JOURNAL OF ECONOMIC THEORY **30**, 74–97 (1983)

Repeated Insurance Contracts and Moral Hazard*

ARIEL RUBINSTEIN AND MENAHEM E. YAARI

The Hebrew University, Jerusalem, Israel

Received May 20, 1981; revised September 28, 1981

An attempt is made to account for the frequently observed phenomenon of insurance companies offering discounts to clients who possess a favorable record of past claims. We argue that such discounts provide a mechanism which enables both insurer and insured to counteract the inefficiency which arises from moral hazard. *Journal of Economic Literature* Classification Number: 026.

1. INTRODUCTION

In this essay, we shall make an attempt to account for the frequently observed phenomenon of insurance companies offering discounts to clients who possess a favorable record of past claims. We shall argue that such discounts provide a mechanism which enables both insurer and insured to counteract the inefficiency which arises from moral hazard.

Moral hazard is an example of economic interaction involving imperfect observability: The insurer cannot observe certain actions taken by the insured, actions which, however, have an effect upon the insurer's payoff. Moreover, this inability of the insurer to observe the actions of the insured creates an incentive for the insured, once insurance is purchased, to act in a manner that is liable to enhance the likelihood of a large claim being filed. As a result, the scope for a mutually advantageous interaction becomes severely hampered and, in particular, fully efficient interaction becomes impossible because efficient contracts are not enforceable.

If interaction takes the form of an isolated contract, then any attempt to correct the inefficiency caused by moral hazard must take place through the specification of what rewards or penalties the insured person would incur for any given level of the commonly observable variables. Much of the literature on moral hazard to date has concentrated on the question of how to design a

* An earlier version of this paper was presented in April of 1979 at a conference on Economics of Uncertainty in honor of Karl Borch. We are indebted to Robert J. Aumann, Abraham Neyman, Roy Radner, and to an anonymous referee for their comments. Financial support from the U.S.–Israel Binational Science Foundation is gratefully acknowledged.

0022-0531/83 $3.00

scheme of such rewards and penalties so that the inefficiency brought about by moral hazard might be alleviated. (See, for example, Holmstrom [6], Pauly [9], and Shavell [12].) Indeed, Mirrlees has shown [8] that, under certain conditions, a scheme of rewards and penalties can be designed under which there exist enforceable contracts such that the loss of efficiency from moral hazard is arbitrarily small. To obtain this result, Mirrlees must allow for penalties which cause the insured person to suffer a negative utility of arbitrarily large magnitude. In general, sophisticated schemes of rewards and penalties will serve to reduce the inefficiency brought about by moral hazard, but they cannot serve to eliminate the inefficiency altogether.

Here, we shall de-emphasize the role of rewards and penalties and concentrate instead on temporal structure. Indeed, in what follows, full-indemnity insurance contracts will be assumed to be offered at a prespecified price, which makes the schedule of rewards and penalties incurred by the insured completely inflexible. The values taken by the commonly observed variables in any given period do not, in any way, affect the insured person's rewards in that period. On the other hand, rewards do change over time, in that the price of a given insurance contract may be different in different periods. In particular, the insurer will be allowed to adjust the price of the contract being offered in accordance with the insured's past record. It will be our purpose, in the sequel, to show that with insurance rates having this kind of temporal flexibility, it becomes possible for insurer and insured to reach an enforceable long-term contract that eliminates the inefficiency caused by moral hazard. This result holds regardless of the shape of the insured's utility function, as long as this function displays risk aversion. In particular, this function need not have the unboundedness property required for Mirrlees's result mentioned above.

The present paper is cast entirely in terms of a simple insurance problem. However, the discussion is clearly applicable, with the appropriate changes of interpretation, to a whole host of principal–agent problems in which the principal is unable to observe some of the actions taken by the agent. In two recent papers, Radner [10] and Rubinstein [11] have studied the intertemporal structure of certain principal–agent relationships along lines similar to those which we intend to pursue in the present essay. All three studies view the principal as being faced with a problem of statistical inference involving a delicate balance between two opposing types of errors. There are, however, some basic differences separating the three studies. In [10], the principal's optimal strategy lacks a certain "perfectness" property which the optimal strategy of the insurer in the present study will have. In [11], the agent's actions are observable and only his intentions are not, which is not the case in our insurance problem. These contrasts and others will be discussed in Section 5.

