



ELSEVIER

Journal of Economic Behavior & Organization
Vol. 50 (2003) 225–247

JOURNAL OF
Economic Behavior
& Organization

www.elsevier.com/locate/econbase

Moral hazard, insurance, and some collusion

Ingela Alger^{a,*}, Ching-to Albert Ma^{a,b,c}

^a *Department of Economics, Boston College, Chestnut Hill, Massachusetts, MA 02467, USA*

^b *Department of Economics, Boston University, Boston, Massachusetts, MA 02215, USA*

^c *Department of Economics, Hong Kong University of Science and Technology,
Clear Water Bay, Kowloon, Hong Kong*

Received 1 July 2000; received in revised form 12 January 2001; accepted 22 January 2001

Abstract

We consider a model of insurance and collusion. Efficient risk sharing requires the consumer to get a monetary compensation in case of a loss. But this in turn implies consumer–provider collusion incentives to submit false claims to the insurer. We assume, however, that only some providers are collusive while some are honest, and determine the optimal contract specifying the treatment, the insurance policy and the reimbursement policy to the provider. Two cases are analyzed: the first allows contract menus that induce self-selection among the providers; the second allows contracts consisting of a single policy. In both cases, deterrence of collusion is optimal only if the probability that the provider is collusive is large. When the contract deters collusion, it is as if the provider was collusive with probability one. The first best is achieved only when the provider is honest with certainty. Furthermore, over-consumption of treatment occurs in many cases, and is sometimes used as a substitute for monetary compensation to the consumer.

© 2003 Elsevier Science B.V. All rights reserved.

JEL classification: G22; I11; I12

Keywords: Insurance; Moral hazard; Optimal contracts; Honesty and dishonesty

1. Introduction

Insurance policies cover consumers for the losses they may suffer due to an accident. The moral hazard problem associated with insurance is well known. Because the true state of the loss may be hard to verify, instead of the indemnity approach, insurance policies often cover fully or partially the expenses consumers may have to insure to recover some of the losses. In disability insurance, the problem is compounded by the potential loss of income after an illness or accident. Disability and health insurance policies are very often simultaneously

* Corresponding author.

E-mail addresses: ingela.alger@bc.edu (I. Alger), ma@bu.edu (C.-t. Albert Ma).

offered by employers. Here, the moral hazard problem becomes more extensive: not only is there a potential for consumers to use medical care excessively, consumers and providers may misreport the extent of injury in order to collect disability insurance payments. In this paper, we study the collusion between consumers and providers in the context of health and disability insurance.¹

In our model, a risk-averse consumer may become ill, and suffer some health loss as well as income loss due to disability. A provider can partially restore these losses by supplying treatment. Insurance is offered by a risk-neutral insurer by means of contracts with the consumer and the provider. These contracts specify payments from the insurer to the provider, from the consumer to the insurer, as well as the amount of treatment. Because treatment only partially recovers the loss, efficient insurance requires a monetary compensation to the consumer.

Whether a loss has actually occurred or not is nonverifiable information, and payments and recovery inputs can only be based on the claim that the insurer receives. In considering how claims are filed, we use a new approach and let some economic agents act honestly. While some providers and consumers may collude to defraud the insurer of disability payments, some do not: these provider–consumer coalitions will report truthfully the private information of the loss.² The assumption that some providers behave honestly is motivated by several considerations. First, an economic agent may have made a conscious decision to be honest, concluding that the potential gain from lying may not be worthwhile. This is supported by several theoretical findings in the literature. The repeated-game literature has explained “sincere” or “altruistic” behavior as optimal responses in long-term interactions. This of course cannot explain all types of cooperation, in particular those that occur in short-term relationships. But honest behavior may also be optimal in short-term relationships.³ An agent may have such a distaste for misreporting the truth that he will forgo any potential gain.

Second, a number of papers have suggested that human behavior is driven by factors other than pure selfishness. Surveying the psychology literature, Rabin (1998) concludes that the standard assumption of selfish behavior in economics may be too narrow, and sometimes even misleading. Extensive experimental evidence in bargaining games has demonstrated that a notion of fairness determines behavior (see Rabin, 1998, p. 21ff). 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, p. 80) “treat[s] preferences as cultural traits, or learned influences of behavior” and includes such examples as never lying and reciprocating dinner invitations. Taken together,

¹ Clearly, our model applies to any environment in which an accident leads to both injury and loss of income, such as home and automobile insurance. We have chosen to use the health and disability example because it illustrates the situation most succinctly.

² Ma and McGuire (1997) also study claims as a result of collusion between the provider and the consumer, but assumes that there are only collusive provider types.

³ In a model where an agent is matched with a new principal each period, Tirole (1996) 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.

these findings strongly indicate that honest behavior is neither uncommon nor irrational. In the tax compliance and evasion literature, it has been recognized that some taxpayers are always honest; Erard and Feinstein (1994) have studied the impact of this.

Our approach does *not* assume that economic agents are nonstrategic. Although honest providers always truthfully file claims with the insurer, they seek to maximize profits, subject to the truthful report of the loss. We have thought about an alternative assumption: that some agents are completely nonstrategic, honestly reporting the health status, as well as always performing any task that is asked of them to do (for example, returning any rent or payoffs above their reservation utility level). We believe that this second assumption is too strong, and hard to justify empirically. Moreover, a consumer's health status has a certain objective connotation, being a piece of information that is potentially verifiable; in fact, we will exploit this property in setting up the collusion subgame between the consumer and the provider. On the other hand, an agent's intention is impossible to ascertain objectively; even under the best of circumstances, one's intention may not be inferred accurately.

We determine the equilibrium insurance and provider payment contracts. As a benchmark, we first assume that the insurer can observe whether the provider is honest or not. If he is honest, the first best is implementable. If he is collusive, the contract is second-best: consumers must bear some risk. Interestingly, the provider does not get any rent: collusion is deterred by eliminating the monetary compensation to the consumer when a loss is claimed. In fact, the treatment is used as a substitute for monetary compensation: it is larger than the first-best level.

When the insurer cannot observe whether the provider is honest or not, we consider the general class of deterministic contracts, each of which consists of a menu of insurance and payment policies and recovery inputs. However, the common methodology of solving for the equilibrium by maximizing an objective function under 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. For completeness, we show in Lemma 1 that some equilibrium allocations cannot be implemented by collusion-proof contracts.⁴ Therefore, 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, and identical to the equilibrium policy in a model where the provider is known to be collusive always (Proposition 2).⁵

This clearly demonstrates that deterring collusion cannot be always optimal. For example, when the likelihood of a collusive provider is very small, a first-best contract, which is not collusion-proof, must perform better. Nevertheless, we show that the first-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

⁴ An honest provider always forgoes any gain from information manipulation. A contract allowing collusion, therefore, rewards *only* the collusive type. Therefore, equilibrium allocations that reward the two types of provider differently must imply equilibrium collusion.

