Moral Hazard, Insurance,

and Some Collusion

By

Ingela Alger
and
Ching-to Albert Ma

DISCUSSION PAPER 318

March 1999

FINANCIAL MARKETS GROUP
AN ESRC RESEARCH CENTRE

LONDON SCHOOL OF ECONOMICS

E·S·R·C
ECONOMIC
& SOCIAL
RESEARCH
COUNCIL

ISSN 0956-8549-318

Moral Hazard, Insurance, and Some Collusion

Ingela Alger* and Ching-to Albert Ma**

August, 1998

Abstract

A risk-averse consumer purchases an insurance policy; if she suffers a loss, she may receive services from a provider to recover some of the loss. Only the consumer and the provider know if the loss has actually occurred. The provider's behavior is uncertain. With some positive probability, the provider is honest, reporting the loss information truthfully to the insurer; with the complementary probability, the provider reports the information strategically, by writing a side-contract with the consumer to maximize the joint surplus of the provider-consumer coalition. We show that there is a loss of generality in considering only collusion-proof contracts, and characterize equilibria implemented by collusion-proof and noncollusion-proof contracts. When the probability of a provider acting collusively is small, the equilibrium contract is not collusion-proof but approximately first-best. When the probability of a provider acting collusively is large, the equilibrium contract is independent of this probability and identical to the equilibrium collusion-proof contract when the provider is collusive with probability 1.

*formerly Ingela Brundin; Financial Markets Group, London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom; E-mail: I.Brundin@lse.ac.uk

**Department of Economics, Boston University, Boston, Massachusetts 02215, United States of America; E-mail: ma@bu.edu

Acknowledgement: The research was conducted while the first author was visiting the Department of Economics at Boston University, and while she was supported by the Sweden-America Foundation. We thank Michael Manove for discussing various issues with us, and Thomas McGuire and seminar participants at a number of universities for their comments.

1 Introduction

Economists are often very good at spotting incentive problems of resource allocation mechanisms. Indeed, for the past two decades or so, the development of incentive theory in economics has proceeded at an exceedingly rapid rate. Many attempts have been made to understand why contractual arrangements are not \incentive compatible," \renegotiation-proof" or \collusion-proof." It is not uncommon, however, to ¯nd economic agents sometimes behaving honestly or nonstrategically. The behavior of economic actors is in°uenced by culture and established rules of thumb, which often persuade economic agents to act honestly. Moreoever, the emergence of cooperative and sincere behavior has been discussed by a large economics literature of repeated games and evolutionary processes. Nevertheless, the implication of honesty on incentive theory has been largely unexplored. In a major departure from the literature, we introduce a model in which some economic agents act honestly.

We develop our theory by a model of contract design between three parties: a risk-averse consumer, a provider, and an insurer. The consumer may su®er a loss, the provider can (partially) recover the loss by supplying some inputs, and the insurer o®ers an insurance contract to the consumer and a payment contract to the provider. This canonical model has been studied extensively in the economics literature, and the fundamental incentive problems are often summarized as moral hazard, hidden action, and nonveri¯able information. We focus on the incentives for the consumer and the provider to misreport their private information about the loss to the insurer. We make the usual assumptions for models of asymmetric information | except that we introduce the possibility of honest behavior for the consumer and the prov honest behavior may also be optimal in short-term relationships, as shown by Tirole (1996). In a model where an agent is matched with a new principal each period, Tirole determines the optimal behavior of opportunistic agents, who co-exist with agents who always cheat and agents who never cheat. He shows that under some conditions, an opportunistic agent maximizes his utility by never cheating the principal.

Second, a number of papers have suggested that human behavior is driven by factors other than pure sel¯shness. Rabin (1998) surveys the psychology literature and concludes that the standard assumption of sel¯sh behavior in economics may be too narrow, and sometimes even misleading. For instance, many people simply dislike being dishonest. Furthermore, extensive experimental evidence in bargaining games has demonstrated that economic agents will respond with sincerity when they think that they have been treated fairly (see Rabin, 1998, p21®). This indicates that factors other than purely economic ones a®ect our behavior. Rabin also quotes the example, originally due to Dawes and Thaler (1988), of farmers leaving fresh produce on a table with a box nearby, expecting customers to pick their purchases and leave money to \complete" the transaction. Similarly, most of us have voluntarily paid for newspapers inside an unlocked box on a sidewalk. In a paper on endogenous preferences, Bowles (1998, p80) \treat[s] preferences as cultural traits, or learned in°uences of behavior" and includes such examples as never lying and reciprocating dinner invitations. Taken together, these ¯ndings strongly indicate that honest behavior is neither uncommon nor irrational. We explore the implication of such behavior on a standard model of contract design under incomplete information. We stress that we do not assume that all economic agents are honest; we only assume that some are.

Once our model is accepted, the common methodology of solving a model with asymmetric information by introducing a set of truth-telling constraints is no longer valid. Indeed, a resource-allocation mechanism that exploits the honesty of economic agents must let those who behave strategically succeed in \gaming" against the system. This may seem ironic at ¯rst sight, but in fact may be the basis for why we sometimes ¯nd actual mechanisms to fail to be incentive compatible or strategy-proof. Equilibria in our model may remain at the usual second best (which obtains when economic agents are always strategic), or it may be better, but it is never first best. Whenever equilibria are better than the second best, the (strategic) provider and the consumer misreport their private information in equilibrium. Equilibria tend to the first best as the probability of the provider being honest goes to one, but whenever this probability is within a neighborhood of zero, equilibria do not change with the probability and are second best.

In our model, a risk-averse consumer may suffer a loss (due to an accident or an illness, for example), and a provider can partially restore this loss by supplying inputs for a recovery. Insurance is offered by a risk-neutral insurer by means of contracts with the consumer and the provider. We consider the general class of deterministic contracts, each of which consists of a menu of insurance and payment policies and recovery inputs. The information of whether a loss has actually occurred or not is not available to the insurer. Payments and recovery inputs can only be based on the claim that the insurer receives.

We assume that this claim is the outcome of a collusive side-contract.[1] A side-contract is an agreement between the consumer and the provider to make a report to the insurer and to arrange a side payment between them. The side-contract allows them to choose the claim which maximizes their joint surplus. However, our central thesis is that only some providers are willing to falsify claims. Therefore, there are two types of providers: those who always make their claims truthfully, and those who collude with consumers and make claims strategically. Under this framework, we determine the equilibrium insurance and provider payment contract.

Our first result (Lemma 1) states that some equilibrium allocations cannot be implemented by collusion-proof contracts. In other words, some equilibrium allocations can only be obtained by letting a collusive provider and the consumer agree to misrepresent the loss information by a side-contract. This is an important point of the analysis, for it says that restricting attention to truthful revelation of the loss information in equilibrium involves a loss of generality. What is the intuition for this result? An honest provider always forgoes any gain from information

[1] Ma and McGuire (1997) and Alger and Salanié (1997) also study claims as a result of collusion between the provider and the consumer. These papers both assume that there are only collusive provider types.

manipulation. A contract allowing collusion therefore rewards only the collusive type. By accommodating collusion, the insurer awards a rent only to the collusive provider, but not the honest one. Therefore, equilibrium allocations that reward the two types of provider di®erently must imply equilibrium collusion.

We must consider two distinct classes of contracts: those which deter collusion, and those which do not. Interestingly, the optimal collusion-proof contract is independent of the likelihood of the provider's type. In fact, it is the same as the equilibrium policy in a model where the provider is known to be collusive always (Propositions 1 and 2). The intuition is easy to understand. When faced with a collusion-proof policy, a collusive provider will not gain from misreporting the loss information. This implies that even if a collusive provider picks a policy that is meant for the honest provider, there will not be any gain from misreporting. So the policy for the honest provider must also be collusion-proof! When collusion is deterred, both types of provider will be treated as if they were collusive. This contract is second-best, recovery inputs being excessive and risk sharing for the consumer imperfect.