2. The Isolated Contract

Consider an asset which is susceptible to damage, with the extent of the damage depending jointly on the *care* exercised by the owner and on random factors. Let the value of the undamaged asset be fixed at unity and assume that damage takes the form of a detraction from this fixed value. The extent of the damage, once it has occurred, is universally observable and unambiguously measurable.

The owner of the asset considers the possibility of insuring it against damage. He can purchase insurance from a price-setting agency, who will be called "the insurer" for short. The insurer offers the owner a single contract on a "take it or leave it" basis: In exchange for a premium π, the insurer undertakes to indemnify the owner to the full extent of the damage. (Thus, contracts offering partial coverage are being ruled out by assumption.) The insurer's action is therefore described completely in terms of selecting a nonnegative real number π, the premium for the contract being offered. Letting the premium be measured in the same units as the value of the asset, we observe that only values of π lying in the unit interval make sense, so the action of the insurer is restricted at the outset to satisfy $\pi \in [0, 1]$.

Given that the insurer offers a contract with premium π, the owner of the asset has two decisions to make in response: He must decide whether or not to accept the offer, and he must determine the level of *damage-preventing care* that he wishes to exercise. Now, care is costly and we shall think of a given level of care in terms of the costs incurred in exercising it, measured in units of value. The action to be taken by the owner of the asset will therefore be described by a pair (b, c), with b taking the values 0 or 1 and with c taking nonnegative real values. Specifically, $(0, c)$ will stand for the decision to *decline* the insurer's offer and spend c on care and $(1, c)$ will stand for the decision to *accept* the insurer's offer and spend c on care. It is clear that values of c satisfying $c > 1$ do not make sense, so the a priori restrictions on the owner's action (b, c) are given by $b \in \{0, 1\}$ and $c \in [0, 1]$.

At any level of care, the actual extent of damage to the asset is a random variable. To formalize this, we shall say that after the insurer picks π and after the owner responds with a choice of (b, c), a *state of nature* ω is selected from some sample space Ω, with a commonly known probability measure, to be denoted μ. The values taken on by c and ω jointly determine the extent of the damage that will have been caused to the asset. Specifically, let $D(c, \omega)$ be the damage to the asset if its owner spends c on care and if the state of nature is ω, so that D is a function defined on $[0, 1] \times \Omega$ with values in $[0, 1]$. Referring to D as the *damage function*, we shall assume that $D(c, \cdot)$ is μ-measurable for all $c \in [0, 1]$ and that $D(\cdot, \omega)$ is continuous for all $\omega \in \Omega$. Also, for any given $c \in [0, 1]$, we shall let $D(c)$ be the random variable describing the damage to the asset at a level of care c, i.e.,

$D(c)(\omega) = D(c, \omega)$. Finally, expected damage at a given level of care will be denoted $\bar{D}(\cdot)$, so that

$$\bar{D}(c) = ED(c) = \int_{\Omega} D(c, \omega)\, d\mu(\omega)$$

for all $c \in [0, 1]$.

The damage function D will be assumed to satisfy the following additional assumptions:

(i) The expected damage function $\bar{D}$ is decreasing and convex.

(ii) The inequality

$$\bar{D}(0) > \bar{D}(c) + c$$

holds for some $c > 0$.

The first of these assumptions states that, on the average, care does indeed reduce damage, but at a nonincreasing rate. The second assumption states that care is economically meaningful in the sense that there exists some $c > 0$ such that, by spending an amount c on care, it is possible to reduce expected damage by more than c.

Let us define the *socially optimal*[1] *levels of care* as those values of c for which $c + \bar{D}(c)$ is at its minimum. Assumption (ii) above says that the care level $c = 0$ is not socially optimal.