⁵ 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.

on the provider's type. Collusion takes the form of the provider–consumer coalition always claiming that a loss has occurred. This leads to a waste in recovery inputs and rents for the provider. The optimal noncollusion-proof policy for the honest provider implements excessive recovery inputs as a substitute for monetary compensation; for the collusive provider, insufficient inputs are used, because they are wasted with some probability. Risk sharing for the consumer is never perfect. More importantly, when compared to the first 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 one. This contrasts with the optimal collusion-proof contract, which is independent of the probability. 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.

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 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.

Our analysis of side-contracts follows Tirole's innovation (Tirole, 1986), but adds the assumption that the provider may not always be willing to manipulate information. Hence, the method of considering only collusion-proof equilibria does not apply. 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) and Tirole (1992) also study models with collusion where some economic agent may not be corrupt. In both papers, it is assumed that the principal offers a single contract; in contrast we allow for a menu of contracts. We find that the contracts meant for the honest and the dishonest provider are different whenever collusion is not deterred. Our results can, therefore, be interpreted to imply that the analyses of Kofman and Lawarrée and Tirole involve a loss of generality.

Alger and Renault (1998, 2000) consider a principal-agent model (without collusion) in which the agent has private information about the surplus generated by the economic exchange with principal (the circumstances). The agent either discloses the information to the principal honestly or he must be induced to do so. Alger and Renault (1998) is closest to the present paper in that two periods are considered: the agent knows only his ethics in the first period and learns the circumstances in the second period.⁶ In the first period, the principal offers a menu of contracts that specify allocations as a function of the circumstances announced by the agent in the second period. Some of our results are similar to theirs: in both papers, the standard second-best contract is optimal when dishonest behavior is sufficiently likely, and ethics are truthfully revealed in the first period; moreover, it is optimal to let

⁶ Alger and Renault (2000) allow for contracts in which the agent announces his ethics and the circumstances simultaneously; they also consider three alternative definitions of honesty, depending on whether the honest agent honestly reveals his ethics or not.

the dishonest type lie in equilibrium when honest behavior is sufficiently likely. This is comforting since it indicates that these results are robust. However, our findings differ for the contracts offered when honest behavior is sufficiently likely. They find that the first-best contract is offered to the honest agent (and to the dishonest agent if the principal's utility does not directly depend on the agent's type, as is the case here). In contrast, we find that the honest and dishonest providers are offered different contracts; moreover, these contracts are never first best. The reason behind this difference will be explained in the text.

The next section presents the model and the first best; we also define the side-contract and the extensive form. The analysis is presented in Section 3. First, we define 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. For completeness, in Section 4, we derive the optimal contract when menus are unavailable. We show that this involves a loss of generality. Finally, concluding remarks are made in the last section. Proofs are contained in the appendix unless otherwise stated.

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 first best. Then we define 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.⁷ The consumer suffers a loss with probability p . The loss can be due to an illness (as in the health insurance case) or an accident (the automobile insurance case). We express the loss in monetary terms and denote it by ℓ . There are, thus, two possible states of nature.⁸ We use the index $i = h, s$ to denote them: $i = s$ is the state when the consumer's loss is ℓ , otherwise, $i = h$. In the health-care market, a consumer may suffer from some symptoms. If those symptoms are actually insignificant, the true state of nature is h (the "healthy" state); otherwise, the state is s (the "sick" state). In the automobile insurance market, the states of nature refer to whether an accident has actually led to a loss. The consumer is risk averse with an increasing and strictly concave utility function U , defined 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. The recovery technology is, however, completely unproductive when the consumer has not experienced a loss. Let m denote the input. 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 through the function f , where f is increasing and strictly concave. We assume that f is bounded from above, and that $\max f(m) < \ell$; in other words, the technology cannot completely recover the consumer's total loss ℓ . Many situations fit this assumption. This is certainly true for illnesses that may

⁷ See Lu (1997) for an example of evidence of misreporting in the health care sector.

⁸ We will introduce another set of states later; they concern the strategic interaction between a consumer and a provider.

lead to disability and loss of income. Moreover, consumers may not be able to recover fully from some illnesses. For the automobile market, it reflects the inconvenience caused by a repair of the car after an accident, the time and effort required to find a replacement, or the unavoidable risk that a repair might have been done improperly.⁹ We assume that the recovery input is verifiable information, and that it can be specified in a contract. The provider is risk neutral, but has limited liability; we normalize his reservation profit to 0.

An insurer offers a policy comprising insurance for the consumer and a reimbursement scheme for the provider. The policy specifies a transfer t_i from the consumer to the insurer, a transfer α_i from the insurer to the provider, and the amount of recovery input m . An insurance-reimbursement policy can, thus, be denoted by $\{(\alpha_h, t_h); (\alpha_s, t_s, m)\}$. If the consumer suffers a loss, and pays an amount t in order to obtain m units of recovery, her utility is $U(W - \ell + f(m) - t)$; if there is no loss and the consumer pays t , her utility is $U(W - t)$. If the provider supplies m units of recovery in return for a remuneration α , his utility is $\alpha - mc$. We suppose that the provider can always refuse to serve the consumer, in which case he obtains his reservation utility. The risk-neutral insurer is assumed to operate in a competitive market, and to set its policy to maximize the consumer's expected utility.

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 is denoted by $\{(\alpha_h^*, t_h^*); (\alpha_s^*, t_s^*, m^*)\}$, and is the solution to the following program: choose $\alpha_h, \alpha_s, t_h, t_s$ and m to maximize

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

subject to

$$\alpha_h \geq 0,$$

$$\alpha_s - mc \geq 0,$$

$$(1 - p)(t_h - \alpha_h) + p(t_s - \alpha_s) \geq 0.$$

The first two constraints are the participation constraints for the provider, and the third one is the insurer's budget constraint. Trivially, the insurer sets $\alpha_h = 0$ and $\alpha_s = mc$. Then, straightforward calculation from the first-order conditions yields:

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

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

The first condition says that all risks are absorbed by the insurer. The 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: $t_h^* > t_s^*$.¹⁰

Having established the first best, we proceed to study the interaction between the provider and the consumer when they can manipulate the information about the consumer's loss. While the information about the true state of nature may be obtained at a cost, frequently the insurer relies on the information reported by the provider or the insured. This is clearly the

⁹ Alternatively, we can assume that the marginal cost of recovery, c , is sufficiently high.