This clearly demonstrates that deterring collusion cannot be always optimal. For example, when the likelihood of a collusive provider is very small, a ¯rst-best contract, which is not collusion-proof, must perform better. Nevertheless, we show that the ¯rst-best contract is not an equilibrium contract even when the provider is almost always honest. Indeed, in Proposition 3 we characterize the optimal noncollusion-proof contract for any given distribution on the provider's type. Collusion takes the form of the provider-consumer coalition always claiming that a loss has occurred. This leads to waste in recovery inputs and rents for the provider. The optimal noncollusion-proof policy for the honest provider implements excessive recovery inputs; for the collusive provider, insu±cient inputs. Risk sharing for the consumer is never perfect. More importantly, when compared to the ¯rst best, the distortion in the honest provider's policy varies strictly monotonically with the probability of the honest provider, and tends to zero as this probability tends to 1. This contrasts with the optimal collusion-proof contract, which is independent of the probability.

Allowing collusion is the only way to exploit honest behavior. If collusion is completely deterred, the optimal contract is always very different from the first best, but permitting collusion allows the contract to be approximately first-best when the provider is almost always honest. Combining these results, we conclude that if the probability of a collusive provider is small, equilibrium contracts allow collusion. Conversely, if this probability is large, the equilibrium contracts must be collusion-proof.

The analysis of side-contracts in recent literature follows Tirole's innovation (1986); the modeling of collusion in our paper fits into that framework. But our model departs significantly from Tirole's paper and the related literature (see the surveys by Tirole (1992) and Laffont and Rochet (1997)) by assuming that the provider may not always be willing to manipulate information. Hence, the method of considering only collusion-proof equilibria does not apply. The result that contracts in our environment must allow collusion parallels that of Alger and Renault (1998), where it is shown that in a general model with a single agent who is honest with some probability, attention cannot be restricted to incentive-compatible contracts. Recently Ma and McGuire (1997) study the effect of collusion on payment and insurance contracts. The model there is quite different, allowing the provider to choose efforts and the consumer to choose quantities. More importantly, in Ma and McGuire (1997) the consumer and provider cannot write any side contract, and collusion may not always maximize their joint surplus.

Kofman and Lawarrée (1996) also use the assumption that some economic agent may not be corrupt. In a model of auditing, they assume that the auditor can be honest or dishonest. Whereas they do not allow different contracts for different types of auditors, we do. In other words, only pooling contracts are considered by Kofman and Lawarrée, but separating contracts are studied here. Our results imply that it is restrictive to assume pooling, since we find that separating contracts may be optimal.[2] Our results can therefore be interpreted to imply that Kofman and Lawarrée's analysis involves a loss of generality.

Another set of papers has explored the implications of costly state falsification; see Lacker

[2] In their model, Kofman and Lawarrée (1996) demonstrate the optimality of collusion when the probability of the dishonest auditor is sufficiently small. But given that contracts are pooling, this is not surprising. Indeed, to deter collusion, the dishonest auditor must be rewarded; pooling implies that the honest auditor also receives this reward, although this is unnecessary for the honest auditor. The expected cost of deterring collusion must outweigh its benefits when the auditor is honest with a high probability.

and Weinberg (1989), Maggi and Rodriguez-Clare (1995) and Crocker and Morgan (1998). These papers assume that an agent possessing private information must incur some costs when misreporting the information. This surprisingly implies that the optimal contract induces the agents to falsify information. Allowing some type of an agent to falsify information raises the cost of lying for other types of the agent.[3] Although our result that allowing collusion can improve welfare appears to be similar, the reasons for our results are quite di®erent. First, we assume that some economic agents do not manipulate information. Second, while researchers working in that literature may apply the revelation principle, we may not. Indeed, in our model, in some equilibrium the collusive provider makes false claims, but there exists no truth-telling equilibrium of a direct relevation mechanism yielding the same equilibrium outcome.

The next section presents the model and the ¯rst best; we also de¯ne the side-contract and the extensive form. The analysis is presented in Section 3. First, we de¯ne a collusion-proof policy, and show that generally both collusion-proof and noncollusion-proof policies must be considered. The following two subsections respectively consider contracts that deter and permit collusion. Our main results are also presented there. Finally, concluding remarks are made in the last section. Proofs are contained in the appendix.

2 The Model

We study the design of contracts between an insurer, a provider and a consumer. We begin by presenting the basic setting and the ¯rst best. Then we de¯ne the side-contracting subgame played by the provider and the consumer, as well as the extensive form of the game played by the three parties. Our model is quite general, but because its application to health and automobile insurance is straightforward, often we illustrate our model by these markets.[4] The consumer su®ers a loss with some probability p. The loss can be due to an illness (as in the

[3] Take the example of sharecropping, where falsi¯cation amounts to hiding part of the crop. By making sharecroppers with a low production outcome conceal some of the crop, the landlord increases the cost of lying for sharecroppers with a high output, since the tenant with a high output will have to destroy a larger amount of crops to mimic the tenant with a low output. Hence, less rents need be given up to him.

[4] See Lu (1997) for an example of evidence of misreporting in the health care sector.

health insurance case) or an accident (respectively, the automobile insurance case). We express the loss in monetary terms and denote it by `. The consumer is risk averse with an increasing and strictly concave utility function U, de¯ned in terms of money. Her initial wealth is W.

The provider has a technology that recovers some of the consumer's loss if that has occurred, and he is risk neutral. The technology is completely unproductive when the consumer has not experienced a loss. Let m denote the input in this production technology of loss recovery. In the health market, m denotes the quantity of treatment; in the automobile market, the repair work. Each unit of the input m costs the provider c. The output is measured in monetary units. If m units of the input is supplied by the provider, the consumer recovers f(m), where f is an increasing and concave function. Nevertheless, we assume that f is bounded from above, and that max f(m) < `; in other words, the technology cannot completely recover the consumer's total loss `. Many situations ¯t this assumption. This is certainly true for illnesses requiring surgery causing scars, or for illnesses which cannot be completely cured. The assumption may also re°ect the pain and discomfort su®ered by the patient during illness. For the automobile market, the assumption may re°ect the inconvenience caused by a repair of the car after an accident, the time and e®ort required to ¯nd a replacement, or the unavoidable risk that a repair might have been done improperly.[5]

If the consumer su®ers a loss, and pays an amount t in order to obtain m units of recovery, her utility is U(W − ` + f(m) − t); if the provider supplies m units of recovery in return for a remuneration ®, his utility is ® − mc. We suppose that the provider can always refuse to serve the consumer, in which case he obtains his reservation utility which is normalized to zero.

An insurer o®ers an insurance policy to the consumer, together with a reimbursement policy to the provider. The risk-netural insurer is assumed to operate in a competitive market, and to set its policy to maximize the consumer's expected utility. As we have de¯ned it, there are two possible states of nature, namely whether the consumer su®ers a loss ` or not.[6] We use the index i = h; s to denote these states: i = s is the state when the consumer's loss is `,

[5] Alternatively, we can assume that the marginal cost of recovery, c, is su±ciently high.

[6] We will introduce another set of states later; they concern the strategic interaction between a consumer and a provider.

otherwise, $i = h$. In the health-care market, a consumer may suffer from some symptoms. If those symptoms are actually insignificant, then the true state of nature is h (the "healthy" state); otherwise, the state is s (the "sick" state), and the loss ℓ occurs. In the automobile insurance market, the states of nature refer to whether an accident has actually led to a loss.

In the first best, whether the consumer has suffered the loss ℓ is verifiable information, and a contract can be based on that. The first-best insurance policy specifies two transfers, t_s and t_h, that the consumer must pay to the insurer when the states are s and h, respectively; in addition, in state s, the amount of recovery input, m, is provided to the consumer. Recovery input is not used in state h, and the payments to the provider are 0 in state h, and mc in state s. The first best is denoted by $(t_h^*; t_s^*; m^*)$, and given by the solution of the following program: choose t_h, t_s and m to maximize

$$(1-p)U(W-t_h) + pU(W-\ell+f(m)-t_s)$$

subject to

$$(1-p)t_h + p(t_s - mc) \geq 0:$$

The constraint is the insurer's break-even condition: the total expected revenue from the consumer is $(1-p)t_h + t_s$, and the expected payment to the provider is pmc.