Now suppose that all three choices have been made: The insurer has picked π, the owner of the asset has picked (b, c), and nature has picked ω. This will determine the monetary proceeds (in units of value) of the two parties as follows:

$$\begin{aligned} \text{insurer's proceeds} &= 0, && \text{if } b = 0,\\ &= \pi - D(c, \omega), && \text{if } b = 1;\\ \text{owner's proceeds} &= -c - D(c, \omega), && \text{if } b = 0,\\ &= -\pi - c, && \text{if } b = 1. \end{aligned}$$

Let us now assume that the insurer's payoff is given simply by the expected value of monetary proceeds. The owner of the asset, on the other hand, will be assumed to calculate expected utility and to be risk averse. (Without risk aversion, there will be no motivation to insure.) When the owner's monetary

[1] It is only appropriate to interpret these levels as being "socially optimal" if it is thought somehow that society as a whole is risk neutral.

proceeds are given by a random variable z, his payoff, to be written $U(z)$, will be given by

$$U(z) = Eu(z) = \int_{\Omega} u(z(\omega))\, d\mu(\omega),$$

where u, a real function defined for all real values, is assumed to be increasing and concave. Under these assumptions, the final payoffs of the two agents, given the choices of π and (b, c), are as follows:

$$\begin{aligned} \text{insurer's payoff} &= 0 && \text{if} \quad b = 0, \\ &= \pi - \bar{D}(c), && \text{if} \quad b = 1; \\ \text{owner's payoff} &= U(-c - D(c)), && \text{if} \quad b = 0, \\ &= u(-\pi - c), && \text{if} \quad b = 1. \end{aligned}$$

If, for some reason, insurance is *not* purchased (so that $b = 0$), then the owner of the asset must set the cost of care at a level, to be denoted $\bar{c}$, for which

$$\max_{0 \leqslant c \leqslant 1} U(-c - D(c))$$

is attained. For convenience, let us define $\bar{U}$ by writing

$$\bar{U} = U(-\bar{c} - D(\bar{c})),$$

so that $\bar{U}$ is the best payoff which the owner of the asset can secure under autarky, without interacting with anyone.

Now, the insurer cannot observe the level of care c, nor the state of nature ω. He cannot plan to separate out the effects of c and ω when he agrees to cover any damage that might be caused to the asset. In our model, there is no way for the insurer to support a claim of negligence against the owner of the asset, nor to vary the premium in accordance with observed damage. Knowing this, the owner of the asset will set the level of care at $c = 0$ whenever he decides to buy insurance (that is, if he sets $b = 1$, then he will set $c = 0$). This is because care is costly and insurance provides complete coverage. Therefore, if the insurer picks a premium π, then, *under moral hazard*, the payoffs to the two agents come out as follows:

$$\begin{aligned} \text{insurer's payoff} &= 0, && \text{if} \quad b = 0, \\ &= \pi - \bar{D}(0), && \text{if} \quad b = 1; \\ \text{owner's payoff} &= \bar{U}, && \text{if} \quad b = 0, \\ &= u(-\pi), && \text{if} \quad b = 1. \end{aligned}$$

At a given premium π, the owner will choose to insure the asset (i.e., he will set $b = 1$) if, and only if, the inequality $u(-\pi) \geqslant \bar{U}$ is satisfied. This inequality $u(-\pi) \geqslant \bar{U}$ is therefore a necessary and sufficient condition for the existence of an active insurance market under moral hazard. In order to be able to compare this situation with what would happen if moral hazard were absent, we find it useful to define two auxiliary functions, to be denoted U^* and U_*, in the following manner: Let x be any real number. Then,

$$U^*(x) = \max_{0 \leqslant c, \pi \leqslant 1} u(-c - \pi) \qquad \text{subject to} \quad \pi - \bar{D}(c) \geqslant x$$

and

$$U_*(x) = u(-\bar{D}(0) - x).$$

The function U^* describes the utility-possibility frontier that would be accessible through the market if moral hazard were absent, while U_* describes the actual utility-possibility frontier, under moral hazard. Suppose that the insurer's payoff (i.e., excess of premium earned over expected payments on claims) were fixed at a given level, x. In the absence of moral hazard, the insurer is able to observe the level of care c and to charge different premiums for different values of c. And since the insurer's payoff is fixed at x, the premium that would have to be charged when the owner sets care at the level c is given by $\bar{D}(c) + x$. In response to this schedule of insurance premiums, the owner of the asset would select c in $[0, 1]$ precisely as described in the maximization problem defining $U^*(x)$. Note that the level of care selected by the owner in this situation is, in fact, a socially optimal level (in the sense discussed above). Now, in the presence of moral hazard, the insurer knows that the owner will set care at $c = 0$, so the premium that he must charge to earn a payoff of size x is given by $\bar{D}(0) + x$. The owner, if he choses to insure, winds up with a payoff of $u(-\bar{D}(0) - x) = U_*(x)$.