¹⁰ Alternatively, one can interpret this lesser payment in the event of a loss as a refund of the premium. That is, the consumer pays the premium t_h upfront, and if there is a loss, $t_h - t_s$ part of the payment will be refunded.

case in the health-care market; although medical records are available and potentially can be used to settle disputes and law suits, almost always a provider simply submits a claim on behalf of the consumer to complete the financial transactions due to the treatments.

As we noted in the previous section, the economic relationship between a provider and a consumer may develop in an uncertain way: while some providers may collude with consumers, others may not. For our purpose of analyzing the effects of collusion, it is unimportant to model explicitly the reasons behind this, although one can easily think of a number of reasons. For example, there may be a chance that collusion results in a prosecution and penalty, and due to differences in wealth or opportunity costs, some providers find collusion unattractive. Alternatively, some providers may have developed a long-term relationship with the insurer and care more about their reputation.

The provider is, therefore, either truthful or collusive; denote these types by the index $j = \tau$ and σ , respectively. The type of the provider is not observable by the insurer. The probability that a provider is collusive is θ . For simplicity, we assume that the consumer colludes whenever the provider is willing to. A truthful provider always submits a truthful report about the consumer's loss. By contrast, a collusive provider may agree to submit a false claim through a side-contract with the consumer, which we now define:

Definition 1 (The side-contract subgame). A side-contract is an offer (x, k) made by the provider to the consumer, where x is a transfer from the consumer to the provider, 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.

Observe that any party in the side-contract subgame can unilaterally enforce truthful reporting. If the true state is i , the provider can simply offer the consumer a transfer of 0 and a report of i ; likewise, a consumer can always reject a side-contract offer from the provider, and the true state will be reported. Hence, for the consumer and provider to agree on a (nontrivial) side-contract, both of them must be at least as well off as when reporting the true state.¹¹

The transfer allows the provider–consumer coalition to maximize their joint surplus. To illustrate, suppose that the true state is h . If they report the true state, the joint surplus is $\alpha_h - t_h$; if they report s , it becomes $\alpha_s - mc - t_s$. Note that since treatment is verifiable, it is carried out although it has no value to the consumer. Suppose that

$$\alpha_s - mc - t_s > \alpha_h - t_h. \quad (3)$$

Then there must exist an x such that

$$t_h - t_s \geq x \geq \alpha_h - \alpha_s + mc,$$

one of the inequalities being strict. Both parties can be made better off by the following side-contract: they report $k = s$ and the consumer pays x to the provider, where $-t_s - x \geq -t_h$, and $\alpha_s - mc + x \geq \alpha_h$. Since the provider makes the side-contract offer, he can extract the

¹¹ As a tie-breaking rule, we assume that if neither the provider nor the consumer gains strictly by misreporting the state, the true state is reported.

whole surplus generated by the misreporting. Thus, if (3) holds, the provider offers a side-contract with the transfer $x = t_h - t_s$, implying a net payoff $\alpha_s - mc + t_h - t_s$ to the provider.

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. We will consider two cases: first, the contract contains two policies for the provider to choose between; second, it consists of one single policy. In stage 2, the provider picks a policy (we will describe this stage in more detail later). 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.

If the provider always reports truthfully ($\theta = 0$), the insurer can rely on the provider’s information and implement the first best. In contrast, if the provider is willing to write side-contracts with a consumer with some probability ($\theta > 0$), the first best is not implementable. Consider the first-best allocation: under state h , there is no recovery input, and the consumer pays t_h^* ; in state s , the consumer receives treatment m^* , pays t_s^* and the provider is reimbursed m^*c . From (1) and (2), we know that $t_h^* > t_s^*$. The provider’s net payoff is independent of the state and is equal to 0. This creates an incentive for the consumer–provider coalition to avoid the higher payment in state h by reporting state s when the true state is h . They will waste the recovery input, but this is not costly, since the provider is completely reimbursed. Explicit attention must, therefore, be paid to the possibility of the formation of side-contracts. As a benchmark, we begin our analysis with the standard assumption, namely, the provider is collusive with probability 1. Results under this assumption will then be compared with those in the more general case ($0 < \theta < 1$), which we analyze subsequently.

When a provider is always collusive, the insurer offers a single policy: $\{(\alpha_h^\sigma, t_h^\sigma); (\alpha_s^\sigma, t_s^\sigma, m^\sigma)\}$. In this case, it can easily be proved that there is no loss of generality in restricting attention to contracts which deter equilibrium side-contracts between the provider and the consumer (see Tirole (1986)).

Definition 2 (Collusion-proof policies). A policy $\{(\alpha_h^\sigma, t_h^\sigma); (\alpha_s^\sigma, t_s^\sigma, m^\sigma)\}$ is said to be collusion-proof if the following conditions are satisfied:

$$\alpha_h^\sigma - t_h^\sigma \geq \alpha_s^\sigma - m^\sigma c - t_s^\sigma, \quad (4)$$

$$\alpha_s^\sigma - m^\sigma c + f(m^\sigma) - t_s^\sigma \geq \alpha_h^\sigma - t_h^\sigma. \quad (5)$$

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. For later use, it is useful to note that both constraints cannot be violated simultaneously: if (4) is violated, then (5) is satisfied, and vice versa.

The insurer maximizes the expected utility of the consumer:

$$(1 - p)U(W - t_h^\sigma) + pU(W - \ell + f(m^\sigma) - t_s^\sigma), \tag{6}$$

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:

$$\alpha_h^\sigma \geq 0, \tag{7}$$

$$\alpha_s^\sigma - m^\sigma c \geq 0, \tag{8}$$

$$(1 - p)(t_h^\sigma - \alpha_h^\sigma) + p(t_s^\sigma - \alpha_s^\sigma) \geq 0. \tag{9}$$

Proposition 1. *When the provider is always collusive, the optimal collusion-proof policy $\{(\alpha_h^\sigma, t_h^\sigma); (\alpha_s^\sigma, t_s^\sigma, m^\sigma)\}$ 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 policy in Proposition 1 will be referred to as *the standard second-best policy*. It has two main features: insurance is imperfect, and the recovery input is excessive compared to the first best. The intuition is as follows. As we pointed out above, given the first best, the consumer and the provider have an incentive to report s when the true state is h. The relevant (and binding) collusion-proofness constraint is, 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 difference between the transfers of the consumer to the insurer $t_h^\sigma - t_s^\sigma$. The first option is suboptimal: 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 transfer to the insurer by the same amount. Hence, the provider must get zero profit, and the optimal way to deter collusion is to reduce $t_h^\sigma - t_s^\sigma$ to zero. In other words, risk sharing through differences in monetary transfers across states h and s is not optimal, and the recovery input is used as a substitute: it is increased from the first-best level, $m^\sigma > m^*$. Now we turn to the case where the provider colludes with the consumer with probability θ , $0 < \theta < 1$.