Straightforward calculation from the first-order conditions yield the following:

$$t_h^* = t_s^* + \ell - f(m^*) \tag{1}$$

$$f'(m^*) = c: \tag{2}$$

The first condition says that all risks are absorbed by the insurer: the consumer's variation of net incomes over the two states is eliminated. The second is the productive efficiency condition: the marginal benefit of recovery $f'(m^*)$ is equal to the marginal cost c. Because $\ell > f(m^*)$, the consumer's payment in the event of a loss is lower: $t_h^* > t_s^*$. The consumer's risk aversion, together with the fact that maximum recovery cannot make up for the loss, implies that a smaller payment from the consumer in the event of a loss is necessary to achieve optimal risk sharing.[7]

Having established the ¯rst best, we proceed to study the interact

De¯nition 1 (The Side-Contract Subgame) Consider an insurance-payment policy consisting of a menu of two items: $\{(®_h; t_h); (®_s; t_s; m)\}$. The ¯rst item corresponds to state h; the second, state s. Within an item, ® is the provider's reimbursement from the insurer; t, the transfer the consumer pays to the insurer; m, the loss recovery input. A side-contract is de¯ned by an o®er $(x; k)$ made by the provider, where x is a side-transfer, and k, $k = h; s$, is the report. If the consumer accepts the side-contract, report k is made, and the transfer x is paid by the consumer to the provider. Otherwise, the true state of nature is reported.

By means of a side-contract, collusion allows the provider and the consumer to lie about the true state of nature. Observe that any party in the side-contract subgame can unilaterally enforce truthful reporting. If the true state is i, then to report truthfully the provider can simply o®er the consumer a transfer of 0 and a report of i; likewise, a consumer can always reject a side-contract o®er from the provider, and the true state will be reported. This implies that for the consumer and provider can agree on a (nontrivial) side-contra and $\alpha_s - mc + x \geq \alpha_h$. The side-contract makes both parties better off than a truthful report. Since the provider makes the side-contract offer, he can extract the whole surplus generated by the misreporting. Thus, if (3) holds, the provider offers a side-contract with the transfer x such that $t_s = t_h - x$, gaining $\alpha_s - mc - t_s - \alpha_h + t_h$.

To complete the description of the setting, we define the extensive form of the game played by the insurer, the consumer, and the provider. In stage 1, the consumer is randomly matched with a provider and "nature" determines the type of the provider; with probability θ the provider is a collusive type. The insurer offers an insurance-payment contract to the provider (without knowing the provider's type): we allow for the most general form of deterministic contracts, so the contract consists of two policies, one for each type of provider; each policy describes the transfers from the consumer to the insurer, and from the insurer to the provider, as well as a recovery input in case of loss. The contract can be written as $[\{(\alpha_h^\tau; t_h^\tau); (\alpha_s^\tau; t_s^\tau; m^\tau)\}; \{(\alpha_h^\sigma; t_h^\sigma); (\alpha_s^\sigma; t_s^\sigma; m^\sigma)\}]$. In stage 2, with the knowledge of whether he will always be truthful about reporting the consumer's loss or may write a side-contract with the consumer, the provider picks a policy $\{(\alpha_h^j; t_h^j); (\alpha_s^j; t_s^j; m^j)\}$, $j = \tau; \sigma$. Next, in stage 3, the consumer suffers the loss ℓ with probability p. Whether the consumer has suffered this loss becomes the consumer's and the provider's private information. In stage 4, the provider and the consumer play the side-contract subgame if the provider is collusive; at the end of this stage, a report on whether the consumer has suffered the loss is made. If the provider is collusive, this report is the result of the side-contract agreement. Otherwise, the report is h if and only if the consumer has not suffered a loss. Finally, payments to the provider, transfers from the consumer, and any recovery inputs are executed according to the selected policy and the report. We assume that the provider has limited liability and can exit the game any time to obtain his reservation profit or utility level, which is set at 0.

Observe that in stage 2, the provider picks the policy which maximizes his expected utility, which is calculated according to what he anticipates will happen in equilibrium in stage 4. That is, if he is a collusive provider, he will have to consider the possibility of side-contracts for the reporting of the information about the consumer's loss. In contrast, if he is an honest provider, he anticipates that he will reveal that information truthfully.

3 Analysis

If the provider always reports truthfully ($\mu = 0$), the insurer can rely on the provider's information and implement th

Inequalities (4) and (5) ensure that the joint surplus from making a truthful report exceeds the joint surplus from lying about the state of nature, in states h and s, respectively. Note that we assume that if the collusive provider is indifferent between writing a side-contract with the consumer and not, he will not do so. Indeed, the inequalities in Definition 2 are not required to hold strictly for the policy to be collusion-proof.

The insurer maximizes the expected utility of the consumer:

$$(6) \qquad (1-p)U(W - t_h^\sigma) + pU(W - \ell + f(m^\sigma) - t_s^\sigma)$$

subject to the collusion-proofness constraints (4) and (5), as well as the participation constraints for the provider and the insurer's break-even constraint:

$$(7) \qquad \alpha_h^\sigma \geq 0$$

$$(8) \qquad \alpha_s^\sigma - m^\sigma c \geq 0$$

$$(9) \qquad (1-p)(t_h^\sigma - \alpha_h^\sigma) + p(t_s^\sigma - \alpha_s^\sigma) \geq 0:$$

Proposition 1 When the provider is always collusive, the optimal collusion-proof contract $\{(\alpha_h^\sigma; t_h^\sigma); (\alpha_s^\sigma; t_s^\sigma; m^\sigma)\}$ has the following properties:

1. The provider obtains zero profit: $\alpha_h^\sigma = \alpha_s^\sigma - m^\sigma c = 0$.
2. The consumer pays the same transfer in states h and s: $t_s^\sigma = t_h^\sigma$.
3. The consumer is imperfectly insured: $W - \ell + f(m^\sigma) - t_s^\sigma < W - t_h^\sigma$.
4. Recovery input is excessive relative to the first best: $m^\sigma > m^*$, or $f'(m^\sigma) < c$.

The intuition for Proposition 1 is as follows.[10] In the first best, the transfer from the consumer to the insurer is reduced when there is a loss; this reduction is necessary to achieve efficient risk sharing, because the recovery cannot fully make up for the loss. Now, as we have pointed out, this creates an incentive for the consumer and provider to report that the state is

[10]The proofs of all results are in the appendix.

s when in fact the true state is h. The relevant (and binding) collusion-proofness constraint is thus the one which ensures that the consumer-provider coalition prefers reporting the truth in state h (constraint (4)). Two options are available to make this constraint hold: either increase the payment to the provider when he reports h, or reduce the di®erence between the transfers of the consumer to the insurer $t_h^{¾} - t_s^{¾}$, compared to the ¯rst best. The ¯rst option cannot be optimal: if the payment to the provider is positive, the expected utility of the consumer can be raised simply by decreasing this payment and decreasing the consumer's transfers to the insurer by the same amount. Hence, the provider must get zero pro¯t, and the optimal way to deter collusion is to reduce $t_h^{¾} - t_s^{¾}$ to zero. As a result, risk sharing through di®erences in monetary transfers across states h and s is impossible. The second-best contract therefore increases the loss recovery input from the ¯rst-best level in order to decrease the risk that the consumer is exposed to. So the marginal bene¯t of the recovery input is lower than its marginal cost.

