon whose quality is known to be high. A contract $\hat{\tau}_T^h$ belongs to Γ_T^h if it satisfies *IR* and *EF* and if there is no other contract τ_T^h that also satisfies *IR* and *EF* and for which expected payment is lower, i.e., $m^h(\tau_T^h, sub_{\emptyset}) < m^h(\hat{\tau}_T^h, sub_{\emptyset})$. Before we proceed and study the properties of Γ_T^h few points are worth mentioning.

1. Note that although the surgeon's quality is observable, there are still problems of moral hazard and adverse selection to solve because the

surgeon's effort and the patient's type are not observable by the principal. Indeed note that if the patient's type is also observable, then a first-best solution can be achieved through a simple contract that promises a payment of $c/\pi_{(h,d)}$ per successful operation on a patient with a major problem ($p_t = d$), and a payment of $c/\pi_{(h,e)}$ per successful operation on a patient with a minor problem ($p_t = e$). Such a contract satisfies *EF* and at the same time brings the surgeon to his *IR* utility. However, when the type of the patient is not observable to the principal and he relies on the surgeon's report of the patients' type, the contract is not incentive-compatible since the surgeon will always report that the patient has a major problem. As a result, when the patient's type is not observable, the optimal contract does leave the surgeon some information rent.

Indeed, if $\hat{\tau}_T^h \in \Gamma_T^h$, then for every history $\omega_{T-1} \in \Omega_{T-1}$, $\hat{\tau}_T^h$ provides incentives for the surgeon to exert effort whenever the patient's problem is major. So consider the following feasible strategy for the surgeon: do not exert effort in all periods $t \in \{1, \dots, T-1\}$ and exert effort in T only if the patient's problem is minor. Note that this strategy guarantees a strictly positive expected payoff since the payment after a sequence of failures is nonnegative and the payment at the last period provides incentives even if $p_T = d$. It allows us to conclude that if $\hat{\tau}_T^h \in \Gamma_T^h$, then $u^h(\tau_T^h, \text{sub}_\emptyset) > 0$.

2. The definition of Γ_T^h implies that expected payment is the same in all contracts in Γ_T^h . Since expected costs are the same in all contracts satisfying *EF* and in particular in all contracts in Γ_T^h , the surgeon is indifferent between all contracts in Γ_T^h .

A 3-periods contract for the high-quality surgeon, where histories containing zeroes are ignored, is described below

The following lemma, proved in Appendix A, lists a few properties that are satisfied by all contracts belonging to Γ_T^h . These properties are then used to characterize the set Γ_T^h of optimal contracts.

Lemma 1 *Properties of Γ_T^h*

1. If $\tau_K^h \in \Gamma_K^h$, then $\exists \tau_{K-1}^h \in \Gamma_{K-1}^h$ s.t. $\forall \omega_{K-1} \in \Omega_{K-1}$, $\tau_K^h(-1, \omega_{K-1}) = \tau_{K-1}^h(\omega_{K-1})$.
2. If $\tau_K^h \in \Gamma_K^h$, then $u^h(\tau_K^h, \text{sub}_1) - u^h(\tau_K^h, \text{sub}_{-1}) = \frac{c}{\pi_{(h,d)}}$.




Figure 2: Three-period Contract for a High-quality Surgeon

3. If $\tau_K^h \in \Gamma_K^h$, then $m_K^h(\tau_K^h, sub_\emptyset) = Kp^h \frac{c}{\pi_{(h,d)}}$.
4. Assume $\tilde{\tau}_T^h$ satisfies IR and EF but $\tilde{\tau}_T^h \notin \Gamma_T^h$. Then, there exists $\tau_T^h \in \Gamma_T^h$ such that for any history $\omega_T \in \Omega_T$, $\tilde{\tau}_T^h(\omega_T) \geq \tau_T^h(\omega_T)$ with strict inequality for at least one history $\omega_T' \in \Omega_T$.

The first property of the lemma refers to the payments restricted to sub_{-1} . In the context of Figure 2 above, it says that if a three-period contract belongs to Γ_3^h then the induced two-period contract in sub_{-1} belongs to Γ_2^h . In other words, a failure in period one is not rewarded and, as a result, from period two on, the surgeon is faces a $K - 1$ -period contract.

Property 2 follows from the fact that effort is not observable and will not be exerted unless incentives are provided. In particular, if the first patient to show up turns out to have a major problem, the surgeon will not exert effort unless the difference in expected payoff between a success and a failure is enough to justify the risk of unsuccessful surgery, an event that occurs with probability $(1 - \pi_{(h,d)})$ if effort is exerted. In a K -period contract, the reward for success in the first period (which occurs with probability $\pi_{(h,d)}$ if the patient's problem is major and effort is exerted) is given by $u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1})$. Thus, exerting effort on a difficult patient is beneficial only if the expected gain is greater than the cost of exerting effort. That is, only if $\pi_{(h,d)}[u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1})] \geq c$. The content of the

second property is that an optimal contract generates, in the first period, the minimal spread between the two subgames, that is needed to provide these incentives.

Property 3 follows from the one-to-one relations between expected utility and expected payment when EF is satisfied, and in particular it implies that Property 2 can be rewritten as

$$m^h \left(\tau_K^h, sub_1 \right) - m^h \left(\tau_K^h, sub_{-1} \right) = \frac{c}{\pi_{(h,d)}}.$$

Of course, the exact same argument holds in every period. That is, in every period incentives to exert effort on a difficult patient must be provided. Thus, for all $t \leq T$ and for every history ω_t the expected reward for success must be at least $\frac{c}{\pi_{(h,d)}}$, and it holds with equality in the first period. Finally, recall that ex ante success occurs with probability $p^h = q\pi_{(h,d)} + (1-q)\pi_{(h,e)}$, and you get the expected payment in a K -period contract specified in Property 3.

The first three properties are employed in the proof of the fourth property, which establishes an important characteristic property of the set Γ_K^h . That is, if a contract is not optimal, then there exists an optimal contract that pays less in *every* possible history. The proof of the following lemma, which is relegated to Appendix A, makes use of the four properties in Lemma 1 to provide a characterization of Γ_T^h and in particular to show that for all T , $\Gamma_T^h \neq \emptyset$.

Lemma 2 *Characterization of Γ_T^h .*

- i. $\tau_1^h \in \Gamma_1^h$ if and only if $\tau_1^h(1) = c/\pi_{(h,d)}$, $\tau_1^h(-1) = 0$, and $\tau_1^h(0) = 0$.
- ii. $\tau_{K+1}^h \in \Gamma_{K+1}^h$ if and only if τ_{K+1}^h can be constructed from contracts in Γ_K^h according to the following procedure:
 - ii.1 The τ_{K+1}^h payments restricted to sub_{-1} are a contract in Γ_K^h .
 - ii.2 The τ_{K+1}^h payments restricted to sub_1 are a contract in Γ_K^h inflated by an expected payment of $c/\pi_{(h,d)}$, which is allocated to the different histories of sub_1 in any way, provided that incentives to exert efforts are not distorted.

Recall that by definition the expected payment is the same in all optimal contracts. This fact together with Lemma 2 yields the following simple corollary and also establishes that the set Γ_T^h is not empty.

Corollary 1 *The set $\Gamma_T^h \neq \emptyset$ and in particular the contract $\hat{\tau}_T^h \in \Gamma_T^h$, where $\hat{\tau}_T^h(\omega_T) = \frac{c}{\pi(h,d)}n(\omega_T)$, and $n(\omega_T)$ is the number of successful operations in ω_T . Thus, a contract is optimal only if it pays in expectation $c/\pi(h,d)$ for every successful operation.*

5.2 Surgeon's Quality is Unobservable

Having characterized the set Γ_T^h we are now ready to study the case where the surgeon's quality is unobservable. Note that now the *IC* constraint must be taken into account since the surgeon will choose the contract that maximizes his expected utility, and not necessarily the one designed for him by the principal. We start by showing that if a contract-pair (τ_T^h, τ_T^l) is optimal, then $\tau_T^h \in \Gamma_T^h$.

Lemma 3 *If (τ_T^h, τ_T^l) is an optimal contract-pair, then $\tau_T^h \in \Gamma_T^h$.*

Proof. Assume by way of contradiction that $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is optimal but $\hat{\tau}_T^h \notin \Gamma_T^h$. Since $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is an optimal contract-pair it is admissible and in particular both contracts satisfy *IR* and *EF*. Hence, Property 4 in Lemma 1 implies that there exists a contract $\tilde{\tau}_T^h \in \Gamma_T^h$ such that for all history $\omega_T \in \Omega_T$, $\hat{\tau}_T^h(\omega_T) \geq \tilde{\tau}_T^h(\omega_T)$ with strict inequality for at least one history. Hence, replacing $\hat{\tau}_T^h$ with $\tilde{\tau}_T^h$ will decrease the expected utility of the low-quality surgeon should he pretend to be a high-quality surgeon by adopting the high-quality surgeon's contract. So consider a contract $\tilde{\tau}_T^l$ that pays $r \geq c/\pi(l,e)$ per success and makes the low-quality surgeon indifferent to the contract $\tilde{\tau}_T^h$. To see that such a contract $\tilde{\tau}_T^l$ always exists, it is enough to note that (i) a low-quality surgeon can always adopt the contract $\tilde{\tau}_T^h$ and then exert no effort to obtain a non-negative utility, and (ii) a contract that pays $c/\pi(l,e)$ per success satisfies *EF* and yields zero expected utility to the low-quality surgeon.