It follows from our assumptions on the utility function u that the functions U^* and U_* are both decreasing and concave. Also, given the properties of the damage function D, we note that the inequality $U^*(x) > U_*(x)$ is true for all values of x. This inequality gives the sense in which moral hazard causes inefficiency: The utility-possibility frontier in the case where moral hazard is absent lies uniformly above the corresponding frontier for the case where moral hazard is present.

By refusing to transact (or, equivalently, by charging a premium $\pi = 1$), the insurer obtains an assured payoff of size 0. Similarly, the owner of the asset can obtain an assured payoff of size $\bar{U}$ without recourse to insurance of any kind. Thus, the pair $(0, \bar{U})$ describes the payoffs to the two agents under autarky. It should be noted, furthermore, that the inequality $\bar{U} \leqslant U^*(0)$ must

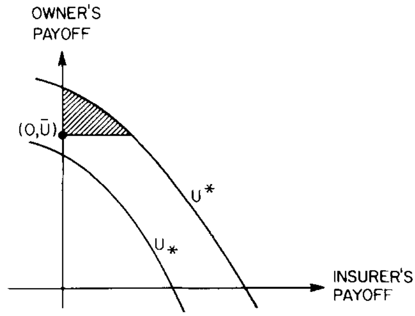

FIG. 1. Here $\bar{U} > U_*(0)$.

always hold. (This follows immediately from Jensen's inequality.) The meaning of this inequality is simply that, in the absence of moral hazard, risk averse individuals find actuarially fair insurance attractive. Now, regarding the relationship between $\bar{U}$ and $U_*(0)$, we note that the inequalities $\bar{U} \geqslant U_*(0)$ and $\bar{U} < U_*(0)$ are both possible. If the former inequality holds—$\bar{U} \geqslant U_*(0)$—then moral hazard has the effect of foreclosing insurance transactions altogether. The two agents end up at the point of autarky, with payoffs $(0, \bar{U})$. If the latter inequality holds—$\bar{U} < U_*(0)$—then some mutually advantageous transactions exist even under moral hazard, but such transactions are very much restricted in comparison with the transactions that would take place if moral hazard were eliminated. These two cases are depicted in Figs. 1 and 2, respectively. In both figures, the shaded areas represent individually rational outcomes, i.e., pairs of payoffs giving each agent at least what the agent could get on his own. Because of moral hazard, many of these outcomes are not attainable, in the sense that contracts leading to these outcomes are not enforceable. Under moral hazard, only the outcomes lying in the heavily shaded area (and, in Fig. 1, only the outcome $(0, \bar{U})$ itself) are both attainable and individually rational.

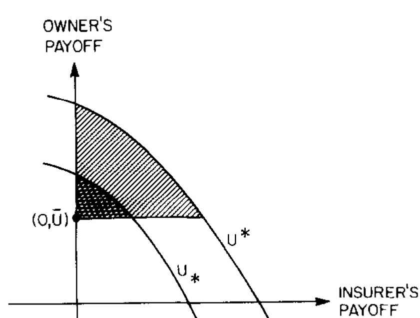

FIG. 2. Here $\bar{U} < U_*(0)$.

3. Repeated Contracts

Now let us imagine a situation where the interaction described in the previous section recurs time and again. Both insurer and insured know, before agreeing on anything, that they will be meeting each other periodically and that there will be an opportunity for concluding a separate contract in each period. Once again, we shall think of the insurer as offering a price for insurance and of the insured as responding to the insurer's offer. However, both the insurer's offer and the insured's response will now take the form of long-term strategies, specifying how each party will act at any given date, depending on the information that will be available to the party at that date. Interaction will be assumed to take place in the following order: The insurer starts out by announcing a strategy of insurance sales. The insured responds by selecting a strategy of insurance purchases and levels of care. After both parties have picked their strategies, "history" begins to evolve through the unfolding of a sequence of random events.