3.2 Uncertain Provider Type

Now we turn to the model in which the provider colludes with the consumer with probability µ, 0 < µ < 1. A contract now takes the form: $[\{(®_h^{¿}; t_h^{¿}); (®_s^{¿}; t_s^{¿}; m^{¿})\}; \{(®_h^{¾}; t_h^{¾}); (®_s^{¾}; t_s^{¾}; m^{¾})\}]$. The two policies in the contract may provide incentives for di®erent types to self-select in stage 2 of the game. But of course, these policies may also be identical. So there is no loss of generality to consider only those equilibria of subgames starting at stage 2 in which the j type provider picks $\{(®_h^{j}; t_h^{j}); (®_s^{j}; t_s^{j}; m^{j})\}$, j = ¿; ¾.[11]

When the provider colludes with the consumer with probability µ, it is no longer possible to consider only collusion-proof policies. The need to consider policies that permit the provider and consumer to lie about the consumer's loss stems from the existence of the honest provider. An honest provider will not exploit a policy which is not collusion-proof. In other words, a policy that can be manipulated by the collusive provider may give a higher pro¯t to the collusive provider than the honest provider. The possibility of \di®erentially" rewarding di®erent types of the provider through a side-contract by the collusive type is the key to our

[11] This can be formally proved, as in Alger and Renault (1998).

analysis.

Consider a contract [{(®$^{¿}_{h}$; t$^{¿}_{h}$); (®$^{¿}_{s}$;

those which prevent collusion, and those which do not.

3.2.1 Optimal contract with collusion deterrence

We ¯rst study contracts that consist of collusion-proof policies for the type ¾ provider. Therefore, we consider those equilibria in which a provider and a consumer jointly report the states h and s truthfully. We now derive the optimal contract $[\{(®^{¿}_{h}; t^{¿}_{h}); (®^{¿}_{s}; t^{¿}_{s}; m^{¿})\}; \{(®^{¾}_{h}; t^{¾}_{h}); (®^{¾}_{s}; t^{¾}_{s}; m^{¾})\}]$ that deters collusion.

In stage 4, the provider of type ¾ must not ¯nd it pro¯table to write a side-contract, so constraints (4) and (5) must hold. Without loss of generality, the policy indexed by ¾ (resp. ¿) is to be chosen by the collusive (resp. truthful) type. Since these policies must allow the provider to obtain his reservation utility, we have the following participation constraints:

$$(10) \qquad ®^{j}_{h} \geq 0 \qquad j = ¾; ¿$$

$$(11) \qquad ®^{j}_{s} - m^{j} c \geq 0 \qquad j = ¾; ¿:$$

Next we state the incentive constraints that guarantee self-selection for the two types of the provider. In stage 2, the type j provider must ¯nd it optimal to select policy $\{(®^{j}_{h}; t^{j}_{h}); (®^{j}_{s}; t^{j}_{s}; m^{j})\}$, $j = ¿; ¾$, anticipating the equilibrium moves in stage 4. Type ¿ provider always reveals the true state in stage 4, regardless of the contract he has chosen. Therefore, the incentive constraint for the truthful provider type is the following:

$$(12) \qquad (1 - p)®^{¿}_{h} + p(®^{¿}_{s} - m^{¿} c) \geq (1 - p)®^{¾}_{h} + p(®^{¾}_{s} - m^{¾} c);$$

The condition for a collusive provider to choose the ¾ policy is slightly more involved. Although the policy indexed by ¾ is collusion-proof, the policy indexed by ¿ may not be. When a type ¾ provider selects a ¿ policy, he will consider any gain from a side-contract at stage 4. As we have seen above, because the provider makes the side-contract o®er, he expropriates all the potential gain from a side-contract. If type ¾ provider does pick the ¿ policy, and if the state turns out to be h, he can either forgo the side-contract to obtain $®^{¿}_{h}$ or use a side-contract to report s and get[13] $®^{¿}_{s} - m^{¿} c + t^{¿}_{h} - t^{¿}_{s}$. Similarly, if the state is s, the provider reports k = s

[13] The total surplus from reporting s is $®^{\tau}_{s} - m^{\tau} c - t^{\tau}_{s}$. Because the consumer must get $-t^{\tau}_{h}$, her transfer from a truthful report, the provider's gain is the di®erence between these two.

(respectively, k = h) when $®^{¿}_{s} - m^{¿}c$ is greater (respectively, smaller) than $®^{¿}_{h} - f(m^{¿}) + t^{¿}_{s} - t^{¿}_{h}$. To summarize, the incentive constraint that ensures that type ¾ provider picks the ¾ policy is:

(13) $$(1-p)®^{¾}_{h} + p(®^{¾}_{s} - m^{¾}c) \geq$$

$$(1-p)\max[®^{¿}_{h}; ®^{¿}_{s} - m^{¿}c + t^{¿}_{h} - t^{¿}_{s}] + p\max[®^{¿}_{s} - m^{¿}c; ®^{¿}_{h} - f(m^{¿}) + t^{¿}_{s} - t^{¿}_{h}]:$$

Finally, the insurer's budget constraint is:

(14) $$\mu[(1-p)(t^{¾}_{h} - ®^{¾}_{h}) + p(t^{¾}_{s} - ®^{¾}_{s})] + (1-\mu)[(1-p)(t^{¿}_{h} - ®^{¿}_{h}) + p(t^{¿}_{s} - ®^{¿}_{s})] \geq 0;$$

and the objective function is:

(15) $$\mu[(1-p)U(W - t^{¾}_{h}) + pU(W - ` + f(m^{¾}) - t^{¾}_{s})] +$$

$$(1-\mu)[(1-p)U(W - t^{¿}_{h}) + pU(W - ` + f(m^{¿}) - t^{¿}_{s})]:$$

An optimal contract deterring collusion maximizes (15) subject to (4), (5), (10), (11), (12), (13), and (14).

Proposition 2 Suppose $0 < \mu < 1$. The optimal contract deterring equilibrium collusion is independent of μ and o®ers the same policy to the truthful and collusive types of provider. Moreover, this policy is the optimal collusion-proof policy when the provider is always collusive ($\mu = 1$), namely the policy in Proposition 1.

Proposition 2 says that if contracts must deter collusion, then the existence of the truthful type of provider is inconsequential. The equilibrium contract is the same as if the provider is collusive with certainty, consisting of a single policy that deters collusion by the ¾ type provider; despite the fact that the ¿ type provider always reports truthfully, he gets the same collusion-proof policy.

The key to understanding this result lies in the pair of inequalities (12) and (13). Given that the contract deters collusion, the ¾ type provider cannot bene¯t from writing a side-contract after choosing the ¾ policy, $\{(®^{¾}_{h}; t^{¾}_{h}); (®^{¾}_{s}; t^{¾}_{s}; m^{¾})\}$. Also, the ¾ provider must not find it attractive to pick the τ policy $(\{(\alpha_h^\tau; t_h^\tau); (\alpha_s^\tau; t_s^\tau; m^\tau)\})$ — even if he can write a side-contract on it: see inequality (13). Now, the τ type provider prefers the τ policy to the σ policy — see inequality (12). But this must mean that when the τ type provider picks the τ policy, even if he could write side-contracts, he would be unable to benefit. Indeed, combining inequalities (13) and (12) yields:

$$(1-p)\alpha_h^\tau + p(\alpha_s^\tau - m^\tau c) \geq$$

$$(1-p)\max[\alpha_h^\tau; \alpha_s^\tau - m^\tau c + t_h^\tau - t_s^\tau] + p\max[\alpha_s^\tau - m^\tau c; \alpha_h^\tau - f(m^\tau) + t_s^\tau - t_h^\tau];$$

which says that the τ policy is collusion proof: the right-hand side is the expected utility for a provider when side-contracts can be written. Therefore, requiring that the σ policy to be collusion-proof implies that the τ policy must also be collusion-proof. When both policies are collusion-proof, obviously the likelihood of the provider being honest is not a determinant of the optimal contract. So the optimal collusion-proof contract simply consists of the policy in Proposition 1 to both provider types.