We next argue that $(\tilde{\tau}_T^h, \tilde{\tau}_T^l)$ is admissible. That is, (i) $r \leq c/\pi(l,d)$, and (ii) the high-quality surgeon prefers the contract $\tilde{\tau}_T^h$ to $\tilde{\tau}_T^l$. Note, however, that $r \leq c/\pi(h,d)$ is sufficient for (i) and (ii). This is because (i) follows from $c/\pi(h,d) < c/\pi(l,d)$ and (ii) from the fact that a contract that pays $c/\pi(h,d)$ per success belongs to Γ_T^h and the high-quality surgeon is indifferent between all contracts in Γ_T^h . Therefore, if $r \leq c/\pi(h,d)$, the contract-pair $(\tilde{\tau}_T^h, \tilde{\tau}_T^l)$ is admissible and generates a lower expected payment to both agents than the pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$, which is a contradiction.

So assume that $r > c/\pi(h,d)$ and observe that a contract-pair that pays $c/\pi(h,d)$ per success to both types of surgeons is admissible and provides

both types of surgeon with lower expected utility than $(\hat{\tau}_T^h, \hat{\tau}_T^l)$, which is again in contradiction to the assumed optimality of the original pair. We conclude that if a contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is optimal, then $\hat{\tau}_T^h \in \Gamma_T^h$. ■

Note that while different contracts in Γ_T^h generate the same expected utility for the high-quality surgeon, they generate different expected utilities for the low-quality one, if he chooses to adopt them. It thus follows from Lemma 3 that a contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is optimal if the contract $\hat{\tau}_T^h$ is the one that, more than any other contract in Γ_T^h , minimizes the expected utility of the low-quality surgeon. In other words, from the high-quality surgeon's point of view, the set Γ_T^h consists of different lotteries between which he is indifferent. However, from the point of view of the low-quality surgeon, these are different lotteries and $\hat{\tau}_T^h$, to be defined below, is the riskiest among them. That is, although the set Γ_T^h for $T > 1$ is not a singleton and contains many contracts, asymmetric information about the surgeon's type pins down the contract that the designer offers to the high-quality surgeon. The theorem also establishes that as $T \rightarrow \infty$ the optimal contract-pair converges to the second-best pair, that is, the contract-pair that is offered when the surgeon's quality is observable.

Prior to presenting the formal statement of the theorem, we describe its content in the simplest possible dynamic context, i.e., when $T = 2$. In this case the theorem postulates that if the high-quality surgeon performs an operation in period one and succeeds, he is compensated for this only if he also performs a successful operation in period two. The compensation in the event that there are two successes in a row, must be high enough to cover the extra risk involved in exerting effort in period one. Specifically;

$$\hat{\tau}_2^h = \begin{cases} \frac{c}{\pi(h,d)} \frac{1+p^h}{p^h} & \text{if } \omega_2 = (1, 1) \\ \frac{c}{\pi(h,d)} & \text{if } \omega_2 = (0, 1) \\ 0 & \text{if } \omega_2 = (1, 0) \\ 0 & \text{if } \omega_2 = (0, 0) \end{cases}$$

Note that while the high-quality surgeon (being risk-neutral) is indifferent between this contract and the one that pays $\frac{c}{\pi(h,d)}$ per success, the low-quality surgeon strictly prefers the latter.

The following theorem characterizes the optimal contract-pair for T periods, while making use of the following definitions:

- (i) Define $A(k)$ recursively by letting $A(0) = 0$ and $A(k) = A(k-1) + \frac{1}{(p^h)^{k-1}}$.

(ii) Let $\tilde{k}(\omega_T)$ be the length of the longest uninterrupted sequence of successful operations in ω_T , starting from period T backward.

Theorem 1 *An optimal contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ has the following properties:*

1. *If ω_T contains an outcome of 0, then $\hat{\tau}_T^h(\omega_T) = 0$. Otherwise, if $\tilde{k}(\omega_T) = k$, then $\hat{\tau}_T^h(\omega_T) = \frac{c}{\pi(h,d)}A(k)$.*
2. *There exists a constant r , such that $\hat{\tau}_T^l(\omega_T) = rn(\omega_T)$, where $n(\omega_T)$ is the number of successful operations in ω_T . Moreover, $\lim_{T \rightarrow \infty} r = \frac{c}{\pi(l,e)}$.*

Proof. We start the proof by showing that the contract $\hat{\tau}_T^h$ described in the theorem minimizes the expected utility of the low-quality surgeon in all the contracts that belong to Γ_T^h . The formal argument follows from Claim 1, setting $\tilde{u} = 0$; the proof of the claim is relegated to Appendix A. First, note that if $\{0\} \in \omega_T$ and $\hat{\tau}_T^h(\omega_T) > 0$, then decreasing this payment will not affect the expected utility of the high-quality surgeon and will decrease (or will not affect) the expected utility of the low-quality surgeon from this contract. Therefore, without loss of generality we can restrict our attention to contracts in Γ_T^h where the payments after histories containing $\{0\}$ are zero. ■

Claim 1 *Let \bar{u}_T^h denotes the expected utility of the high-quality surgeon from any contract in Γ_T^h . Assume that, the principal is asked to provide the high-quality surgeon with an additional expected utility of $\tilde{u} \geq 0$ (in excess of \bar{u}_T^h), but in a way that preserves incentives to exert effort, while at the same time minimizing the expected utility of the low-quality surgeon should he adopt this contract. This is achieved by amending the contract $\hat{\tau}_T^h$ described in Theorem 1-1 and adding a payment of $\tilde{u}/(p^h)^T$ after a sequence of T successful operations.*

Proof. (continued) We proceed by constructing the contract $\hat{\tau}_T^l$ described in Theorem 1. The constant r in $\hat{\tau}_T^l$ is chosen so that the low-quality surgeon is indifferent between choosing $\hat{\tau}_T^l$ and $\hat{\tau}_T^h$. Since the expected utility of the low-quality surgeon from $\hat{\tau}_T^h$ is positive (one possible strategy for him is to invest only in period T and only if the patient's problem is minor), we have $r \geq c/\pi(l,e)$. Moreover, since a contract that pays $c/\pi(h,d)$ per success belongs to Γ_T^h , Claim 1 implies that the utility of the low-quality surgeon from this contract is higher than in $\hat{\tau}_T^h$, which in turn implies that

$r \leq c/\pi_{(h,d)}$. Therefore, since $c/\pi_{(l,e)} \leq r \leq c/\pi_{(h,d)}$, the contract $\hat{\tau}_T^l$ generates the right incentives for the low-quality surgeon.

We complete the proof by showing the limit result. Note that to establish this result it is sufficient to show that the expected utility of the low-quality surgeon from the contract $\hat{\tau}_T^h$ stays bounded as $T \rightarrow \infty$. To show the last statement, it is enough to show, that as $T \rightarrow \infty$, the low-quality surgeon who adopts $\hat{\tau}_T^h$ exerts efforts in a finite number of (last) periods. Denote by K the first period at which the surgeon begins exerting effort conditional on the patient having a minor problem. It is sufficient to show that as $T \rightarrow \infty$, the optimal strategy for the low-quality surgeon who adopts $\hat{\tau}_T^h$, is to start exerting effort only if $t \geq T - K$, where K remains bounded even if $T \rightarrow \infty$. Assume by way of contradiction that this is not the case and instead $K \rightarrow \infty$ as $T \rightarrow \infty$. Observe however that whenever the surgeon exerts effort, it affects his utility only if it is followed by an uninterrupted sequence of successes. That is, if the surgeon succeeds in all remaining K periods (starting from period $(T - K)$ till the end of the contracting period) he will, according to $\hat{\tau}_T^h$, receive a payment of

$$\frac{c}{\pi_{(h,d)}} A(K) = \frac{c}{\pi_{(h,d)}} \frac{\left(\frac{1}{p^h}\right)^K - 1}{\frac{1}{p^h} - 1}$$

and zero otherwise. Recall that for any strategy of the low-quality surgeon, the probability of success in K operations is less than or equal to $(p^l)^K$. Since $p^l < p^h$, the expected utility of the low-quality surgeon from any strategy in which he starts exerting effort in period $T - K$ is bounded by

$$-c + \frac{c}{\pi_{(h,d)}} \frac{\frac{\pi_{(l,e)}}{p^h} \left(\frac{p^l}{p^h}\right)^{K-1} - (p^l)^{K-1}}{\frac{1}{p^h} - 1}$$