3. Self-selection among providers

In this section, we assume that the insurer offers a contract consisting of a menu of two policies: $\{(\alpha_h^\tau, t_h^\tau); (\alpha_s^\tau, t_s^\tau, m^\tau)\}, \{(\alpha_h^\sigma, t_h^\sigma); (\alpha_s^\sigma, t_s^\sigma, m^\sigma)\}$. In stage 2, with the knowledge of whether he is truthful or collusive, the provider picks a policy $\{(\alpha_h^j, t_h^j); (\alpha_s^j, t_s^j, m^j)\}$, $j = \tau, \sigma$. We assume that, independently of his type, the provider picks the policy that 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. The interpretation in the health insurance case is straightforward: although a physician may have a high ethical standard while working (i.e. report the truth and avoid prescribing treatment when unnecessary), he chooses the highest expected payment for doing so.

The two policies may be distinct and provide incentives for different types to self-select in stage 2 of the game. But of course, the 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 $\{(\alpha_h^j, t_h^j); (\alpha_s^j, t_s^j, m^j)\}$, $j = \tau, \sigma$.

A first and crucial observation is that when the provider colludes with the consumer with probability θ , it is no longer possible to consider only collusion-proof policies. Consider a contract $\{[(\alpha_h^\tau, t_h^\tau); (\alpha_s^\tau, t_s^\tau, m^\tau)], [(\alpha_h^\sigma, t_h^\sigma); (\alpha_s^\sigma, t_s^\sigma, m^\sigma)]\}$ that defines a subgame at stage 2, and a continuation equilibrium, and suppose that $\{(\alpha_h^\sigma, t_h^\sigma); (\alpha_s^\sigma, t_s^\sigma, m^\sigma)\}$ is *not* collusion-proof. Then the continuation equilibrium of the subgame may not be replicated by a collusion-proof policy for the collusive type.

If $\{(\alpha_h^\sigma, t_h^\sigma); (\alpha_s^\sigma, t_s^\sigma, m^\sigma)\}$ is not collusion-proof, the collusive provider and the consumer will find it optimal to lie in some state. At the beginning of stage 2, the collusive provider then anticipates that his expected utility, denoted x , is strictly greater than his expected utility if he was to report truthfully: $x > y \equiv (1 - p)\alpha_h^\sigma + p(\alpha_s^\sigma - m^\sigma c)$. Because a type τ provider does not write a side-contract, his expected utility is just y if he picks the σ policy. Of course the τ policy will have to give him at least y in expected utility; say it gives him exactly y . Now consider a collusion-proof contract that mimics the outcome of the noncollusion-proof contract and, therefore, gives the same expected utilities to types σ and τ of the provider. Since the new policy is collusion-proof, the type σ gets expected utility x by reporting truthfully. So the type τ provider will also get expected utility x in equilibrium (he can always pick the policy meant for type σ in stage 2). Therefore, the collusion-proof policy cannot mimic the equilibrium of the subgame of the noncollusion-proof policy. We have, thus, proved:¹²

Lemma 1. *For $\theta \in (0, 1)$, a (continuation) equilibrium of a contract may not be a (continuation) equilibrium of a contract consisting only of 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 profit to the collusive than to the honest provider. The possibility of "differentially" rewarding different types of the provider through a side-contract by the collusive type is the key to our analysis.¹³

Lemma 1 implies that we must consider contracts consisting of policies that may not be collusion-proof. In most models in information economics, an incentive problem typically

¹² Alger and Renault (1998) contains a similar argument in a different model.

¹³ Lemma 1 does not depend on the assumption that the provider gets all the potential gain from misreporting. Indeed, if the provider earns any strictly positive share of the gain from a side-contract, the lemma continues to hold.

results in deterrence of misreporting of information. In contrast, our analysis must consider allocations that are generated by the misrepresentation of information. Again, this is due to our assumption that a player in our model may not want to exploit the superior information he has in some subgame. In the next two subsections, we analyze two classes of contracts: those which prevent collusion, and those which do not.

3.1. Optimal contract with collusion deterrence

Here we derive the optimal contract $\{[(\alpha_h^\tau, t_h^\tau); (\alpha_s^\tau, t_s^\tau, m^\tau)], [(\alpha_h^\sigma, t_h^\sigma); (\alpha_s^\sigma, t_s^\sigma, m^\sigma)]\}$ that deters collusion. In stage 4, the provider of type σ must not find it profitable to write a side-contract, so constraints (4) and (5) must hold. Since the policies must allow the provider to obtain his reservation utility in both states, we have the following participation constraints:

$$\alpha_h^j \geq 0, \quad j = \sigma, \tau, \tag{10}$$

$$\alpha_s^j - m^j c \geq 0, \quad j = \sigma, \tau. \tag{11}$$

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 find it optimal to select policy $\{(\alpha_h^j, t_h^j); (\alpha_s^j, t_s^j, m^j)\}$, $j = \tau, \sigma$, anticipating the equilibrium moves in stage 4. Type τ provider always reveals the true state in stage 4, regardless of the policy he has chosen. Therefore, the incentive constraint for the truthful provider type is the following:

$$(1 - p)\alpha_h^\tau + p(\alpha_s^\tau - m^\tau c) \geq (1 - p)\alpha_h^\sigma + p(\alpha_s^\sigma - m^\sigma c), \tag{12}$$

The condition for a collusive provider to choose the σ policy is slightly more involved. Although the policy indexed by σ is collusion-proof, the policy indexed by τ may not be. When a collusive provider selects the τ policy, he will consider any gain from a side-contract at stage 4. Thus, if the state turns out to be h , he can either forgo the side-contract to obtain α_h^τ or use a side-contract to report s and get $\alpha_s^\tau - m^\tau c + t_h^\tau - t_s^\tau$. Similarly, if the state is s , the provider reports $k = s$ (respectively, $k = h$) when $\alpha_s^\tau - m^\tau c$ is greater (respectively, smaller) than $\alpha_h^\tau - f(m^\tau) + t_s^\tau - t_h^\tau$. To summarize, the incentive constraint that ensures that type σ provider picks the σ policy is:

$$(1 - p)\alpha_h^\sigma + p(\alpha_s^\sigma - m^\sigma c) \geq (1 - p)\max[\alpha_h^\tau, \alpha_s^\tau - m^\tau c + t_h^\tau - t_s^\tau] + p\max[\alpha_s^\tau - m^\tau c, \alpha_h^\tau - f(m^\tau) + t_s^\tau - t_h^\tau]. \tag{13}$$

Finally, the insurer’s budget constraint is:

$$\theta[(1 - p)(t_h^\sigma - \alpha_h^\sigma) + p(t_s^\sigma - \alpha_s^\sigma)] + (1 - \theta)[(1 - p)(t_h^\tau - \alpha_h^\tau) + p(t_s^\tau - \alpha_s^\tau)] \geq 0, \tag{14}$$

and the objective function is:

$$\theta[(1 - p)U(W - t_h^\sigma) + pU(W - \ell + f(m^\sigma) - t_s^\sigma)] + (1 - \theta)[(1 - p)U(W - t_h^\tau) + pU(W - \ell + f(m^\tau) - t_s^\tau)]. \tag{15}$$

An optimal contract deterring collusion maximizes (15) subject to (4), (5), (10)–(14).