3.2.2 Optimal contract without collusion deterrence

In this section we consider contracts $[\{(\alpha_h^\tau; t_h^\tau); (\alpha_s^\tau; t_s^\tau; m^\tau)\}; \{(\alpha_h^\sigma; t_h^\sigma); (\alpha_s^\sigma; t_s^\sigma; m^\sigma)\}]$ for which the policy $\{(\alpha_h^\sigma; t_h^\sigma); (\alpha_s^\sigma; t_s^\sigma; m^\sigma)\}$ is not collusion-proof; the consumer and the collusive provider may profit from misreporting in equilibrium. Contracts that are not collusion-proof may be one of two classes. The first class consists of contracts for which in equilibrium the provider-consumer coalition always reports state h; the second, state s. Formally, the first class corresponds to contracts for which inequality (5) of definition 2 is violated; the second, inequality (4) is. Clearly, it is unnecessary to consider contracts of the first class because, without loss of generality, such a contract will never prescribe any recovery input for the collusive provider, and insurance breaks down completely; a collusion-proof contract dominates it.[14] Therefore, the joint surplus when the consumer-provider coalition reports state s must

[14] This is without loss of generality because if the contract did prescribe recovery input, the continuation equilibrium would be the same as one in the second class.

be greater than when it reports truthfully in state h, so we have:

$$\alpha_h^\sigma - t_h^\sigma < \alpha_s^\sigma - m^\sigma c - t_s^\sigma. \tag{16}$$

This inequality implies that

$$\alpha_h^\sigma - t_h^\sigma < \alpha_s^\sigma - m^\sigma c + f(m^\sigma) - t_s^\sigma$$

which says that the joint surplus is higher when the coalition reports truthfully in state s. So condition (16) is sufficient to ensure that the coalition always reports s.

Given the noncollusion-proof contract, in state h the provider offers a side-contract with the transfer x from the consumer to the provider defined by $t_s^\sigma = t_h^\sigma - x$. Therefore, the provider's utility in state h is $\alpha_s^\sigma - m^\sigma c - t_s^\sigma + t_h^\sigma$, and the consumer's utility in state h remains at $U(W - t_h^\sigma)$. In state s, the information is reported truthfully; hence the provider's and the consumer's utilities are $\alpha_s^\sigma - m^\sigma c$ and $U(W - \ell + f(m^\sigma) - t_s^\sigma)$, respectively. Therefore, the objective function is the same as in the previous section:

$$\theta[(1-p)U(W - t_h^\sigma) + pU(W - \ell + f(m^\sigma) - t_s^\sigma)] \tag{17}$$

$$+(1-\theta)[(1-p)U(W - t_h^\tau) + pU(W - \ell + f(m^\tau) - t_s^\tau)]$$

As before, the provider must obtain his reservation utility so the participation constraints are:

$$\alpha_h^\tau \geq 0 \tag{18}$$

$$\alpha_s^j - m^j c \geq 0 \quad j = \sigma; \tau. \tag{19}$$

Notice that since in equilibrium α_h^σ is not chosen by type σ provider in state h, we do not impose a lower bound on it. The budget constraint for the insurer is modified:

$$\theta(t_s^\sigma - \alpha_s^\sigma) + (1-\theta)[(1-p)(t_h^\tau - \alpha_h^\tau) + p(t_s^\tau - \alpha_s^\tau)] \geq 0. \tag{20}$$

When the provider's type is σ, the provider-consumer coalition always reports state s (recovery input will be used even when there has not been any loss); this explains the first term of (20), which says the transfer collected from the consumer is always t_s^σ while the payment to the provider is always α_s^σ.

Without loss of generality, the policy indexed by σ (resp. τ) is to be chosen by the collusive (resp. truthful) type. For the truthful provider to prefer the τ policy, we must have:

$$(21) \qquad (1-p)\alpha_h^\tau + p(\alpha_s^\tau - m^\tau c) \geq (1-p)\max[\alpha_h^\sigma, 0] + p(\alpha_s^\sigma - m^\sigma c).$$

Because we have not required $\alpha_h^\sigma \geq 0$, we allow for the possibility that type τ provider refuses to accept the transfer in state h if he has picked the σ policy. Next, using the information on the σ type provider's equilibrium utility from the optimal side-contract, we can write down his expected utility if he selects the σ policy:

$$(1-p)(\alpha_s^\sigma - m^\sigma c - t_s^\sigma + t_h^\sigma) + p(\alpha_s^\sigma - m^\sigma c).$$

The incentive constraint for type s provider therefore is:

$$(22) \qquad \alpha_s^\sigma - m^\sigma c + (1-p)(t_h^\sigma - t_s^\sigma) \geq$$

$$(1-p)\max[\alpha_h^\tau, \alpha_s^\tau - m^\tau c + t_h^\tau - t_s^\tau] + p\max[\alpha_s^\tau - m^\tau c, \alpha_h^\tau - f(m^\tau) + t_s^\tau - t_h^\tau]$$

Similar to (13), the above inequality takes into account the possibility of collusion when the σ type provider picks the policy that is meant for the τ type.

The contract $[\{(\alpha_h^\tau, t_h^\tau), (\alpha_s^\tau, t_s^\tau, m^\tau)\}, \{(\alpha_h^\sigma, t_h^\sigma), (\alpha_s^\sigma, t_s^\sigma, m^\sigma)\}]$ which maximizes (17) subject to constraints (18)-(22) is described in the following proposition.

Proposition 3 Given that the policy $\{(\alpha_h^\sigma, t_h^\sigma), (\alpha_s^\sigma, t_s^\sigma, m^\sigma)\}$ is not collusion-proof, the optimal contract has the following properties:

1. The truthful provider obtains zero profit: $\alpha_h^\tau = \alpha_s^\tau - m^\tau c = 0$.

2. In state s, the collusive provider reports k = s and obtains zero profit: $\alpha_s^\sigma - m^\sigma c = 0$; in state h, he reports k = s and obtains profit $t_h^\sigma - t_s^\sigma \geq 0$ through a side-payment from the consumer.

3. The consumer is imperfectly insured, whether she is matched with a truthful or a collusive provider: $W - \ell + f(m^j) - t_s^j < W - t_h^j$, $j = \sigma, \tau$.

4. When the provider is truthful, recovery input is excessive relative to the first best: $m^\tau > m^*$; when the provider is collusive, the recovery input is smaller than in the first best: $m^\sigma < m^*$.

The intuition for Proposition 3 is as follows. Given that the policy for the σ type permits a side-contract, the provider-consumer coalition always reports that a loss has occurred. As a result, recovery inputs are used even when a loss has not occurred. So prescribing recovery inputs for the collusive provider becomes more costly than in the first best. For this reason, the level of recovery input must be reduced from the first best. For the same reason, maintaining full insurance under a reduced recovery input is too costly, so a consumer who is matched with a collusive provider must face some risk. On the other hand, this misreporting incentive is not exploited by the truthful provider. Nevertheless, first-best risk-sharing is still not optimal for a consumer who is matched with a truthful provider. This is due to the fact that the collusive provider earns a rent. Indeed, this rent is $t_h^\sigma - t_s^\sigma$ which is equal to $t_h^\tau - t_s^\tau$ by the binding incentive constraint (22). Hence, to limit the rent to the collusive provider, $t_h^\tau - t_s^\tau$ is reduced from the first best. Given the lack of full insurance, it is optimal to raise the recovery input from the first-best level to reduce the amount of risk faced by the consumer with a truthful provider.

For the noncollusion-proof contract to lead to a side-contract, condition (16) must hold as a strict inequality.[15] For small values of θ (the likelihood of the collusive type being small), we can show that $t_h^\tau - t_s^\tau = t_h^\sigma - t_s^\sigma > 0$, so that (16) indeed is satisfied. Furthermore, in stark contrast with the collusion-proof contract of Proposition 2, the values of the variables in the τ-policy in proposition 3 vary with θ (this is readily seen by examining the first-order conditions). In fact, as θ goes to zero from above, the optimal policy for the truthful type tends towards the first-best policy.

[15] Recall that the provider reveals the information truthfully if he is indifferent between lying and telling the truth.

Corollary 1 As θ tends to 0, the optimal policy for the truthful type provider tends to the first-best policy: $\lim_{\theta \to 0}\{(\alpha_h^\tau; t_h^\tau); (\alpha_s^\tau; t_s^\tau; m^\tau)\} = \{(\alpha_h^*; t_h^*); (\alpha_s^*; t_s^*; m^*)\}$. Moreover, $t_h^\tau > t_s^\tau$ (so that (16) is satisfied).