Since

$$\lim_{K \rightarrow \infty} \frac{c}{\pi_{(h,d)}} \frac{\frac{\pi_{(l,e)}}{p^h} \left(\frac{p^l}{p^h}\right)^{K-1} - (p^l)^{K-1}}{\frac{1}{p^h} - 1} = 0,$$

we are done. ■

Observe that when $T = 1$ (the static problem) the optimal contract-pair is actually pooling. That is,

$$\hat{\tau}_1^h(\omega) = \hat{\tau}_1^l(\omega) = \begin{cases} \frac{c}{\pi_{(h,d)}} & \text{if } \omega = \{1\} \\ 0 & \text{otherwise} \end{cases}.$$

When $T > 1$ the contract-pair in which this pooling payment scheme is repeated, satisfies *IC* and *EF*. Theorem 1, however, shows that the dynamic structure alleviates the screening problem of the principal and allows us to decrease the low-quality surgeon's information rents. The optimal contract uses the fact that some histories are more likely to occur when the contract is chosen by the high-quality surgeon, rather than the low-quality one, for any choice of effort. Increasing the payments assigned to these histories at the expense of the payments assigned to the other histories makes this contract much less attractive to the low-quality surgeon.

Non-uniqueness

The optimal contract for the high-quality surgeon described in Theorem 1, where he is compensated for success in period t only if he succeeds in all periods following period t , looks rather extreme, especially when T is large. Not only does this contract become very risky, it also has the unpleasant feature that no payment is guaranteed until T is reached. Of course, when the agent and the principal are risk-neutral and both discount the future at the same rate, these features are irrelevant. Yet, it is comforting to note that for T large enough, the contract described in Theorem 1 is not the unique optimal contract and there are others in which these two features are relaxed dramatically. The following corollary presents such an optimal contract when only during the last stages of the contract the extreme form described above is used. In particular, we construct a contract for the high-quality surgeon in which during the first $T/2$ periods, he is compensated for a success in period t if following that period the success rate is at least $\alpha \in (0.5, 1]$, while during the last $T/2$ periods the contract is as described in Theorem 1⁶. We start with a definition.

Definition 4 *A proportional contract $\tau_T^h(\alpha)$ is a contract that pays $\tau_T^h(\alpha, \omega_T)$ after a history ω_T , where $\tau_T^h(\alpha, \omega_T)$ consists of two parts. The first part defines the part of the payment that is due to successes in the first half of the contract, i.e., up to period $T/2$, while the second part defines the part of the payment that is due to successes in periods $T/2$ onward. In particular,*

$$\tau_T^h(\alpha, \omega_T) = \sum_{m=1}^{\lceil T/2 \rceil} \mathbf{1}_{\{o(m)=1, n(T)-n(m) \geq \alpha(T-m-1)\}} B(m) + \frac{c}{\pi(h,d)} A(\tilde{K})$$

where the $\mathbf{1}_{\{D\}}$ in the first part is just an index function, $o(m)$ is the outcome of period m , $n(t)$ is the number of successful surgeries from the beginning of

⁶Of course, the choice of $T/2$ is arbitrary and as T gets larger; the part in which the original contract is used can be reduced further.

the contract until period t , and

$$B(m) = \frac{c}{\pi(h,d)} \frac{1}{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^h)^j (1-p^h)^{T-m-1-j}}.$$

where

$$\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^h)^j (1-p^h)^{T-m-1-j}$$

is the ex ante probability of a high-quality surgeon to obtain a success rate of at least α in all future periods, conditional on exerting effort in all $T-m-1$ remaining periods. Finally, note that the second part of $\tau_T^h(\alpha, \omega_T)$ is just the optimal contract defined in Theorem 1 restricted to periods $T/2$ onward, where $\tilde{K} = \min\{K, T/2\}$ and K is the longest uninterrupted sequence of successes in ω_T starting from T backward.

Note that this contract mimics the one defined in Theorem 1 for $\alpha = 1$, while for $\alpha = 0$, this contract pays $\frac{c}{\pi(h,d)}$ per success in every period during the first $T/2$ periods, regardless of the outcomes in other periods.

Corollary 2 For all $\alpha \in (0.5, 1)$ for which

$$(p^l)^\alpha < (1-\alpha)^{1-\alpha} \alpha^\alpha p^h \quad (1)$$

there exists $\hat{T}(\alpha)$ and a constant r such that, for all $T > \hat{T}(\alpha)$, the contract-pair $(\tau_T^h(\alpha), \tau_T^l)$, where $\tau_T^h(\alpha)$ is the proportional contract described in the definition and $\tau_T^l(\omega_T) = rn(\omega_T)$, is an optimal contract-pair.

The proof is provided at the end of Appendix A. Notice that since $(1-\alpha)^{1-\alpha} \alpha^\alpha$ is a monotone increasing function with $\lim_{\alpha \rightarrow 1} (1-\alpha)^{1-\alpha} \alpha^\alpha = 1$ for any p^l and p^h such that $p^h > p^l$ there exists $\alpha^* \in (1/2, 1)$ such that for any $\alpha \in (\alpha^*, 1)$ the inequality (1) holds.

6 Conclusion and Extensions

Risk aversion

The nature of the optimal contract of the high-quality agent described above, where a success occurs in, say, period one, is rewarded only if it is

followed by an uninterrupted sequence of T successes. This exact form is of course an artifact of the assumption that the agent and the principal are risk-neutral. Of course, the same forces are at play when this assumption is relaxed, but now the principal faces a trade-off. Making the contract of the high-quality agent riskier enables the principal to lower the expected payment promised to the low-quality agent, but it comes at the cost of increasing the expected payment offered to the high-quality agent in order to compensate for the extra risk. The exact characterization will now depend on the agent's degree of risk aversion as well as the principal's priors of the likelihood that the agent is of high quality. The assumed risk neutrality allows us to obtain a clean characterization, which highlights the important forces that are at play.

Learning

It is worth noting that the model presented here can accommodate some form of learning and, in particular, it conveniently admits the possibility of a bad agent improving to become a good one by learning. Assume for example that the low-quality agent can invest some amount e per period and undergo additional training programs. Let us assume that an agent is enrolled in a training program for κ periods, that by the end of the program he becomes a high-quality agent with probability $\beta(\kappa)$, and the outcome of the program is the agent's private information. Finally, it is reasonable to assume that $\beta(\kappa)$ is increasing and concave in κ . Given the assumed preferences of our principal, it is in his interest to allow an agent who has become a high-quality agent to switch at any $t \leq T$ from the contract he signed to one that is tailored to his new quality level. Indeed, if we allow an agent who in period $t = 0$ signed a contract designed for a low-quality agent to switch at any $t \leq T$ to a high-quality contract for the remaining $T - t$ periods, then incentive compatibility conditions are not violated. Of course the longer the program is, the shorter is the new contract, and the lower is the payoff from the course. So the optimal κ , the length of the training program, is determined optimally by the cost per period e , the function $\beta(\cdot)$, and T . It is easy to see that in this set up, it might be in the planner's interest to affect this learning process by, say, subsidizing the training.

Pooling

All along we maintained the assumption that $\pi_{(l,e)} > \pi_{(h,d)}$, and left the rather similar analysis of the case where $\pi_{(l,e)} < \pi_{(h,d)}$ to be dealt with in Appendix B (hereafter cases (i) and (ii) respectively). It is, however, worth describing the main result of case (ii) and providing some intuition for the sharp differences between the two cases and in particular for the fact that in case (ii) the optimal contract-pair is pooling, in the sense that regardless

of the surgeon's type, he is paid a fixed amount $c/\pi_{(h,d)}$ per success, as is shown in Theorem 2 in Appendix B. Recall that in case (i) the contract that is offered to the high-quality surgeon is the one that is offered to him when his quality is observable, and it is the low-quality who enjoys some information rent (which converges to zero as T gets larger). As we establish in Appendix B, case (ii) is different. First, it is the high-quality surgeon who enjoys the information rent, and second the repeated nature of the relation is not helpful.

To obtain some insights into the differences between the two cases, assume first that in every $t \in \{1, \dots, T\}$ the principal is constrained to propose a short-term one-period contract only. It is easy to see that in both cases the only contract-pair that satisfies EF and IC is a pooling one in which regardless of the surgeon's quality he is paid a fixed amount per success: $c/\pi_{(h,d)}$ in case (i), and $c/\pi_{(l,e)}$ in case (ii). Adopting the terminology developed above for long-term contracts, and letting $n(\omega_T)$ denotes the number of successful operations in ω_T , these contract can be written as

$$\text{Case (i): } \bar{\tau}_T^h(\omega_T) = \bar{\tau}_T^l(\omega_T) = \frac{c \cdot n(\omega_T)}{\pi_{(h,d)}}$$

and

$$\text{Case (ii): } \bar{\tau}_T^h(\omega_T) = \tilde{\tau}_T^l(\omega_T) = \frac{c \cdot n(\omega_T)}{\pi_{(l,e)}}.$$

Note that $\bar{\tau}_T^h(\omega_T) \in \Gamma_T^h$, and $\tilde{\tau}_T^l(\omega_T) \in \Gamma_T^l$, which implies that in case (i) the expected utility of the high-quality surgeon is at its lower bound (at its level when his quality is observable), while the expected utility of the low-quality surgeon is above its lower bound. The reverse, however, is true in case (ii), where the expected utility of the low-quality surgeon is at its lower bound.