Proposition 2. *Suppose $0 < \theta < 1$. The optimal contract deterring equilibrium collusion is independent of θ and offers the same policy to the truthful and collusive types of provider. In fact, this policy is the standard second-best policy of Proposition 1.*

Proposition 2 says that if the contract must deter collusion, then the existence of the truthful type of provider is inconsequential: the equilibrium contract is the same as if the provider was collusive with certainty. Despite the fact that the τ type provider always reports truthfully, he gets the optimal 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 σ type provider cannot benefit from writing a side-contract after choosing the σ policy, $\{(\alpha_h^\sigma, t_h^\sigma); (\alpha_s^\sigma, t_s^\sigma, m^\sigma)\}$. Also, the collusive 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 truthful provider prefers the τ policy to the σ policy—see inequality (12). But this must mean that when the truthful 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 the σ policy to be collusion-proof implies that the τ policy must also be collusion-proof. Hence, the optimal collusion-proof contract simply consists of the policy in Proposition 1 to both provider types.

3.2. Optimal contract without collusion deterrence

In this subsection, 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 σ policy is *not* collusion-proof. As we noted in the discussion of Definition 2, under a noncollusion-proof policy the provider–consumer coalition either always reports state h or always reports state s. We will consider only those policies in which the coalition always reports s, and verify later that this is optimal. So consider those σ policies for which the joint surplus from reporting state s is greater than reporting state h truthfully:

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

In state h the collusive provider offers a side-contract with the transfer $x = t_h^\sigma - t_s^\sigma$ from the consumer to the provider. Therefore, the provider’s utility is $\alpha_s^\sigma - m^\sigma c + t_h^\sigma - t_s^\sigma$, and the consumer’s utility 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. The objective function is the same as in the previous section:

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

As before, the provider must obtain his reservation utility:

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

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

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 recovery input will be used with probability 1; this explains the first term of (20), which says that the transfer collected from the consumer is always t_s^σ while the payment to the provider is always α_s^σ .

Finally, the self-selection constraints for the provider are as follows. For the truthful provider:

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

Because we have not required $\alpha_h^\sigma \geq 0$, we allow for the possibility that the truthful provider refuses the transfer in state h if he has picked the σ policy. Next, using the information on the collusive provider’s equilibrium utility from the optimal side-contract, we can determine 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 σ provider is:

$$\begin{aligned} &\alpha_s^\sigma - m^\sigma c + (1 - p)(t_h^\sigma - t_s^\sigma) \\ &\geq (1 - p)\max[\alpha_h^\tau, \alpha_s^\tau - m^\tau c + t_h^\tau - t_s^\tau] + p \max[\alpha_s^\tau - m^\tau c, \alpha_h^\tau - f(m^\tau) + t_s^\tau - t_h^\tau]. \end{aligned} \tag{22}$$

Similar to (13), the right-hand side takes into account the possibility of collusion when the collusive provider picks the policy that is meant for the truthful provider.

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 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: $m^\sigma < m^*$.*

Given that the collusive provider always reports s ,¹⁴ recovery inputs are used even when they have no value. For this reason, the level of recovery input must be reduced from the first best when the consumer is matched with a collusive provider. And 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. The truthful provider does not exploit the misreporting incentive. 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 $t_h^\sigma - t_s^\sigma$ 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 matched with a truthful provider. This is reminiscent of the regular second-best policy, where treatment is also used as a substitute for monetary compensation (see Proposition 1).

As mentioned in the introduction, the result in the above proposition differs from that in Alger and Renault (1998): given that the principal's utility does not depend directly on the agent's type, they find that the principal offers the first-best contract to both the honest and the dishonest agents when the dishonest agent's type is not revealed in equilibrium. In their model, when the dishonest agent lies in equilibrium, it is as if the dishonest agent was of one particular type with probability 1. As a result, the principal sets the decision variable to its first-best level conditional on that type: offering the first-best contract also to the honest agent then does not violate the dishonest agent's incentive constraint. In our model it is not optimal to do so, because of the insurance motive: even though the dishonest provider always claims that the patient is in the sick state, the consumer's utility is not the same in the healthy and the sick states. Moreover the treatment is unproductive when the consumer is healthy. Therefore, it is not optimal to set the treatment level and the payments to their first-best levels when the consumer is matched with a collusive provider; in order to satisfy the incentive constraint (22) the contract for the honest provider then also has to be distorted.

3.3. The equilibrium contract

To determine the equilibrium contract, we compare the consumer's expected utilities from the optimal collusion-proof and noncollusion-proof contracts. In 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.

¹⁴ Note that condition (16) must hold for this to be true. 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. Note further that a simple revealed preference argument can be used to verify that the insurer could not do better by offering some contract $\{(\alpha_h^\tau, t_h^\tau); (\alpha_s^\tau, t_s^\tau, m^\tau)\}, \{(\alpha_h^\sigma, t_h^\sigma); (\alpha_s^\sigma, t_s^\sigma, m^\sigma)\}$ such that the collusive provider always reports h . Note that for such a contract, no recovery input is used when the consumer is in state s and is matched with a collusive provider. Hence, it must be dominated by the one considered above: indeed, in the contract above, $m^\sigma = 0$ could have been chosen, but was not.