Corollary 1 and Proposition 3 together say that the distortion of the τ policy changes in a strictly monotonic way as θ, the probability for a collusive provider, begins to increase from 0. When $\theta = 0$, the insurer offers the first-best policy to the truthful provider, the insurer bears all the risks, and production is efficient. As θ increases from 0, it is optimal to depart from the first best in order to reduce the collusive provider's rent and production inefficiency that result from the side-contract between the collusive provider and the consumer. This follows from the Envelope Theorem. Because risk sharing and production are efficient in the first-best policy, reducing the difference between t_s^τ and t_h^τ and increasing the recovery input m^τ slightly leads to a second-order loss, but this results in a first-order gain because the collusive provider's incentive constraint is relaxed. When θ begins to increase from 0, the equilibrium contract must begin to adjust. Therefore, although the first-best policy is feasible for the truthful provider, it is not offered in equilibrium. However, the equilibrium policy for the truthful provider must be approximately first-best when θ is in the neighborhood of zero. Note that this implies that allowing collusion must outperform deterring collusion when θ is small: the expected utility for the consumer must be higher when collusion is allowed than when it is not as θ becomes sufficiently small.

Next we turn to the case when θ is close to 1. Here, the rent to the collusive provider becomes large, as does the waste of the recovery input due to the misreporting in state h. Furthermore, the expected benefit to the consumer due to the truthful provider's behavior becomes small. The overall benefit from allowing collusion thus becomes small. Not surprisingly, when θ is large enough, it is better to deter collusion than to allow it.

Corollary 2 For all θ sufficiently close to 1, the optimal collusion contract in Proposition 3 has $t_h^\sigma - t_s^\sigma = t_h^\tau - t_s^\tau = 0$. That is, the optimal collusion contract satisfies the collusion-proofness constraints in Definition 2.

Corollary 2 implies that for θ sufficiently close to 1, the optimal noncollusion-proof policy must give the consumer a lower expected utility than the optimal collusion-proof policy. This means that for these values of μ, the equilibrium contract must be the collusion-proof contract in Proposition 2. Consequently, (and in contrast with the equilibrium contract for μ close to 0) for μ close to 1 the equilibrium contract does not vary with μ and must be the collusion-proof contract in Proposition 2. The intuition for why the equilibrium contract is independent of μ once it is close to 1 is this. Suppose μ is equal to 1, clearly, the equilibrium contract is collusion-proof. Now, let μ decrease from 1 slightly. If the contract is now changed slightly, it will not be collusion-proof any more: recall that with the collusion-proof contract, the collusive type is just indi®erent between writing a side-contract and not doing so. But now if the contract is no longer collusion-proof, a side-contract will be written and the collusive type will always report state s. This results in a waste of the treatment in state h. Put di®erently, if the policy is not collusion-proof for μ close to 1, then there is a discrete decrease in the payo® compared to the equilibrium payo®s at $\mu = 1$. There is indeed a discontinuity at $\mu = 1$. So relaxing the collusion-proof contract for μ close to 1 is suboptimal. As a result, the policy for the collusion type provider must remain collusion-proof for μ close to 1. To summarize, we have:

Proposition 4 When the probability of the provider acting collusively is su±ciently small, the equilibrium contract consists of policies that do not deter collusion; in equilibrium, the collusive provider writes a side-contract with the consumer, while the truthful provider is given a policy that is approximately ¯rst best. When the probability of the provider acting collusively is su±ciently high, the equilibrium contract consists of the equilibrium collusion-proof policy (independent of μ as long as μ is su±ciently close to 1) as if the provider were always collusive, namely the contract in Proposition 2.

We have been unable to show that the value of the objective function at the solution of the program for the optimal noncollusion-proof contract is monotonic in μ. So for intermediate values of μ (those not in the neighborhood of 0 or 1), we cannot characterize conditions for which the equilibrium contract is the collusion-proof contract in Proposition 2, or the noncollusion-proof contract in Proposition 3. Nevertheless, we suspect that the noncollusion-proof contract is the equilibrium contract if and only if μ is below a certain threshold.

4 Conclusion

In this paper, we have characterized the equilibrium insurance-payment contract under the assumption that some, but not all, providers may collude with consumers to ¯le false claims with the insurer. For this environment, we discover a basic tradeo® facing a contract designer. On the one hand, deterring collusion limits the rent the collusive provider may get, but leads to a high ine±ciency in risk sharing and production. On the other hand, allowing collusion mitigates the production and risk ine±ciency but increases the rent to the collusive provider. Allowing misreporting by the collusive provider is the only way to exploit the sincere behavior of the honest provider. We show that the equilibrium contract entails collusion and misreporting of information when the probability of the provider being collusive is su±ciently small.

Many of the assumptions of the model have been used to simplify the analysis. We have assumed that there are only two states of nature for the consumer (either a loss has occurred or not). Allowing for several levels of losses substantially complicates the analysis without adding to the important insights of the model. We are con¯dent that collusion will still be allowed in equilibrium when the probability of a collusive provider is su±ciently small. Nevertheless, the way collusion can happen in a model with many levels of losses is very complicated, because the number of ways information can be manipulated quickly becomes large.

We have simply assumed that there are some collusive and some honest economic agents in our static model. Clearly, in an environment where assymmetric information is important, contracting parties will prefer to transact with honest agents. There may also be an incentive for a principal to encourage agents to become honest, or for honest agents to signal their types to potential principals. These are very interesting issues for our future work.

Appendix

Proof of Proposition 1

Obviously, (9) is binding at the optimum; otherwise, t_s^σ and t_h^σ could be reduced. Also, (7) binds; otherwise, α_h^σ and t_h^σ could be reduced by the same amount. This would affect none of the other constraints but would increase the value of the objective function. An identical argument establishes that (8) binds. We have proven 1 in Proposition 1.

So with $\alpha_h^\sigma = 0 = \alpha_s^\sigma - m^\sigma c$, and with (9) binding, the optimal collusion-proof contract now is the solution of the following problem: choose t_h^σ, t_s^σ, and m^σ to maximize (6) subject to the following:

$$(1 - p)t_h^\sigma + p(t_s^\sigma - m^\sigma c) = 0 \tag{23}$$

$$-t_h^\sigma \geq -t_s^\sigma \tag{24}$$

$$f(m^\sigma) - t_s^\sigma \geq -t_h^\sigma . \tag{25}$$

Constraint (24) binds. If it did not, t_h^σ could be increased and t_s^σ decreased (in proportions such that (23) is unaffected). Such a change would not violate (25), but would decrease the difference between $W - t_h^\sigma$ and $W - \ell + f(m^\sigma) - t_s^\sigma$, since $f(m^\sigma) < \ell$. As U is concave, this would increase the expected utility. This proves 2 of the proposition. Now when $t_h^\sigma = t_s^\sigma$, 3 of the proposition follows from the assumption of $f(m) < \ell$.

Next, the binding constraint (24) implies that (25) is satisfied as a strict inequality, and that it can be ignored. Substituting the binding constraint (24) into (23), we have $t_h^\sigma = pm^\sigma c$. Putting this into the objective function, we have reduced the problem into the maximization of

$$(1 - p)U(W - t_h^\sigma) + pU(W - \ell + f(m^\sigma) - t_h^\sigma)$$

subject to

$$t_h^\sigma = pm^\sigma c.$$

Let λ be the Lagrange multiplier. The first-order conditions for t_h^σ and m^σ are, respectively:

$$(1-p)U'(W - t_h^\sigma) + pU'(W - \ell + f(m^\sigma) - t_h^\sigma) = \lambda$$

$$pU'(W - \ell + f(m^\sigma) - t_h^\sigma)f'(m^\sigma) = \lambda pc$$

Combining the above, we have

$$\frac{U'(W - \ell + f(m^\sigma) - t_h^\sigma)}{(1-p)U'(W - t_h^\sigma) + pU'(W - \ell + f(m^\sigma) - t_h^\sigma)} f'(m^\sigma) = c$$

Since $\ell < f(m^\sigma)$ and U is concave, the fraction in the above expression is greater than 1. Therefore $f'(m^\sigma) < c$. This proves 4 of the proposition. Q.E.D.