As we show in the analysis of case (i) above, the important effect of long-term contracts is the availability of other contracts in Γ_T^h which, from the low-quality surgeon's point of view, are worse than $\bar{\tau}_T^h(\omega_T)$. The optimal contract-pair exploits this by assigning the high-quality surgeon the contract in Γ_T^h that is the least attractive to the low-quality surgeon. This enables the principal to then assign to the low-quality surgeon a contract that yields a lower expected payment than the repeated short-term contract $\bar{\tau}_T^l(\omega_T)$. In case (ii) it is the high-quality surgeon who is receiving a level of expected utility above his lower bound. But unlike in case (i) where the short-term contract $\bar{\tau}_T^h(\omega_T)$ was, to the low-quality surgeon, the best in Γ_T^h , now the short-term contract $\tilde{\tau}_T^l(\omega_T)$ is the worst in Γ_T^l to the high-quality surgeon.

It follows that in case (ii) the short-term contract is the best the principal can achieve when the low-quality surgeon is already at his IR , because any other contract τ_l^T that satisfies EF would yield the high-quality surgeon an even higher expected utility.

7 Appendix A: Proofs for the Separating Contract Case

Proof of Lemma 1:

Property 1: Assume that this property is false. Since $\tau_K^h \in \Gamma_K^h$, τ_K^h provides sufficient incentives in all subgames, and in particular in sub_{-1} (the subgame following a failure in the first period). Consider replacing τ_K^h with $\tilde{\tau}_K^h$, where $\tilde{\tau}_K^h$ is obtained by amending the contract τ_K^h and replacing the payments in all histories that belong to sub_{-1} , adopting instead the payments in one of the optimal $K - 1$ -period contracts in Γ_{K-1}^h . That is, $\tilde{\tau}_K^h(-1, \omega_{K-1}) = \tilde{\tau}_{K-1}^h(\omega_{K-1})$. Clearly, the proposed change does not affect incentives in sub_1 . Also, because an optimal $K - 1$ -period contract provides incentives in the $K - 1$ -period problem, incentives are provided in sub_{-1} .

Since the new payment scheme in sub_{-1} , is a contract in Γ_{K-1}^h , it minimizes expected payment in all schemes that provide incentives. That is,

$$m^h(\tilde{\tau}_K^h, sub_{-1}) < m^h(\tau_K^h, sub_{-1}) \quad (2)$$

because otherwise the τ_K payments restricted to sub_{-1} is a contract from Γ_{K-1}^h . Since incentives are provided by τ_K^h to exert effort in period one on a patient with a major problem, it must be the case that

$$u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1}) \geq \frac{c}{\pi_{(h,d)}}$$

and in particular

$$m^h(\tau_K^h, sub_1) - m^h(\tau_K^h, sub_{-1}) \geq \frac{c}{\pi_{(h,d)}}. \quad (3)$$

This together with (2) implies that

$$u^h(\tau_K^h, sub_1) - u^h(\tilde{\tau}_K^h, sub_{-1}) > \frac{c}{\pi_{(h,d)}}$$

which guarantees that incentives to exert effort in period one are preserved and in general incentives are provided in the revised K -period contract.

Finally, note that since this revision decreases the expected utility of the surgeon after failure in the first period and keeps the expected utility after success in the first period, it decreases the expected payment, which is in contradiction to the claimed optimality of the original contract. ■

Property 2: Assume by way of contradiction that $\tau_K^h \in \Gamma_K^h$ but

$$u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1}) \neq \frac{c}{\pi(h,d)}$$

and recall that since an optimal contract provides incentives to exert effort, it must be the case that

$$u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1}) > \frac{c}{\pi(h,d)},$$

Therefore, let us revise τ_K^h to $\tilde{\tau}_K^h$ so that $\tilde{\tau}_K^h(1, \omega_{K-1}) = \tau_K^h(-1, \omega_{K-1}) + \frac{c}{\pi(h,d)}$. Note that incentives to exert efforts in $\tilde{\tau}_K^h$ are kept and that $u^h(\tau_K^h, sub_1)$ is now decreased to $u^h(\tau_K^h, sub_{-1}) + \frac{c}{\pi(h,d)}$ so that expected payment is decreased, which is in contradiction to τ_K^h being optimal. ■

Property 3: The simple proof is done by induction. Observe first that for $T = 1$ we have $\tau_1^h(1) = \frac{c}{\pi(h,d)}$ and $\tau_1^h(-1) = 0$, which implies that $m^h(\tau_1^h, sub_\emptyset) = p^h \frac{c}{\pi(h,d)}$. Next assume that if $\tau_{K-1}^h \in \Gamma_{K-1}^h$ then $m^h(\tau_{K-1}^h, sub_\emptyset) = (K-1)p^h \frac{c}{\pi(h,d)}$. From Properties 1 and 2 it follows that

$$\begin{aligned} m^h(\tau_K^h, sub_\emptyset) &= (1-p^h)(K-1)p^h \frac{c}{\pi(h,d)} + p^h[(K-1)p^h \frac{c}{\pi(h,d)} + \frac{c}{\pi(h,d)}] = \\ &= Kp^h \frac{c}{\pi(h,d)}, \end{aligned}$$

which is the desired final step of the proof. ■

Note that Properties 1, 2, and 3, together imply that in any optimal contract we have

$$m^h(\tau_K^h, sub_1) = (K-1)p^h c / \pi(h,d) + c / \pi(h,d).$$

Denote by \bar{u}_T^h the expected utility of the high-quality surgeon from any contract in Γ_T^h . That is,

$$\bar{u}_T^h = Tp^h \frac{c}{\pi(h,d)} - Tc.$$

Property 4: This property is an immediate consequence of the following claim:

Claim 2 *If a T -period contract τ_T^h satisfies EF and IR and generates expected utility $u > \bar{u}_T^h$, then for any $\tilde{u} \in [\bar{u}_T^h, u)$ there exists another T -period contract $\tilde{\tau}_T^h$ that satisfies EF and IR and generates an expected utility of \tilde{u} and for all $\omega_T \in \Omega_T$, $\tilde{\tau}_T^h(\omega_T) \leq \tau_T^h(\omega_T)$ with at least one strict inequality.*

Proof. The proof is done by inducting on the contract's length, T . Assume that $T = 1$ and observe that since τ_1^h satisfies efficiency, we have

$$\tau_1^h(1) - \tau_1^h(-1) \geq \frac{c}{\pi(h,d)}.$$

Moreover,

$$p^h \tau_1^h(1) + (1 - p^h) \tau_1^h(-1) - c = u.$$

Consider two cases.

Case 1. $\tau_1^h(-1) \geq u - \tilde{u}$. In this case, we set $\tilde{\tau}_1^h(\omega_1) = \tau_1^h(\omega_1) - (u - \tilde{u})$ for $\omega_1 \in \{1, -1\}$. It can be easily verified that the new contract satisfies EF and IR and generates an expected utility of \tilde{u} , and for any $\omega_1 \in \{1, -1\}$ holds $\tilde{\tau}_1^h(\omega_1) < \tau_1^h(\omega_1)$.

Case 2. $\tau_1^h(-1) < u - \tilde{u}$. In this case set $\tilde{\tau}_1^h(-1) = 0$ and $\tilde{\tau}_1^h(1) = \frac{\tilde{u} + c}{p^h} < \frac{u + c - (1 - p^h)\tau_1^h(-1)}{p^h} = \tau_1^h(1)$, where the inequality follows from the fact that in this case $u - \tilde{u} > \tau_1^h(-1)$. Since $\tilde{u} \geq \bar{u}_1^h$, incentives are preserved. Moreover, the contract $\tilde{\tau}_1^h$ generates an expected utility of \tilde{u} and for any $\omega_1 \in \{1, -1\}$ we have $\tilde{\tau}_1^h(\omega_1) \leq \tau_1^h(\omega_1)$. This complete the proof for $T = 1$.

Having established the claim for $T = 1$, we proceed by assuming the statement holds for $T = K - 1$ periods and show that it holds for $T = K$ periods. Assume that there exists a K -period contract τ_K^h for which $u^h(\tau_K^h, \text{sub}_\emptyset) > \bar{u}_K^h$. As in the case of $T = 1$, we consider two cases.