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

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 of 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 contract in Proposition 3 has $t_h^\sigma - t_s^\sigma = t_h^\tau - t_s^\tau = 0$. That is, it 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. Therefore, for these values of θ , the equilibrium contract must be the collusion-proof 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 θ . So for intermediate values of θ (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 θ is below a certain threshold. To summarize our results, we have:

Proposition 4. *When the probability of the provider acting collusively is sufficiently 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 first best. When the probability of the provider acting collusively is sufficiently high, the equilibrium contract consists of the equilibrium collusion-proof policy (independent of θ) as if the provider were always collusive, namely the contract in Proposition 1.*

The intuition for why the equilibrium contract is independent of θ once it is close to 1 is this. Suppose θ is equal to 1, clearly, the equilibrium contract is collusion-proof. Now, let θ decrease from 1 slightly. If the contract is now changed slightly, so that it 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 differently, changing from the collusion-proof to a noncollusion-proof policy for θ close to 1 results in a discrete decrease in the payoff. So departing from the collusion-proof contract for θ close to 1 is suboptimal.

4. Two types of provider, one single policy

For completeness and comparison, we now consider a regime in which the insurer offers one single policy, $\{(\alpha_h, t_h); (\alpha_s, t_s, m)\}$. Screening menus are ruled out in this section. The analysis turns out to be much simpler. Clearly, the policy must be either collusion-proof or not. If it is collusion-proof, the optimal policy is the standard second-best policy. We only need to determine the optimal noncollusion-proof policy, and again, we only need to consider those policies where the consumer–provider coalition always reports s . So condition (16) must hold. The consumer’s expected utility is independent of the provider type and it is:

$$(1 - p)U(W - t_h) + pU(W - \ell + f(m) - t_s). \tag{23}$$

The participation constraints for the provider and the insurer’s budget constraint are:

$$\alpha_h \geq 0, \tag{24}$$

$$\alpha_s - mc \geq 0, \tag{25}$$

$$(1 - \theta)(1 - p)(t_h - \alpha_h) + [\theta + (1 - \theta)p](t_s - \alpha_s) \geq 0. \tag{26}$$

All three constraints are binding, and they can be solved to replace α_h , α_s and t_h in the objective function. The first-order conditions with respect to m and t_s are, respectively:

$$\left(p + \frac{\theta}{1 - \theta}\right)U'(W - t_h)c = pU'(W - \ell + f(m) - t_s)f'(m), \tag{27}$$

$$\left(p + \frac{\theta}{1 - \theta}\right)U'(W - t_h) = pU'(W - \ell + f(m) - t_s). \tag{28}$$

Simplification shows that these conditions imply:

Proposition 5. *If a single noncollusion-proof policy is offered, the optimal policy satisfies:*

1. *Unless $\theta = 0$ the consumer is imperfectly insured: $W - \ell + f(m) - t_s < W - t_h$.*
2. *The recovery input is at the first-best level: $m = m^*$.*
3. *The truthful provider obtains zero profit: $\alpha_h = \alpha_s - mc = 0$.*
4. *In state s , the collusive provider reports $k = s$ and obtains zero profit: $\alpha_s - mc = 0$; in state h , he reports $k = s$ and obtains profit $t_h - t_s \geq 0$ through a side-payment from the consumer.*

In the previous section, we found that with a noncollusion-proof menu of policies, there was under-treatment if the provider was collusive and over-treatment if he was honest. Here, it turns out that the level of treatment is always first best. It can be explained by two differences: here, the consumer's utility does not depend on whether the provider is honest or dishonest; furthermore, there is no leeway in adjusting the amount of treatment according to the type of provider. However, in order to limit the payment to the collusive provider, the difference between the payments t_h and t_s is reduced compared to the first-best policy, resulting in imperfect insurance.¹⁵

To determine the equilibrium contract, we compare the consumer's expected utilities from the optimal collusion-proof and noncollusion-proof contracts. With the former, the expected utility is independent of θ . Furthermore:

Corollary 3. *If the policy is not collusion-proof, the expected utility of the consumer is strictly decreasing in θ .*

Proof. Since $m = m^*$ for all θ , the expected transfer from the consumer to the insurer increases with θ . Furthermore, condition (28) implies that the difference between $W - t_h$ and $W - \ell + f(m) - t_s$ increases as θ increases. \square

Moreover, when θ is close to one, the standard second-best policy dominates the optimal noncollusion-proof policy. Thus, we have proved:

Proposition 6. *If a single policy is offered, there exists $\hat{\theta}$ such that:*

- (i) *if $\theta \leq \hat{\theta}$, the equilibrium policy is the optimal noncollusion-proof policy in Proposition 5;*
- (ii) *if $\theta > \hat{\theta}$, the equilibrium policy is the standard second-best collusion-proof policy in Proposition 1.*

5. Conclusion

In this paper, we have characterized the equilibrium insurance-payment contracts when the consumer-provider coalition may sometimes file false claims with the insurer. Our focus has been on the risk sharing and treatment efficiency properties. We have explored two possibilities: either the insurer offers a menu of policies to screen the prospective providers, or it offers a single policy. Restricting attention to single policies leads to a loss of generality, and lower welfare. In both cases, however, allowing collusion turns out to be optimal when the coalition will falsify claims with small probabilities. When menus are used by the insurer, the equilibrium treatment level is higher than the first best when it is supplied by the honest provider; lower, when supplied by the dishonest provider. Insurance is never perfect, and the consumer must face some risks. The use of monetary transfers to

¹⁵ An interesting implication of this result is that production efficiency obtains when insurance is provided with positive loading: indeed, the problem can be interpreted as a standard insurance problem with positive loading. To see this, choose $\lambda \equiv (\theta(1 - p)/p)$ as the loading factor.

compensate the consumer's loss not recovered by treatment only occurs when collusion is tolerated. When an insurance and payment contract is designed, it is important that the interaction between consumers and providers be understood properly. As we have shown, the nature of this interaction—collusive or honest—determines the optimal design critically.

When the insurer is unable to screen the providers by offering a menu of policies, a simple monotonicity result obtains: collusion is tolerated if and only if the provider is dishonest with a sufficiently low probability (see the very last proposition). We have been unsuccessful in obtaining a similar monotonicity result when the insurer is able to use menus of policies. There remains the possibility that there are multiple equilibrium switches between collusion-proof and noncollusion-proof policies as the probability of the provider being dishonest varies, although we have been unable to construct such an example.

Our paper is one of the few in the insurance and health economics literature that explicitly models agents who are not completely self-interested. It seems likely that this framework will shed light on many institutions and phenomena, and future research in this direction may prove fruitful.

Acknowledgements

The research was conducted while the first author was visiting the Department of Economics at Boston University and supported by the Sweden–America Foundation. We thank Michael Manove for discussing various issues with us. For their comments, we thank Thomas McGuire and seminar participants at Boston College, Boston, Queen's, Rice, Vanderbilt and Washington Universities, CIDE, University College London, and the Universities of Dundee, Munich, Namur, and Southampton.