Proof of Lemma 1

Consider a contract $[\{(\alpha_h^\tau; t_h^\tau); (\alpha_s^\tau; t_s^\tau; m^\tau)\}; \{(\alpha_h^\sigma; t_h^\sigma); (\alpha_s^\sigma; t_s^\sigma; m^\sigma)\}]$, where $(\alpha_h^\tau; t_h^\tau) = (\alpha_h^\sigma; t_h^\sigma) = (0; t_h)$, $(\alpha_s^\tau; t_s^\tau; m^\tau) = (\alpha_s^\sigma; t_s^\sigma; m^\sigma) = (mc; t_s; m)$, and $t_h > t_s$. The σ policy is not collusion-proof: with $\alpha_h^\sigma = \alpha_s^\sigma - m^\sigma c = 0$, $t_h > t_s$ implies that inequality (4) in definition 2 is violated, so that the collusive provider will write a side-contract with the consumer and falsely report state s when the true state is h. Notice that inequality (5) does hold, so that the consumer and the provider will not misreport when the true state is s. Therefore in equilibrium the truthful provider gets 0 in states h and s, whereas the collusive provider reports truthfully if and only if the state is s, and gets 0 in state s and $t_h - t_s > 0$ in state h. Does there exist a contract $[\{(\hat\alpha_h^\tau; \hat t_h^\tau); (\hat\alpha_s^\tau; \hat t_s^\tau; \hat m_s^\tau)\}; \{(\hat\alpha_h^\sigma; \hat t_h^\sigma); (\hat\alpha_s^\sigma; \hat t_s^\sigma; \hat m_s^\sigma)\}]$ consisting only of collusion-proof policies that gives rise to the same equilibrium? If such a contract exists, it must have $\hat\alpha_h^\sigma = t_h - t_s$ (> 0) to give the same utility to the collusive provider in state h. But since at stage 2 the truthful provider picks the policy which maximizes his expected utility, he must also obtain $(1-p)(t_h - t_s) > 0$ from this new contract. This therefore must result in a different equilibrium. So such a collusion-proof contract cannot exist.[16] This is a contradiction. Q.E.D.

[16] We assume that the provider makes the side-contract offer, giving him all the potential gain from misreporting by the consumer-provider coalition. Lemma 1 does not depend on that assumption. Indeed, if the provider earns any strictly positive share of the gain from a side-contract, the lemma continues to hold.

Proof of Proposition 2

First, constraints (12) and (13) together imply:

$$\alpha_h^* - t_h^* \geq \alpha_s^* - m^* c - t_s^* \tag{26}$$

$$\alpha_s^* - m^* c + f(m^*) - t_s^* \geq \alpha_h^* - t_h^* \tag{27}$$

(which are equivalent to collusion-proofness constraints for the truthful provider). So adding these two constraints into the program is inconsequential.

Next, we relax the program by dropping (12) and (13); we will show that they are satisfied at the solution of the relaxed program. So the relaxed program is the maximization of (15) subject to (4), (5), (26), (27), (10), (11) and the budget constraint (14),

For the relaxed program, constraints (10) and (11) bind; this follows a similar argument as in the proof of Proposition 1. So we can substitute (10) and (11) as equalities to the other constraints and eliminate the α variables. Furthermore, the missing constraints (12) and (13) are satisfied once the constraints (10) and (11) have been shown to be binding. So we know that the solution to the relaxed program is the solution to the original problem. After having substituted for the α variables, we now rewrite all the constraints again:

$$-t_h^\sigma \geq -t_s^\sigma \tag{28}$$

$$f(m^\sigma) - t_s^\sigma \geq -t_h^\sigma. \tag{29}$$

$$-t_h^* \geq -t_s^* \tag{30}$$

$$f(m^*) - t_s^* \geq -t_h^*. \tag{31}$$

$$\mu[(1-p)t_h^\sigma + p(t_s^\sigma - m^\sigma c)] + (1-\mu)[(1-p)t_h^* + p(t_s^* - m^* c)] \geq 0 \tag{32}$$

Using a similar argument as in the proof of Proposition 1, constraints (28) and (30) are binding; as a result, constraints (29) and (31) are satisfied with slack. Now use (28) and (30) as equalities to substitute into (32) and the objective function. So we simplify the program into choosing t_h^σ, m^σ, t_h^τ, and m^τ to maximize

$$\theta[(1-p)U(W - t_h^\sigma) + pU(W - \ell + f(m^\sigma) - t_h^\sigma)] +$$

$$(1-\theta)[(1-p)U(W - t_h^\tau) + pU(W - \ell + f(m^\tau) - t_h^\tau)]$$

subject to

$$\theta(t_h^\sigma - pm^\sigma c) + (1-\theta)(t_h^\tau - pm^\tau c) \geq 0$$

The proposition follows from comparing the first-order conditions of this program to those in the proof of Proposition 1. Q.E.D.

Proof of Proposition 3.

We begin by assuming that for the τ policy $\alpha_h^\tau - t_h^\tau < \alpha_s^\tau - m^\tau c - t_s^\tau$, so that the collusive provider always reports s if he has chosen the policy meant for the truthful provider. We will later consider the other possibility. Then, constraint (22) reduces to:

$$\alpha_s^\sigma - m^\sigma c + (1-p)(t_h^\sigma - t_s^\sigma) \geq \alpha_s^\tau - m^\tau c + (1-p)(t_h^\tau - t_s^\tau) \tag{33}$$

We now show that the participation constraints (18) and (19) bind. First note that we can set $\alpha_h^\sigma = 0$. Then we show that (19) for $j = \tau$ is binding. Suppose that it does not bind. Then we can reduce α_s^τ, while increasing α_h^τ so as to leave the left-hand sides of (20) and (21) unaffected. Constraint (33) has been relaxed by the decrease of α_s^τ. Hence, t_h^σ can be decreased. Since t_h^σ is absent from the other constraints, these are not affected. This results in an increase of the expected utility. So (19) for $j = \tau$ must bind. Next, we show that (19) for $j = \sigma$ is binding. Suppose that it was not binding. Then, we can increase m^σ. For instance, increase $m^\sigma c$ by $\epsilon > 0$ and sufficiently small. This relaxes constraint (21), so that α_h^τ can be decreased by $\frac{p}{1-p}\epsilon$. By (20), t_h^τ can then also be decreased by $\frac{p}{1-p}\epsilon$. This in turn implies that the right-hand side of constraint (33) decreases by ϵ. Since the left-hand side was also decreased by ϵ, this constraint remains unaffected. But the expected utility has increased, since $m^{¾}$ has been increased, and $t^{¿}_{h}$ has been decreased. Therefore, (19) for j = ¾ is binding. Given that the constraints (19) for j = ¿ and for j = ¾ are binding, constraint (21) is redundant. Last, we show that constraint (18) must bind. Indeed, if it did not, $®^{¿}_{h}$ could be decreased. By (20), $t^{¿}_{h}$ can then be decreased, thus raising the expected utility. We have therefore proven 1 and 2.

Using (18) and (19) as equalities to substitute for the ® variables, we simplify the budget constraint (20) and the incentive constraint for the collusive provider (33) to:

µ(t

which says that $f'(m^\sigma) > c$, or $m^\sigma < m^*$.

Next we show that $m^\tau > m^*$: there is overproduction when the provider is truthful. First, adding (37) and (38) yields:

$$pU'(W - \ell + f(m^\tau) - t_s^\tau) + (1-p)U'(W - t_h^\tau) = \lambda \tag{41}$$

This and (36) imply:

$$\frac{U'(W - \ell + f(m^\tau) - t_s^\tau)}{pU'(W - \ell + f(m^\tau) - t_s^\tau) + (1-p)U'(W - t_h^\tau)} f'(m^\tau) = c \tag{42}$$

Getting an expression for λ from (38) and plugging it into (41), we obtain:

$$p[U'(W - \ell + f(m^\tau) - t_s^\tau) - U'(W - t_h^\tau)] = \frac{\mu}{1-\mu} U'(W - t_h^\sigma) > 0$$

So we conclude that $U'(W - \ell + f(m^\tau) - t_s^\tau) > U'(W - t_h^\tau)$. Therefore, (42) implies that $m^\tau > m^*$. Thus, we have proven 4.