We have chosen *not* to model the interaction between insurer and insured in terms of a full-fledged supergame, because we felt that the kind of Stackelberg view proposed in the previous paragraph—with one player announcing a strategy and the opponent picking a best response—is more suitable for our problem. In particular, we do not allow the insurer ever to take actions which deviate from his announced strategy, so that it is not necessary for us to consider what the insured's response to such deviations might be and whether or not it is possible for the insured to force the insurer to abide by his own announcement. We shall return to this question briefly in Section 5.

As before, we think of the insured as owning an asset of unit value. The number $D(c, \omega)$ will, once again, stand for the extent of the damage to the asset in any period, given that the owner spends an amount c ($c \in [0, 1]$) on care and that the state of nature is ω ($\omega \in \Omega$). Damage takes the form of a detraction from the value of the asset, and the damage function D is assumed to be the same in all periods and to satisfy all the assumptions introduced in the previous section. We assume also that, at the beginning of each period, the insurable asset is restored to unit value, regardless of any damage which it may have suffered in the preceding period, and regardless of whether or not it had been insured in the preceding period. We do not inquire into the source of the funds that the owner of the asset must put forth in order to pay insurance premiums or otherwise to restore the asset to its original value. As in the previous section, the only insurance contracts to be considered are full-indemnity contracts where, in return for a premium which is independent of observed damage, the insurer undertakes to cover *all* losses. Finally, letting ω^t be the state of nature in period t, we assume that $\{\omega^1, \omega^2, \ldots\}$ is a sequence of independent observations from a sample space Ω, with the probability law given in each period by the same probability measure μ.

Now we must state precisely what *information* the two parties will have in any given period. Let us consider the insurer first. At the end of any given period, the insurer will have obtained answers to the following two questions: First, has the asset been insured in this period? Second, *if* the asset has been insured, what was the size of the claim filed by the owner, if any. (By our assumption, if insurance is purchased, then the owner always submits a claim for the full extent of the damage, the latter being commonly observable.) At any given moment of time, the insurer's information will therefore consist of a history of answers to these two questions. If we let the insurer's *one-period* information set be denoted I_1, then we can write

$$I_1 = \{(0, 0)\} \cup \{1\} \times [0, 1],$$

i.e., I_1 consists of the pair $(0, 0)$ and of all the pairs of the form $(1, d)$, with $d \in [0, 1]$. Here, $(0, 0)$ means "asset has not been insured" and $(1, d)$ means "asset has been insured and a claim of size d has been filed." If $\theta_1 \in I_1$, then we shall write

$$\theta_1 = (\theta_B, \theta_D)$$

with $\theta_B \in \{0, 1\}$ telling whether or not insurance has been bought and $\theta_D \in [0, 1]$ telling, in case insurance has been bought, how much damage has been claimed. For $t = 1, 2, \ldots$, let I_1^t be the t-fold Cartesian product of I_1 with itself, i.e., the set of all t-histories of pairs of the form (θ_B, θ_D). Under the assumption of perfect memory, I_1^t describes all the states of information in which the insurer can possibly be at the end of the tth period. We shall refer to I_1^t as the insurer's information set in period $t + 1$.

As for the insured, in the course of any given period he too obtains answers to two questions: First, what was the premium being charged by the insurer in this period? Second, what was the extent of the damage to the asset in this period? One would think, therefore, that the owner's information at any given moment is a history of answers to these two questions. However, as has already been asserted above, our framework here differs from that of a full-fledged supergame in that the insurer is assumed never to deviate from the initially announced strategy. Consequently, the owner of the asset is always in a position to *compute* the premium being charged in any given period from the insurer's pre-announced strategy, and there is no need to count this premium as part of the information gathered by the owner in the course of the period. Thus, the one-period information set of the owner of the asset is given simply by

$$I_2 = [0, 1],$$

with $d \in I_2$ standing for "damage this period has been of size d." As before, we let I_2^t be the t-fold Cartesian product of I_2 with itself, i.e., the set of all possible t-histories of damage to the asset. Then I_2^t is the owner's information set in period $t+1$.