Appendix A

A.1. 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. An identical argument establishes that (8) binds. We have proved 1 in Proposition 1. With $\alpha_h^\sigma = 0 = \alpha_s^\sigma - m^\sigma c$, and with (9) binding, the problem becomes: choose t_h^σ , t_s^σ , and m^σ to maximize (6) subject to:

$$(1 - p)t_h^\sigma + p(t_s^\sigma - m^\sigma c) = 0, \quad (\text{A.1})$$

$$-t_h^\sigma \geq -t_s^\sigma, \quad (\text{A.2})$$

$$f(m^\sigma) - t_s^\sigma \geq -t_h^\sigma. \quad (\text{A.3})$$

Constraint (A.2) binds. If it did not, t_h^σ could be increased and t_s^σ decreased (in proportions such that (A.1) is unaffected). Such a change would not violate (A.3), but would decrease the difference between $W - t_h^\sigma$ and $W - \ell + f(m^\sigma) - t_s^\sigma$, since $f(m^\sigma) < \ell$. This proves

2 of the proposition. Now when $t_h^\sigma = t_s^\sigma$, 3 of the proposition follows from the assumption of $f(m) < \ell$.

Further, the binding constraint (A.2) implies that (A.3) is satisfied as a strict inequality. Substituting the binding constraint (A.2) into (A.1), we have $t_h^\sigma = pm^\sigma c$. Putting this into the objective function, we have reduced the problem to 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 >1 . Therefore, $f'(m^\sigma) < c$. This proves 4 of the proposition.

A.2. Proof of Proposition 2

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

$$\alpha_h^\tau - t_h^\tau \geq \alpha_s^\tau - m^\tau c - t_s^\tau, \tag{A.4}$$

$$\alpha_s^\tau - m^\tau c + f(m^\tau) - t_s^\tau \geq \alpha_h^\tau - t_h^\tau, \tag{A.5}$$

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), (A.4), (A.5), (10) and (11) and the budget constraint (14).

For the relaxed program, constraints (10) and (11) bind; this follows an argument similar to the one 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 rewrite the constraints:

$$-t_h^\sigma \geq -t_s^\sigma, \tag{A.6}$$

$$f(m^\sigma) - t_s^\sigma \geq -t_h^\sigma, \tag{A.7}$$

$$-t_h^\tau \geq -t_s^\tau, \tag{A.8}$$

$$f(m^\tau) - t_s^\tau \geq -t_h^\tau, \tag{A.9}$$

$$\theta[(1 - p)t_h^\sigma + p(t_s^\sigma - m^\sigma c)] + (1 - \theta)[(1 - p)t_h^\tau + p(t_s^\tau - m^\tau c)] \geq 0. \tag{A.10}$$

Following a now familiar argument, constraints (A.6) and (A.8) are binding; as a result, constraints (A.7) and (A.9) are satisfied with slack. Now use (A.6) and (A.8) as equalities to substitute into (A.10) and the objective function. So we simplify the program into choosing $t_h^\sigma, m^\sigma, t_h^\tau,$ and m^τ to maximize

$$\begin{aligned} &\theta[(1 - p)U(W - t_h^\sigma) + pU(W - \ell + f(m^\sigma) - t_h^\sigma)] \\ &+ (1 - \theta)[(1 - p)U(W - t_h^\tau) + pU(W - \ell + f(m^\tau) - t_h^\tau)], \end{aligned}$$

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.

A.3. Proof of Proposition 3

We begin by assuming that $\alpha_h^\tau - t_h^\tau < \alpha_s^\tau - m^\tau c - t_s^\tau$. We will later consider the other possibility. Then, constraint (22) reduces to:

$$\alpha_s^\sigma - m^\sigma c + (1 - p)(t_h^\sigma - t_s^\sigma) \geq \alpha_s^\tau - m^\tau c + (1 - p)(t_h^\tau - t_s^\tau). \tag{A.11}$$

First note that we can set $\alpha_h^\sigma = 0$. We next 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 (A.11) 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. Next, we show that for $j = \sigma$ (19) is binding. Suppose that it is 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 $(p/(1 - p))\epsilon$. By (20), t_h^τ can then also be decreased by $(p/(1 - p))\epsilon$. This in turn implies that the right-hand side of constraint (A.11) 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^σ has been increased, and t_h^τ has been decreased. Given that the constraints (19) for $j = \tau$ and for σ are binding, constraint (21) is redundant. Last, we show that constraint (18) must bind. Indeed, if it did not, α_h^τ could be decreased. By (20), t_h^τ could then be decreased, thus, raising the expected utility. We have, therefore, proved 1 and 2.

Using (18) and (19) as equalities to substitute for the α variables, we simplify the budget constraint (20) and the incentive constraint for the collusive provider (A.11) to:

$$\theta(t_s^\sigma - m^\sigma c) + (1 - \theta)[(1 - p)t_h^\tau + p(t_s^\tau - m^\tau c)] \geq 0, \tag{A.12}$$

$$t_h^\sigma - t_s^\sigma \geq t_h^\tau - t_s^\tau. \tag{A.13}$$

Constraint (A.13) must bind, since otherwise t_h^σ could be decreased; this does not affect the budget constraint and increases the expected utility. Similarly, the budget constraint is binding. Using (A.13) to substitute for t_h^σ and letting λ be the multiplier for constraint (A.12), the solution is characterized by the first-order conditions with respect to m^τ , t_s^τ , t_h^τ , m^σ and t_s^σ :

$$p(1 - \theta)U'(W - \ell + f(m^\tau) - t_s^\tau)f'(m^\tau) = \lambda(1 - \theta)pc, \tag{A.14}$$

$$p(1 - \theta)U'(W - \ell + f(m^\tau) - t_s^\tau) - \theta(1 - p)U'(W - t_h^\sigma) = \lambda(1 - \theta)p, \tag{A.15}$$

$$(1 - \theta)(1 - p)U'(W - t_h^\tau) + \theta(1 - p)U'(W - t_h^\sigma) = \lambda(1 - \theta)(1 - p), \tag{A.16}$$

$$\theta pU'(W - \ell + f(m^\sigma) - t_s^\sigma)f'(m^\sigma) = \lambda\theta c, \tag{A.17}$$

$$\theta(1 - p)U'(W - t_h^\sigma) + \theta pU'(W - \ell + f(m^\sigma) - t_s^\sigma) = \lambda\theta. \tag{A.18}$$

Conditions (A.17) and (A.18) together imply:

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

which implies that $f'(m^\sigma) > c$, or $m^\sigma < m^*$. Next we show that $m^\tau > m^*$. Adding (A.15) and (A.16) yields:

$$pU'(W - \ell + f(m^\tau) - t_s^\tau) + (1 - p)U'(W - t_h^\tau) = \lambda. \tag{A.19}$$

This and (A.14) imply:

$$\frac{U'(W - \ell + f(m^\tau) - t_s^\tau)}{pU'(W - \ell + f(m^\tau) - t_s^\tau) + (1 - p)U'(W - t_h^\tau)} f'(m^\tau) = c. \tag{A.20}$$

Getting an expression for λ from (A.16) and substituting it into (A.19), we obtain:

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

So we conclude that $U'(W - \ell + f(m^\tau) - t_s^\tau) > U'(W - t_h^\tau)$, implying that the consumer is imperfectly insured against the risk of losing ℓ when he is matched with a truthful provider. Furthermore, this together with (A.19) implies that $m^\tau > m^*$. Thus, we have proved 4.