Moreover, the last inequality implies that the consumer is imperfectly insured against the risk of losing ℓ when he is matched with a truthful provider: his net wealth is greater in case of no loss than in case of a loss. We now compare the differences of the consumer's net wealth between states h and s, when she is matched with a collusive provider, and when she is matched with a truthful provider. These differences are, respectively:

$$W - t_h^\sigma - W + \ell - f(m^\sigma) + t_s^\sigma = \ell - f(m^\sigma) + t_s^\sigma - t_h^\sigma;$$

and

$$W - t_h^\tau - W + \ell - f(m^\tau) + t_s^\tau = \ell - f(m^\tau) + t_s^\tau - t_h^\tau:$$

Since $m^\sigma < m^\tau$, and $t_s^\tau - t_h^\tau = t_s^\sigma - t_h^\sigma$, the former expression is greater than the latter. So we have

$$\ell - f(m^\sigma) + t_s^\sigma - t_h^\sigma > \ell - f(m^\tau) + t_s^\tau - t_h^\tau > 0;$$

where the last inequality follows from the fact that $U'(W - \ell + f(m^\tau) - t_s^\tau) > U'(W - t_h^\tau)$. So we have proven 3 of the proposition.

Finally, we consider the assumption of setting $®^{¿}_{h} - t^{¿}_{h} < ®^{¿}_{s} - m^{¿}c - t^{¿}_{s}$. So consider the opposite case: $®^{¿}_{h} - t^{¿}_{h} \geq ®^{¿}_{s} - m^{¿}c - t^{¿}_{s}$. In words, the joint surplus in state h is higher when the truth is reported. Without loss of generality, we can assume that in state s, a collusive provider will report truthfully if he has chosen the policy meant for type ¿: $®^{¿}_{s} - m^{¿}c + f(m^{¿}) - t^{¿}_{s} \geq ®^{¿}_{h} - t^{¿}_{h}$. In other words, this opposite case corresponds to the case of the ¿ policy being collusion-proof. The incentive constraint for the collusive provider (22) becomes:

$$(43) \qquad ®^{¾}_{s} - m^{¾}c + (1 - p)(t^{¾}_{h} - t^{¾}_{s}) \geq (1 - p)®^{¿}_{h} + p(®^{¿}_{s} - m^{¿}c)$$

It is straightforward to show that all the individual rationality constraints (18) and (19) bind. Hence, (43) becomes $t^{¾}_{h} - t^{¾}_{s} \geq 0$. Since this must be binding at any optimum (otherwise $t^{¾}_{h}$ could be decreased without a®ecting any constraint), the collusive provider's policy is in fact collusion-proof, contradicting the assumption that it is not collusion-proof, i.e., that $t^{¾}_{h} - t^{¾}_{s} > 0$.

Q.E.D.

Proof of Corollary 1

Setting µ = 0 in (36), (37), and (38), we have

$$U'(W - ` + f(m^{¿}) - t^{¿}_{s}) = U'(W - t^{¾}_{h})$$

and

$$f'(m^{¿}) = c:$$

Furthermore, the budget constraint becomes $(1-p)t^{¿}_{h} + p(t^{t}_{s} - m^{¿}c) = 0$, so that the policy is exactly ¯rst best. Obviously (16) is satis¯ed. The corollary follows from the upper-semicontinuity of the solutions to the constrained maximization program. Q.E.D.

Proof of Corollary 2

Let $t^{¿}_{h} - t^{¿}_{s} = t^{¾}_{h} - t^{¾}_{s} \equiv k \geq 0$. The program for Proposition 3 can be rewritten as the maximization of

$$µ[(1 - p)U(W - k - t^{¾}_{s}) + pU(W - ` + f(m^{¾}) - t^{¾}_{s})] +$$

$$(1-\theta)[(1-p)U(W-k-t_S^{\tau})+pU(W-\ell+f(m^{\tau})-t_S^{\tau})]$$

subject to

$$\theta(t_S^{\sigma}-m^{\sigma}c)+(1-\theta)(t_S^{\tau}+(1-p)k-pm^{\tau}c)=0:$$

The first-order conditions with respect to t_S^{σ}, t_S^{τ}, k, m^{σ} and m^{τ} are respectively:

$$(1-p)U'(W-k-t_S^{\sigma})+pU'(W-\ell+f(m^{\sigma})-t_S^{\sigma})=\lambda$$

$$(1-p)U'(W-k-t_S^{\tau})+pU'(W-\ell+f(m^{\tau})-t_S^{\tau})=\lambda$$

$$\theta(1-p)U'(W-k-t_S^{\sigma})+(1-\theta)(1-p)U'(W-k-t_S^{\tau})=\lambda(1-\theta)(1-p)$$

$$pU'(W-\ell+f(m^{\sigma})-t_S^{\sigma})f'(m^{\sigma})=\lambda c$$

$$U'(W-\ell+f(m^{\tau})-t_S^{\tau})f'(m^{\tau})=\lambda c:$$

As θ tends to 1, the first-order condition with respect to k must imply that $U'(W-k-t_S^{\sigma})$ tends towards 0. So for θ sufficiently close to 0, the value of k must be set at 0. Q.E.D.

References

Bowles, S. (1998). \Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions," Journal of Economic Literature, 36:75-111.

Alger, I. and R. Renault (1998). \Honest and Opportunistic Agents: Optimal Mechanisms without Information Revelation," mimeo, Boston University and GREMAQ Toulouse.

Alger, I. and F. Salanié (1997). \Claim In°ation in Health Insurance: A Rationale for Managed Health Care," mimeo, GREMAQ and INRA, Toulouse.

Crocker, K.J. and J. Morgan (1998). \Is Honesty the Best Policy? Curtailing Insurance Fraud through Optimal Incentive Contracts," Journal of Political Economy, 106:355-375.

Dawes, R.M. and R.H. Thaler (1988). \Anomalies: Cooperation," Journal of Economic Perspectives, 2:187-197.

Frank, R.H. (1987). \If Homo Economicus Could Choose His Own Utility Function, would He Want One with a Conscience?" American Economic Review, 77:593-604.

Kofman, F. and J. Lawarrée (1996). \On the Optimality of Allowing Collusion," Journal of Public Economics, 61:383-407.

Lacker, J.M. and J.A. Weinberg (1989). \Optimal Contracts under Costly State Falsi¯-cation," Journal of Political Economy, 97:1345-1363.

Laffont, J-.J. and J-.C. Rochet (1997). \Collusion in Organizations," Scandinavian Journal of Economics, xx:487-495.

Lu, Mingshan (1997). \Separating the `True E®ect' from `Gaming' in Incentive-Based Contracts in Health Care," mimeo, University of Calgary.

Ma, C.-T. A. and T. McGuire (1997). \Optimal Health Insurance and Provider Payment," American Economic Review, 87:685-704.

Maggi G. and A. Rodriguez-Clare (1989). \Costly Distortion of Information in Agency Problems," Rand Journal of Economics, 26:675-689.

Rabin, M. (1998). \Psychology and Economics," Journal of Economic Literature, 36:11-46.

Tirole J. (1986). \Hierarchies and Bureaucracies: On the Role of Collusion in Organizations," Journal of Law, Economics, and Organizations, 2:181-214.

Tirole, J. (1992). \Collusion and the Theory of Organizations,," in J.-J. La®ont, ed., Advances in Economic Theory: Proceedings of the Sixth World Congress of the Econometric Society. Cambridge; Cambridge University Press.151-206

Tirole, J. (1996). \A Theory of Collective Reputation (with applications to the persistence of corruption and to ¯rm quality)," Review of Economic Studies, 63:1-22.