What are the *actions* to be taken by the two parties? The insurer, for his part, knows that in period t he will have to pick a *premium* π that he will charge for insuring the asset on a full-indemnity basis. This premium will be some real number lying in the unit interval. Writing P for the set of actions open to the insurer in any given period, we therefore have

$$P = [0, 1].$$

As for the owner of the asset, his action in period t will be twofold: First, he will either insure the asset or refrain from doing so. Second, he will pick an amount c to be spent on damage-preventing care. Therefore, we can write the set of actions open to the owner of the asset in any given period as the product $B \times C$, where

$$B = \{0, 1\}, \qquad C = [0, 1]$$

so that $(0, c)$ will stand for the action "do not insure and spend c on care" and $(1, c)$ will stand for "insure and spend c on care."

It is now possible to define the notion of a strategy for either party. Intuitively, a strategy is a rule which maps information into action. Formally, we have:

A *strategy for the insurer* is a sequence of the form $f = (f^1, f^2, \ldots)$ such that $f^1 \in P$ and such that, for $t = 1, 2, \ldots$, f^{t+1} is a Borel-measurable[2] function defined on I_1^t and having values in P.

Similarly, a *strategy for the owner of the asset* is a sequence of the form $g = (g^1, g^2, \ldots)$, such that $g^1 \in B \times C$ and such that, for $t = 1, 2, \ldots$, g^{t+1} is a Borel-measurable function defined on I_2^t with values in $B \times C$. For each t, g^t can be decomposed in a natural manner by writing $g^t = (g_B^t, g_C^t)$.

Let S_1 and S_2 be the sets of all strategies for the insurer and the owner, respectively, and define $S = S_1 \times S_2$. Given a pair $(f, g) \in S$, let $\theta_1^t(f, g)$ and $\theta_2^t(f, g)$ be the *prospective* (or anticipated) state of information of the insurer and the owner, respectively, at the end of period t, on the assumption that the strategies f and g will be followed. Then $\theta_1^t(f, g)$ and $\theta_2^t(f, g)$ are random variables which depend on how the various random events occurring up to period t will have turned out and which take their values, respectively, in I_1^t and I_2^t. Formally, $\theta_1^t(f, g)$ and $\theta_2^t(f, g)$ must be defined inductively, as follows: Let $(f, g) \in S$, with $f = (f^1, f^2, \ldots)$ and $g = (g^1, g^2, \ldots)$. We shall

[2] Borel-measurability is required for f^{t+1} and g^{t+1} below in order to ensure that random variables appearing below are well defined.

write $\theta_1^t = ((\theta_B^1, \theta_D^1),..., (\theta_B^t, \theta_D^t))$. For $t = 1, 2...$, we may now define $\theta_1^t(f, g)$ and $\theta_2^t(f, g)$ in the following manner: First, we write

$$\theta_B^1(f, g) = g_B^1, \qquad \theta_D^1(f, g) = g_B^1 \times D(g_C^1, \omega^1),$$

where $\times$ denotes multiplication, and

$$\theta_2^1(f, g) = D(g_C^1, \omega^1).$$

Then, for $t = 2, 3,...$, we have

$$\theta_B^t(f, g) = g_B^t(\theta_2^{t-1}(f, g))$$
$$\theta_D^t(f, g) = g_B^t(\theta_2^{t-1}(f, g)) \times D(g_C^t(\theta_2^{t-1}(f, g), \omega^t)),$$

and

$$\theta_2^t(f, g) = (\theta_2^{t-1}(f, g), D(g_C^t(\theta_2^{t-1}(f, g), \omega^t))).$$

In all of the above, ω^t stands for the state of nature observed in period t.[3]

For a given pair of strategies $(f, g) \in S$, the prospective (or anticipated) *actions* of the insurer and the owner of the asset in period t are given, respectively, by $f^t(\theta_1^{t-1}(f, g))$ and $g^t(\theta_2^{t-1}(f, g))$—and by f^1 and g^1 in the case $t = 1$. These two quantities are random variables taking values, respectively, in the insurer's action space P and in the owner's action space $B \times C$. Having noted this, it is now possible to go ahead and write down the prospective (or anticipated) *single-period payoffs* of the parties. Let $h_1^t(f, g)$ be the insurer's prospective random payoff in period t and let $h_2^t(f, g)$ be the owner's prospective random payoff in period t, both calculated for the given choice of strategies $(f, g) \in S$. Then, we have, except for an obvious modification when $t = 1$, that