Case 1. $u^h(\tau_K^h, \text{sub}_{-1}) - \bar{u}_{K-1}^h \geq u - \tilde{u}$. In this case consider two $K - 1$ -period contracts that satisfy EF and IR $\tau_{K-1,-1}^h$ and $\tau_{K-1,1}^h$ such that $u^h(\tau_{K-1,-1}^h, \text{sub}_\emptyset) = u^h(\tau_K^h, \text{sub}_{-1}) - (u - \tilde{u})$ and for which we have $\tau_{K-1,-1}^h(-1, \omega_{K-1}) \geq \tau_{K-1,-1}^h(\omega_{K-1})$ (since $u^h(\tau_K^h, \text{sub}_{-1}) - (u - \tilde{u}) \geq \bar{u}_{K-1}^h$, the induction argument guarantees the existence of such a contract) and $u^h(\tau_{K-1,1}^h, \text{sub}_\emptyset) = u^h(\tau_K^h, \text{sub}_1) - (u - \tilde{u})$ and for which we have $\tau_{K-1,1}^h(1, \omega_{K-1}) \geq \tau_{K-1,1}^h(\omega_{K-1})$ (since $u^h(\tau_K^h, \text{sub}_1) - (u - \tilde{u}) \geq \bar{u}_{K-1}^h$, the induction argument guarantees the existence of such a contract). Construct a contract $\tilde{\tau}_K^h$, such that $\tilde{\tau}_K^h(1, \omega_{K-1}) = \tau_{K-1,1}^h(\omega_{K-1})$ and $\tilde{\tau}_K^h(-1, \omega_{K-1}) = \tau_{K-1,-1}^h(\omega_{K-1})$ for any $\omega_{K-1} \in \Omega_{K-1}$. First, notice that by construction, the incentives in any subgame after the first period are guaranteed. Second,

since $u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1}) = u^h(\tilde{\tau}_K^h, sub_1) - u^h(\tilde{\tau}_K^h, sub_{-1})$, the first-period incentives are preserved. The expected utility of the high-quality surgeon is given by $p^h u^h(\tilde{\tau}_K^h, sub_1) + (1 - p^h) u^h(\tilde{\tau}_K^h, sub_{-1}) - c = \tilde{u}$. Finally, by construction, for all $\omega_K \in \Omega_K$ we have $\tilde{\tau}_K^h(\omega_K) \leq \tau_K^h(\omega_K)$, where the inequality is strict for at least one $\omega_K \in \Omega_K$.

Case 2. $u^h(\tau_K^h, sub_{-1}) - \bar{u}_{K-1}^h < u - \tilde{u}$. Consider two $K - 1$ -period contracts that satisfy *EF* and *IR* $\tau_{K-1,-1}^h$ and $\tau_{K-1,1}^h$ such that $u^h(\tau_{K-1,-1}^h, sub_\emptyset) = \bar{u}_{K-1}^h$ and $u^h(\tau_{K-1,1}^h, sub_\emptyset) = \frac{\tilde{u} + c - (1 - p^h)\bar{u}_{K-1}^h}{p^h} < \frac{u + c - (1 - p^h)u^h(\tau_K^h, sub_{-1})}{p^h} = u^h(\tau_K^h, sub_1)$ where the inequality follows from the fact that in this case $u^h(\tau_K^h, sub_{-1}) - \bar{u}_{K-1}^h < u - \tilde{u}$. Since

$$p^h u^h(\tau_{K-1,1}^h, sub_\emptyset) + (1 - p^h) u^h(\tau_{K-1,-1}^h, sub_\emptyset) = \tilde{u} + c > \bar{u}_{K-1}^h + c$$

and $u^h(\tau_{K-1,-1}^h, sub_\emptyset) = \bar{u}_{K-1}^h$ we get that $u^h(\tau_{K-1,1}^h, sub_\emptyset) > \bar{u}_{K-1}^h$. The induction argument guarantees the existence of the contracts with the required properties. As in the previous case we construct a contract $\tilde{\tau}_K^h$ from two $K - 1$ -period contracts, $\tau_{K-1,-1}^h$ and $\tau_{K-1,1}^h$, such that $\tilde{\tau}_K^h(1, \omega_{K-1}) = \tau_{K-1,1}^h(\omega_{K-1})$ and $\tilde{\tau}_K^h(-1, \omega_{K-1}) = \tau_{K-1,-1}^h(\omega_{K-1})$ for any $\omega_{K-1} \in \Omega_{K-1}$. The rest of the proof is similar to the proof of Case 1. ■

Proof of Lemma 2:

Consider first the set Γ_1^h of one-period optimal contracts. Because effort is not verifiable, incentives must be provided to induce effort-exerting even when $p_1 = d$, where the probability of success is low. It follows that incentives to exert effort on all types of patients are provided if and only if $\tau_1^h(1) - \tau_1^h(-1) \geq c/\pi(h,d)$. We conclude that Γ_1^h is a singleton and $\tau_1^h \in \Gamma_1^h$ if and only if $\tau_1^h(1) = c/\pi(h,d)$ and $\tau_1^h(-1) = 0$, which establishes (i) in the statement of the lemma. To complete the proof, note that Property 1 in Lemma 1 shows (ii.1) and to establish (ii.2) it is enough to show that for every $\tau_K^h \in \Gamma_K^h$ there exists a contract $\tilde{\tau}_{K-1}^h \in \Gamma_{K-1}^h$ such that for any history ω_{K-1} we have $\tilde{\tau}_{K-1}^h(\omega_{K-1}) \leq \tau_K^h(1, \omega_{K-1})$. This, however, follows from Property 4 in Lemma 1.

Proof of Claim 1. The proof is done by induction on T , the length of the contract. For $T = 1$, the statement holds trivially. We assume then that the statement holds for $T = K$ and next prove it for $T = K + 1$. Denote by $\tilde{\tau}_{K+1}^h$ a contract that yields a utility of $\bar{u}_{K+1}^h + \tilde{u}$ to the high-quality surgeon and the induced contract on sub_1 and sub_{-1} by $\tilde{\tau}_{K+1}^h$ are as described in point 1 of Theorem 1 amended by some non-negative extra payments μ_1

in sub_1 and μ_{-1} in sub_{-1} , that are paid after a history of K uninterrupted successes. First note that it is always possible to find μ_1 and μ_{-1} such that: (i) $u^h(\tilde{\tau}_{K+1}^h, sub_\emptyset) = \bar{u}_{K+1}^h + \tilde{u}$, (ii) the incentives for the high-quality surgeon are preserved (for example by choosing $\mu_{-1} = 0$), and (iii) by the induction argument $\tilde{\tau}_{K+1}^h$ minimizes the expected utility of the low-quality surgeon in each of these subgames in all contracts that generate an expected utility of $\bar{u}_K^h + (p^h)^K \mu_1$ and $\bar{u}_K^h + (p^h)^K \mu_{-1}$, respectively.

It is left for us to show that $\mu_{-1} = 0$. Assume by way of contradiction that $\mu_{-1} > 0$ and consider decreasing μ_{-1} (the payment after a failure following a sequence of K successes) by $\varepsilon > 0$ and increasing μ_1 (the payment after a sequence of $K + 1$ successes) by $\varepsilon \frac{p^h}{1-p^h}$. Note that this change does not affect the expected utility of the high-quality surgeon and preserves his incentives. Moreover, this change decreases the expected utility of the low-quality surgeon because for any strategy of the low-quality surgeon, his utility decreases with ε and hence, the same is true of the strategy that maximizes his utility. ■

Proof of Corollary 2. We first show that the contract described in Definition 4 is an optimal contract for the high-quality surgeon. To see this, observe that in this contract the expected compensation per success in each period is $\frac{c}{\pi(h,d)}$. It follows that this contract provides the high-quality surgeon with the same expected utility as the original contract described in Theorem 1 and it generates the efficient incentives. We next show that there exists a threshold $\hat{T}(\alpha)$ so that if $T > \hat{T}(\alpha)$, this contract yields the low-quality surgeon the same expected utility as the one for the high-quality surgeon, described in Theorem 1. To show this, it is enough to establish that if the low-quality surgeon adopts this contract his best strategy then is to exert no effort during the first $T/2$ periods. To show this, we now show that for any period $m \in \{1, \dots, T/2\}$ if the low-quality surgeon does not succeed in all $t < m$, his expected utility is higher if he does not exert effort in period m . Since his probability of success in any period is bounded by p^l , the change in his expected utility if he exerts effort at period $m \in \{1, \dots, T/2\}$ is bounded by

$$-c + \pi_{(l,e)} \frac{c}{\pi(h,d)} \frac{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^l)^j (1-p^l)^{T-m-1-j}}{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^h)^j (1-p^h)^{T-m-1-j}}.$$

To show that for T big enough this expression is negative, it is enough to

show that

$$\lim_{T \rightarrow \infty} \frac{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^l)^j (1-p^l)^{T-m-1-j}}{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^h)^j (1-p^h)^{T-m-1-j}} = 0.$$

Note that

$$\begin{aligned} & \frac{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^l)^j (1-p^l)^{T-m-1-j}}{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^h)^j (1-p^h)^{T-m-1-j}} \\ & \leq \frac{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^l)^j (1-p^l)^{T-m-1-j}}{(p^h)^{T-m-1}} \\ & \leq \frac{\binom{T-m-1}{\lceil \alpha(T-m-1) \rceil} (p^l)^{\lceil \alpha(T-m-1) \rceil} (1-\alpha)(T-m-1)}{(p^h)^{T-m-1}} \end{aligned}$$

where the first inequality follows from the fact that $(p^h)^{T-m-1}$ is just one of the elements in summation, while the second inequality follows since for $\alpha \geq 1/2$ the monotonicity of the binomial coefficient implies that $\binom{T-m-1}{\lceil \alpha(T-m-1) \rceil} \geq \binom{T-m-1}{j}$ for any $j \in \{\lceil \alpha(T-m-1) \rceil, \dots, T-m-1\}$ and $p^l \in (0, 1)$.