We now compare the differences of the consumer's net wealth between states h and s, when he is matched with a collusive provider, and when he 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 proved 3 of the proposition.

Finally, we verify that we need not consider τ policies that do not verify $\alpha_h^\tau - t_h^\tau < \alpha_s^\tau - m^\tau c - t_s^\tau$ (see the assumption at the beginning of the proof). Consider the opposite case: $\alpha_h^\tau - t_h^\tau \geq \alpha_s^\tau - m^\tau c - t_s^\tau$. In words, the joint surplus in state h is higher when the truth is reported. If also $\alpha_s^\tau - m^\tau c + f(m^\tau) - t_s^\tau \geq \alpha_h^\tau - t_h^\tau$, then the τ policy is collusion-proof, and trivially, the optimal contract would then be the one in Proposition 2 (the optimal collusion-proof contract).

Hence, we need only consider contracts such that the $\alpha_h^\tau - t_h^\tau \geq \alpha_s^\tau - m^\tau c - t_s^\tau$ and $\alpha_s^\tau - m^\tau c + f(m^\tau) - t_s^\tau < \alpha_h^\tau - t_h^\tau$: the collusive provider would always report h if he were to pick the τ policy. The incentive constraint for the collusive provider (22) would be:

$$\alpha_s^\sigma - m^\sigma c + (1 - p)(t_h^\sigma - t_s^\sigma) \geq \alpha_h^\tau \tag{A.21}$$

It is straightforward to show that all the individual rationality constraints (18) and (19) bind. Hence, (A.21) becomes $t_h^\sigma - t_s^\sigma \geq 0$. Since this must be binding at any optimum (otherwise t_h^σ could be decreased without affecting any constraint), the collusive provider's policy is in fact collusion-proof, contradicting the assumption that it is not collusion-proof, i.e. that $t_h^\sigma - t_s^\sigma > 0$.

A.4. Proof of Corollary 1

Setting $\theta = 0$ in (A.14)–(A.16), we have

$$U'(W - \ell + f(m^\tau) - t_s^\tau) = U'(W - t_h^\tau),$$

and

$$f'(m^\tau) = c.$$

Furthermore, the budget constraint becomes $(1 - p)t_h^\tau + p(t_s^\tau - m^\tau c) = 0$, so that the policy is exactly first best. Obviously (16) is satisfied. The corollary follows from the upper-semicontinuity of the solutions to the constrained maximization program.

A.5. Proof of Corollary 2

Let $t_h^\tau - t_s^\tau = t_h^\sigma - t_s^\sigma \equiv k$. The program for Proposition 3 can be rewritten as the maximization of

$$\begin{aligned} &\theta[(1 - p)U(W - k - t_s^\sigma) + pU(W - \ell + f(m^\sigma) - t_s^\sigma)] \\ &+ (1 - \theta)[(1 - p)U(W - k - t_s^\tau) + pU(W - \ell + f(m^\tau) - t_s^\tau)], \end{aligned}$$

subject to $k \geq 0$ and

$$\theta(t_s^\sigma - m^\sigma c) + (1 - \theta)(t_s^\tau + (1 - p)k - pm^\tau c) = 0.$$

Letting μ and λ be the multipliers of the two constraints, we obtain the first-order conditions with respect to t_s^σ , t_s^τ , k , m^σ and m^τ , respectively:

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

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

$$\theta(1-p)U'(W-k-t_s^\sigma) + (1-\theta)(1-p)U'(W-k-t_s^\tau) = \lambda(1-\theta)(1-p) + \mu,$$

$$pU'(W-\ell + f(m^\sigma) - t_s^\sigma) f'(m^\sigma) = \lambda c,$$

$$U'(W-\ell + f(m^\tau) - t_s^\tau) f'(m^\tau) = \lambda c.$$

For the contract to be noncollusion-proof, the constraint $k \geq 0$ must not bind, i.e. we must have $\mu = 0$. Replacing λ of the third constraint by the left-hand side of the second constraint, and rearranging we get:

$$\begin{aligned} & \theta(1-p)U'(W-k-t_s^\sigma) \\ & = (1-\theta)(1-p)p[U'(W-\ell + f(m^\tau) - t_s^\tau) - U'(W-k-t_s^\tau)] + \mu. \end{aligned}$$

As θ tends to 1, the first term on the right-hand side tends to 0. The left-hand side being positive, we must have $\mu > 0$.

References

- Alger, I., Renault, R., 1998. Honest Agents and Equilibrium Lies. Boston College and Université de Cergy-Pontoise, Mimeo.
- Alger, I., Renault, R., 2000. Screening Among Agents with Heterogeneous Ethics. Boston College and Université de Cergy-Pontoise, Mimeo.
- Bowles, S., 1998. Endogenous preferences: the cultural consequences of markets and other economic institutions. *Journal of Economic Literature* 36, 75–111.
- Dawes, R.M., Thaler, R.H., 1988. Anomalies: cooperation. *Journal of Economic Perspectives* 2, 187–197.
- Erard, B., Feinstein, J.S., 1994. Honesty and evasion in the tax compliance game. *Rand Journal of Economics* 25, 1–19.
- Kofman, F., Lawarrée, J., 1996. On the optimality of allowing collusion. *Journal of Public Economics* 61, 383–407.
- Lu, M., 1997. Separating the ‘True Effect’ from ‘Gaming’ in Incentive-Based Contracts in Health Care, University of Calgary, Mimeo.
- Ma, C.-T.A., McGuire, T., 1997. Optimal health insurance and provider payment. *American Economic Review* 87, 685–704.
- 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: Laffont, J.-J. (Ed.), *Proceedings of the Sixth World Congress of the Econometric Society on Advances in Economic Theory*. Cambridge University Press, Cambridge, pp. 151–206.
- Tirole, J., 1996. A theory of collective reputation (with applications to the persistence of corruption and to firm quality). *Review of Economic Studies* 63, 1–22.