From Stanica (2001; Corollary 2.3) it follows that for $\alpha \geq 1/2$ the binomial coefficient is bounded by

$$\binom{T-m-1}{\lceil \alpha(T-m-1) \rceil} \leq \frac{1}{\sqrt{2\pi\alpha(1-\alpha)(T-m-1)}} \frac{\left(\frac{1}{\alpha}\right)^{T-m-1}}{\left(\frac{1}{\alpha}-1\right)^{(T-m-1)(1-\alpha)}}.$$

Plugging this bound into the previous expression yields

$$\begin{aligned} & \frac{\binom{T-m-1}{\lceil \alpha(T-m-1) \rceil} (p^l)^{\lceil \alpha(T-m-1) \rceil} (1-\alpha)(T-m-1)}{(p^h)^{T-m-1}} \\ & \leq \frac{\sqrt{(1-\alpha)(T-m-1)}}{\sqrt{2\pi\alpha}} \left(\frac{\frac{1}{\alpha} (p^l)^\alpha}{\left(\frac{1}{\alpha}-1\right)^{(1-\alpha)} p^h} \right)^{T-m-1} \\ & = \frac{\sqrt{(1-\alpha)(T-m-1)}}{\sqrt{2\pi\alpha}} \left(\frac{(p^l)^\alpha \left(\frac{1}{\alpha}-1\right)^\alpha}{(1-\alpha)p^h} \right)^{T-m-1} \end{aligned}$$

Therefore, for $(p^l)^\alpha \left(\frac{1}{\alpha} - 1\right)^\alpha < (1 - \alpha)p^h$ we have

$$\lim_{T \rightarrow \infty} \frac{\sqrt{(1 - \alpha)(T - m - 1)}}{\sqrt{2\pi\alpha}} \left(\frac{(p^l)^\alpha \left(\frac{1}{\alpha} - 1\right)^\alpha}{(1 - \alpha)p^h} \right)^{T - m - 1} = 0$$

which completes the proof. ■

8 Appendix B: The Pooling Contract-pair

In this appendix we turn our attention to the second case, where $\pi_{(l,e)} \leq \pi_{(h,d)}$. In solving for the optimal contract we take a route as similar as possible to the one used in solving for the first case, where $\pi_{(l,e)} > \pi_{(h,d)}$. That is, we start by assuming that the surgeon is known to be a low-quality one, and define the set of optimal contracts Γ_T^l . After characterizing Γ_T^l , we drop the assumption that the surgeon is known to be a low-quality one and show, that when surgeon's quality is unobservable the contract-pair (τ_T^h, τ_T^l) is optimal only if $\tau_T^l \in \Gamma_T^l$. Equipped with this result, it is rather easy to characterize the optimal contract pair $(\bar{\tau}_T^h, \bar{\tau}_T^l)$ and show that $\bar{\tau}_T^h \equiv \bar{\tau}_T^l$.

8.1 Surgeon's Quality is Known to be Low

First note that unlike the contract for the high-quality surgeon who is expected to operate on all types of patients the contract for a low-quality surgeon imposes no such requirement. It is thus necessary to consider also payments after histories along which at some t the surgeon's choice was not to operate. A two-period contract for a low-quality surgeon is depicted below.

One of the differences between the current case and the previous one is that the optimal mechanism provides no information rents to the surgeon. In particular, note that a contract that pays a constant sum of $\frac{c}{\pi_{(l,e)}}$ per success provides efficient incentives and generates an expected utility of zero to the surgeon, which, in particular, implies that it is optimal. The next lemma provides a characterization of the set Γ_T^l of T -periods *optimal* contracts when the surgeon is known to be a low-quality one.

Lemma 4 *Properties of Γ_T^l*

1. If a contract $\tau_T^l \in \Gamma_T^l$ then, $u^l(\tau_T^l, \text{sub}_1) - u^l(\tau_T^l, \text{sub}_{-1}) = \frac{c}{\pi_{(l,e)}}$ and $u^l(\tau_T^l, \text{sub}_0) = u^l(\tau_T^l, \text{sub}_{-1})$.




Figure 3: Two-period Contract for a Low-quality Surgeon

2. If a contract $\tau_T^l \in \Gamma_T^l$ then $\exists \tau_{T-1}^l \in \Gamma_{T-1}^l$ s.t. $\forall \omega_{T-1} \in \Omega_{T-1}$, $\tau_T^l(-1, \omega_{T-1}) = \tau_{T-1}^l(\omega_{T-1})$. Also $\exists \tau_{T-1}^l \in \Gamma_{T-1}^l$ s.t. $\forall \omega_{T-1} \in \Omega_{T-1}$, $\tau_T^l(0, \omega_{T-1}) = \tau_{T-1}^l(\omega_{T-1})$.
3. If a contract $\tau_T^l \in \Gamma_T^l$ then $m^l(\tau_T^l, sub_\emptyset) = T \cdot c \cdot p^l \frac{1-q}{\pi(l,e)}$.
4. Assume that a T -period contract τ_T^l satisfies EF and IR and for which $u^l(\tau_T^l, sub_\emptyset) = u > 0$. Then for any $\tilde{u} \in [0, u)$ there exists another T -period contract $\tilde{\tau}_T^l$ that also satisfies EF and IR and for which $u^l(\tilde{\tau}_T^l, sub_\emptyset) = \tilde{u}$. Moreover, for any history $\omega_T \in \Omega_T$ we have $\tau(\omega_T) \geq \tilde{\tau}(\omega_T)$ with at least one strict inequality.

Proof: Property 1. First, observe that if $u^l(\tau_T^l, sub_1) - u^l(\tau_T^l, sub_{-1}) < \frac{c}{\pi(l,e)}$ the low-quality surgeon will not exert effort even if an easy patient arrives in the first period. Also note that $u^l(\tau_T^l, sub_0) \geq u^l(\tau_T^l, sub_{-1})$ since otherwise the surgeon will perform surgery without exerting effort when a difficult patient arrives in the first period. Assume then that $\tau_T^l \in \Gamma_T^l$ but $u^l(\tau_T^l, sub_1) - u^l(\tau_T^l, sub_{-1}) > \frac{c}{\pi(l,e)}$. Consider then the following changes: in sub_0 adopt the same payment as in sub_{-1} and in sub_1 add a payment of $\frac{c}{\pi(l,e)}$ to every history of sub_{-1} . Note that this changes preserve incentives and decreases the expected payment, in contradiction that $\tau_T^l \in \Gamma_T^l$.

Property 2. Assume that this property is false. Since $\tau_T^l \in \Gamma_T^l$, τ_T^l provides sufficient incentives in all subgames, and in particular in sub_{-1} . Consider revising the contract τ_T^l to $\tilde{\tau}_T^l$ by replacing the payments in all histories that belong to sub_{-1} , adopting instead the payments in one of the optimal $T - 1$ -period contracts $\tilde{\tau}_{T-1}^l \in \Gamma_{T-1}^l$ (that is, $\tilde{\tau}_T^l(-1, \omega_{T-1}) = \tilde{\tau}_{T-1}^l(\omega_{T-1})$) and adjusting the contracts in other subgames correspondingly (that is, adopting in sub_0 the same payments as in sub_{-1} , while adopting in sub_1 the same payments as in sub_{-1} and adding $\frac{c}{\pi_{(l,e)}}$ to every history). Clearly, the proposed change preserves incentives to invest in all subgames after the first period and generates efficient incentives in the first period.

Since the new payment scheme in sub_{-1} is a contract in Γ_{T-1}^l it minimizes expected payment in all schemes that provide incentives. It follows that the proposed change strictly decreases $u^l(\tau_T^l, sub_{-1})$ because otherwise the τ_T^l payments restricted to sub_{-1} is a contract from Γ_{T-1}^l . Property 1 of the lemma implies that this changes also decreases $u^l(\tau_T^l, sub_0)$ and $u^l(\tau_T^l, sub_1)$, in contradiction to $\tau_T^l \in \Gamma_T^l$. The same argument also establishes that the payments in sub_0 is a contract in Γ_{T-1}^l .

Property 3. Consider the contract that pays $c/\pi_{(l,e)}$ per success (i.e., pays $\frac{nc}{\pi_{(l,e)}}$ after a history of n successful operations). Note that this is an optimal contract even when the designer observes the type of the arriving patient and the effort exerted by the surgeon. Therefore, it is an optimal contract when the patient's type and the surgeon's effort are not observable. Since in this contract $m^l(\tau_T^l, sub_\emptyset) = T \cdot c \cdot p^l \frac{1-q}{\pi_{(l,e)}}$ any optimal contract should pay the same expected payment as the one described above.

Property 4. The proof is done by induction on the contract's length T . Start with $T = 1$ and observe that since τ_1^l satisfies EF we have that

$$\begin{aligned} \tau_1^l(0) &\geq \pi_{(l,d)}\tau_1^l(1) + (1 - \pi_{(l,d)})\tau_1^l(-1) - c & (4) \\ \tau_1^l(0) &\geq \tau_1^l(-1) \\ \tau_1^l(0) &\leq \pi_{(l,e)}\tau_1^l(1) + (1 - \pi_{(l,e)})\tau_1^l(-1) - c \\ \tau_1^l(-1) &\leq \pi_{(l,e)}\tau_1^l(1) + (1 - \pi_{(l,e)})\tau_1^l(-1) - c \end{aligned}$$

There are two cases to consider.

Case 1 $\tau_1^l(0) \geq \tilde{u}$. From (4) we get that

$$\pi_{(l,e)}\tau_1^l(1) + (1 - \pi_{(l,e)})\tau_1^l(-1) - c \geq \tau_1^l(0) \geq \tilde{u}. \quad (5)$$

So first set $\tilde{\tau}_1^l(0) = \tilde{u}$. If $\tau_1^l(-1) \geq \tilde{u}$ then set $\tilde{\tau}_1^l(0) = \tilde{\tau}_1^l(-1) = \tilde{u}$ and $\tilde{\tau}_1^l(1) = \tilde{u} + \frac{c}{\pi_{(l,e)}} \leq \tau_1^l(1)$, where the last inequality follows from the

fact that the original payment satisfied EF , which in particular implies that $\tau_1^l(1) \geq \tau_1^l(-1) + \frac{c}{\pi(l,e)}$. Now note that since the expected payment are strictly lower in $\tilde{\tau}_1^l$ and both contracts satisfy EF , there exists at least one history where the payment in $\tilde{\tau}_1^l$ is strictly lower than in τ_1^l . If $\tau_1^l(-1) < \tilde{u}$ then (5) and the fact that expected utility in τ_1^l is u imply that $\tau_1^l(1) > \frac{c+\tilde{u}-(1-\pi(l,e))\tau_1^l(-1)}{\pi(l,e)}$. Setting $\tilde{\tau}_1^l(1) = \frac{c+\tilde{u}-(1-\pi(l,e))\tau_1^l(-1)}{\pi(l,e)}$, and $\tilde{\tau}_1^l(-1) = \tau_1^l(-1)$. Recall that $\tilde{\tau}_1^l(0) = \tilde{u}$ and observe that $\tilde{\tau}_1^l$ generates an expected utility of \tilde{u} and satisfies EF .

Case 2 $\tau_1^l(0) < \tilde{u}$. We start by setting $\tilde{\tau}_1^l(0) = \tau_1^l(0)$ and proceed by decreasing the utility from exerting effort by $\frac{u-\tilde{u}}{1-q}$, which will generate for $\tilde{\tau}_1^l$ an expected utility of \tilde{u} . If $\tau_1^l(-1) \geq \frac{u-\tilde{u}}{1-q}$ then set $\tilde{\tau}_1^l(\omega_1) = \tau_1^l(\omega_1) - \frac{u-\tilde{u}}{1-q}$ for $\omega_1 \in \{1, -1\}$. If however $\tau_1^l(-1) < \frac{u-\tilde{u}}{1-q}$ then set $\tilde{\tau}_1^l(-1) = 0$ and $\tilde{\tau}_1^l(1) = \frac{\tilde{u}+c-q\tilde{\tau}_1^l(0)}{(1-q)\pi(l,e)} < \tau_1^l(1)$, where the last inequality follows from the fact that when $\tau_1^l(-1) < \frac{u-\tilde{u}}{1-q}$, decreasing the utility of the agent from exerting effort by $\frac{u-\tilde{u}}{1-q}$ implies that the payment conditional on success should be decreased by more than the amount decreased in the case where $\tau_1^l(-1) \geq \frac{u-\tilde{u}}{1-q}$. However, since $\tilde{u} - q\tilde{\tau}_1^l(0) > 0$, the payments in $\tilde{\tau}_1^l$ satisfy EF and all payments are lower.

Having established the claim for $T = 1$ we proceed by assuming that the statement holds for $T = K - 1$ periods and show that it holds for $T = K$ periods. Assume that there exists τ_K^l for which $u^l(\tau_K^l, sub_\emptyset) = u > 0$. Similarly to the proof for $T = 1$, there are two cases to consider.

Case 1 $u^l(\tau_K^l, sub_0) \geq \tilde{u}$. We start by replacing the payment in sub_0 with a $K - 1$ -period contract $\tilde{\tau}_{(K-1)_0}^l$ that generates an expected payment of \tilde{u} and for which $\tilde{\tau}_{(K-1)_0}^l(\omega_{K-1}) \leq \tau_K^l(0, \omega_{K-1}) \forall \omega_{K-1} \in \Omega_{K-1}$. Such a contract exists by the induction argument. If $u^l(\tau_K^l, sub_{-1}) \geq \tilde{u}$ then we replace the payments in sub_{-1} with a $K - 1$ -periods contract $\tilde{\tau}_{(K-1)_{-1}}^l(\omega_{K-1})$ that generates an expected payment of \tilde{u} and for which we have $\tilde{\tau}_{(K-1)_{-1}}^l(\omega_{K-1}) \leq \tau_K^l(-1, \omega_{K-1}) \forall \omega_{K-1} \in \Omega_{K-1}$ (again, such a contract exists by the induction argument). We complete this part of the argument by replacing the payments in sub_1 with a $K - 1$ -periods contracts $\tilde{\tau}_{(K-1)_1}^l(\omega_{K-1})$ that generates an expected payment of $\tilde{u} + \frac{c}{\pi(l,e)}$ and for which $\tilde{\tau}_{(K-1)_1}^l(\omega_{K-1}) \leq \tau_K^l(1, \omega_{K-1}) \forall \omega_{K-1} \in \Omega_{K-1}$ (again, such contract exists by the induction argument). We still have to show that there exists an history ω_K for which the inequality is strict. However, since the new contract $\tilde{\tau}_K^l$ generated from the three contracts $\tilde{\tau}_{(K-1)_z}^l$ for $z = -1, 0, 1$

generates a strictly lower expected payment and both contracts satisfy efficiency, there must exist at least one history for which the inequality is strict.

The proof of the case where $u^l(\tau_K^l, sub_0) < \tilde{u}$ is proved similarly. ■

Remark 4 *An immediate consequence of the lemma and in particular of Property 3 is that for all $\tau_T^l \in \Gamma_T^l$, we have $u^l(\tau_T^l, sub_\emptyset) = 0$.*

8.2 Surgeon's Quality is Unobservable

We are now ready to characterize the optimal contract-pair when $\pi_{(l,e)} \leq \pi_{(h,d)}$, which is shown to have a very simple structure. Namely, the two contracts are the same and they pay a fixed compensation per successful operation. Moreover, this contract belongs to the set of optimal contracts when the surgeon is known to be a low-quality surgeon. We start by establishing the latter.

Lemma 5 *When $\pi_{(l,e)} \leq \pi_{(h,d)}$, (τ_T^h, τ_T^l) is an optimal contract-pair, only if $\tau_T^l \in \Gamma_T^l$.*

Proof. Assume by way of contradiction that (τ_T^h, τ_T^l) is an optimal contract-pair but $\tau_T^l \notin \Gamma_T^l$. Since (τ_T^h, τ_T^l) is optimal the contract-pair is admissible and in particular satisfies *IR* and *EF*. Since $\tau_T^l \notin \Gamma_T^l$, Remark 4 implies that $u^l(\tau_T^l, sub_\emptyset) > 0$. Hence, Property 4 of Lemma 4 implies that there exists a contract $\hat{\tau}_T^l \in \Gamma_T^l$ such that for every history $\omega_T \in \Omega_T$, $\tau_T^l(\omega_T) \geq \hat{\tau}_T^l(\omega_T)$ with strict inequality for at least one $\omega_T \in \Omega_T$. Consider replacing (τ_T^h, τ_T^l) with the pair $(\tau_T^h, \hat{\tau}_T^l)$. To verify that this contract-pair satisfies *EF*, note first that since $\pi_{(h,d)} > \pi_{(l,e)}$ *EF* is satisfied for the high-quality surgeon whenever it is satisfied for the low-quality one, and the latter holds since $\hat{\tau}_T^l \in \Gamma_T^l$. Obviously, *IC* holds as well for this new contract-pair $(\tau_T^h, \hat{\tau}_T^l)$. By definition, the expected payment to the low-quality surgeon are now lower, and the same (with weak inequality) also holds for the high-quality surgeon. That is,

$$(i) \quad m^l(\hat{\tau}_T^l, sub_\emptyset) < m^l(\tau_T^l, sub_\emptyset) \quad \text{and} \quad (ii) \quad m^h(\hat{\tau}_T^l, sub_\emptyset) \leq m^h(\tau_T^h, sub_\emptyset).$$

To verify *(ii)*, recall that the original contract-pair (τ_T^h, τ_T^l) was incentive-compatible, which in particular implies that the high-quality surgeon prefers the contract τ_T^h to τ_T^l . By Property 4 of Lemma 4 the new contract $\hat{\tau}_T^l$ generates for the high-quality surgeon an even lower expected utility than τ_T^l . Since this contract satisfies *EF*, the monotonicity relation between

expected payment and expected utility implies (ii). This establishes the contradictions to the statement that the original contract-pair (τ_T^h, τ_T^l) was optimal. ■

Theorem 2 *When $\pi_{(l,e)} \leq \pi_{(h,d)}$, the optimal contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is*

$$\hat{\tau}_T^h(\omega_T) = \hat{\tau}_T^l(\omega_T) = \frac{c \cdot n(\omega_T)}{\pi_{(l,e)}}$$

where $n(\omega_T)$ is the number of successful operations in ω_T .

The proof of the theorem is a simple consequence of the following claim and hence will be provided after the proof of the claim.

Claim 3 *Assume that the principal is asked to provide the low-quality surgeon with an expected utility of $\bar{u} \geq 0$ but in a way that preserves incentives to exert efficient effort while at the same time minimizing the expected utility of the high-quality surgeon if he adopts this contract. This is achieved by amending the contract $\hat{\tau}_T^l$ described in Theorem 2 and adding a payment of \bar{u} after every history. That is,*

$$\hat{\tau}_T^l(\omega_T) = \frac{cn(\omega_T)}{\pi_{(l,e)}} + \bar{u} \quad \text{for all } \omega_T \in \Omega_T.$$

Proof. We prove by induction on the length of the contract T . Start with one period. Recall that in this case the only optimal contract for the low-quality surgeon is $\tau_1^l(1) = c/\pi_{(l,e)}$, $\tau_1^l(-1) = 0$, and $\tau_1^l(0) = 0$. Denote by $u(\omega_1)$ the additional payment above $\tau_1^l(\omega_1)$ for $\omega_1 \in \{-1, 0, 1\}$. First, note that $u(1) \geq u(-1)$, because otherwise the surgeon will not exert effort when an easy patient arrives. In addition, observe that $u(1) \geq u(0)$ as otherwise, the low-quality surgeon will not conduct surgery even if an easy patient arrives. Also note that $u(0) \geq u(-1)$, because otherwise the surgeon will conduct the surgery (maybe without exerting effort) even when the arriving patient has a major problem, $p_t = d$. Recall that since $\pi_{(l,e)} \leq \pi_{(h,d)}$, if incentives are provided for the low-quality surgeon to exert effort on $p_t = e$, then the high-quality surgeon will operate on all types of patients, if he faces the same contract. Moreover, $\pi_{(h,d)} \geq \pi_{(l,e)}$ implies that specifying $u(1) = u(0) = u(-1) = \bar{u}$ necessarily minimizes the utility of the high-quality surgeon from all contracts that generate efficient incentives for the low-quality surgeon and provides him with the additional utility of \bar{u} .

We assume that the statement holds for $T = K - 1$ periods and proceed to the proof of the statement for $T = K$ periods. Consider a contract τ_K^l that yields a utility of \bar{u} to the low-quality surgeon and minimizes the expected utility of the high-quality surgeon. We first show that the induced contract on sub_1 , sub_0 , and sub_{-1} by τ_K^l are as described in the statement of the claim. The reason for that is as follows, assume by way of contradiction that the above statement is false and note that: (i) it is always possible to construct a contract $\tilde{\tau}_K^l$ such that the induced contracts on sub_1 , sub_0 , and sub_{-1} by $\tilde{\tau}_K^l$ are as described in the statement of the claim and for which there are \bar{u}_1 , \bar{u}_{-1} , and \bar{u}_0 such that the low-quality surgeon is indifferent between $\tilde{\tau}_K^l$ and τ_K^l , (ii) the incentives for the low-quality surgeon are preserved in $\tilde{\tau}_K^l$, and (iii) by the induction argument, in each of the subgames, the amended contract $\tilde{\tau}_K^l$ decreases the expected utility of the high-quality surgeon. We still need to show that $\bar{u}_1 = \bar{u}_{-1} = \bar{u}_0 = \bar{u}$. However, this proof is identical to the proof of the one-period case. ■

Proof of Theorem 2. First observe that $\hat{\tau}_T^l \in \Gamma_T^l$ and that $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ satisfies *EF* and *IR*. It follows that if we prove that $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is *optimal* we are done. As in Theorem 1, we need to show that the contract $\hat{\tau}_T^l$ described in the theorem minimizes the expected utility of the high-quality surgeon from all contracts belonging to Γ_T^l but the rest of the proof follows from the previous claim for $\bar{u} = 0$. ■

References

- [1] Baron, D. & Besanko, D. (1984): “Regulation and Information in a Continuing Relationship,” *Information Economics and Policy*, 1, 267-302.
- [2] Biais, B., Mariotti, T., Rochet, J.-C. & Villeneuve, S. (2010): “Large Risks, Limited Liability, and Dynamic Moral Hazard,” *Econometrica* 78, 73-118.
- [3] Dranove, D., Kessler, D., McClellan, M., & Satterthwaite, M. (2003): “Is More Information Better? The Effect of “Report Cards” on Health Care Providers,” *Journal of Political Economy* 111, 555-588.
- [4] Dulleck, U. & Kerschbamer, R. (2006): “On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods,” *Journal of Economic Literature* XLIV, 5-42.
- [5] Emons, W. (1997): “Credence Goods and Fraudulent Experts,” *Rand Journal of Economics* 28, 107-119.

- [6] Fong, K. (2009): "Evaluating Skilled Experts: Optimal Scoring Rules for Surgeons," memo.
- [7] Fudenberg, D., Holmstrom, B. & Milgrom, P. (1990): "Short-Term Contracts and Long-Term Agency Relationships," *Journal of Economic Theory* 51, 1-31.
- [8] Gruber, J., Kim, J. & Mayzlin, D. (1999): "Physician Fees and Procedure Intensity: The Case of Cesarean Delivery," *Journal of Health Economics* 18, 473-490.
- [9] Holmstrom, B. (1982): "Managerial Incentive Problems: A Dynamic Perspective," *Essays in Economics and Management in Honor of Lars Wahlbeck*. Helsinki, Swedish School of Economics.
- [10] Holmstrom, B. & Milgrom, P. (1987): "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica* 55, 303-328.
- [11] Hughes, D. & Yule, B. (1992): "The Effect of Per-item Fees on the Behavior of General Practitioners," *Journal of Health Economics* 11, 413-437.
- [12] Laffont, J. & Tirole J. (1988): "The Dynamic of Incentive Contracts," *Econometrica* 56, 1153-1175.
- [13] Malcomson, J. & Spinnewyn, F. (1988): "The Multiperiod Principal-Agent Problem," *The Review of Economic Studies* 55, 391-407.
- [14] Rey, P. & Salanie, B. (1990): "Long-Term, Short-Term, and Renegotiation: On the Value of Commitment in Contracting," *Econometrica* 58, 597-619.
- [15] Rubinstein, A. & Yaari, M. (1983): "Repeated Insurance Contracts and Moral Hazard," *Journal of Economic Theory* 30, 74-97.
- [16] Sannikov, Y. (2007): "Agency Problems, Screening and Increasing Credit Lines," memo
- [17] Stanica, P. (2001): "Good Lower and Upper Bounds on Binomial Coefficients", *Journal of Inequalities in Pure and Applied Mathematics*, Vol 2(3), 1-5.
- [18] Wolinsky, A. (1993): "Competition in a Market for Informed Experts' Services," *Rand Journal of Economics* 24, 380-398.
- [19] Wolinsky, A. (1995): "Competition in Markets for Credence Goods," *Journal of Institutional and Theoretical Economics* 151, 117-131.