$$h_1^t = 0, \qquad \text{if } g_B^t(\theta_2^{t-1}) = 0,$$
$$= f^t(\theta_1^{t-1}) - D(g_C^t(\theta_2^{t-1}), \omega^t), \qquad \text{if } g_B^t(\theta_2^{t-1}) = 1;$$

and

$$h_2^t = u(-g_C^t(\theta_2^{t-1}) - D(g_C^t(\theta_2^{t-1}), \omega^t)), \qquad \text{if } g_B^t(\theta_2^{t-1}) = 0,$$
$$= u(-g_C^t(\theta_2^{t-1}) - f^t(\theta_1^{t-1})), \qquad \text{if } g_B^t(\theta_2^{t-1}) = 1,$$

where, for typographical reasons, we have suppressed the arguments (f, g) throughout the equations. Here $h_1^t(f, g)$ and $h_2^t(f, g)$ are random variables

[3] Note that $\theta_1^t(f, g)$ and $\theta_2^t(f, g)$ actually do not depend on f. However, we have chosen to retain the symbol f in our notation.

whose expected values are the respective single-period payoffs of the two parties, as defined in the previous section.

Both the insurer and the owner of the asset are assumed to be interested in the long-run behavior of the payoffs which they receive. In particular, we shall assume that both parties evaluate infinite streams of prospective random payoffs by attempting to calculate the long-run arithmetic average along these streams. Consider a pair of strategies $(f, g) \in S$. We shall say that the pair (f, g) is *averageable* if there exist two real numbers, to be denoted $H_1(f, g)$ and $H_2(f, g)$, such that, as $T \to \infty$,

$$\frac{1}{T} \sum_{t=1}^{T} h_1^t(f, g) \to H_1(f, g) \qquad \text{a.s.}$$

and

$$\frac{1}{T} \sum_{t=1}^{T} h_2^t(f, g) \to H_2(f, g) \qquad \text{a.s.}$$

The subset of S consisting of all averageable pairs of strategies will be denoted $\tilde{S}$. Note that the definition of averageability requires convergence almost surely, whereas the structure of the problem indicates that mean convergence should suffice. However, it will turn out that even the more restricted class of strategies $\tilde{S}$ is broad enough to give us our results.

At this point we must write down, in a precise way, what is to be meant by a pair of strategies $(\tilde{f}, \tilde{g}) \in \tilde{S}$ being *enforceable*. Recall the structure of interaction being assumed here, namely the insurer moving first, followed by the owner of the asset picking a response. Therefore, it is reasonable to regard a pair $(\tilde{f}, \tilde{g}) \in \tilde{S}$ as enforceable if $\tilde{g}$ is, in a suitably defined sense, a best response on the part of the owner to an offer of $\tilde{f}$ on the part of the insurer. This leads to the following definition:

Let $(\tilde{f}, \tilde{g}) \in \tilde{S}$. We shall say that $\tilde{g}$ is a *best response* to $\tilde{f}$ if there does not exist a $g \in S_2$ having the property that, for some $\varepsilon > 0$, the event

$$\frac{1}{T} \sum_{t=1}^{T} h_2^t(\tilde{f}, g) \geqslant H_2(\tilde{f}, \tilde{g}) + \varepsilon \qquad \text{for infinitely many values of } T$$

has positive probability. In other words, we assume that a sufficient cause for the owner of the asset to reject $\tilde{g}$ as a possible response to $\tilde{f}$ would be the finding of a strategy $g \in S_2$ for which

$$\limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} (h_2^t(\tilde{f}, g) - h_2^t(\tilde{f}, \tilde{g})) > 0$$

has positive probability.

Let $(\tilde{f}, \tilde{g}) \in \tilde{S}$. We shall say that the insurer *can enforce* $\tilde{g}$ *by announcing* $\tilde{f}$ (or, more concisely, that the pair $(\tilde{f}, \tilde{g})$ is *enforceable*) if $\tilde{g}$ is a best response to $\tilde{f}$.

It is our objective in the present essay to see whether, by adopting a suitable "no-claims discounts" strategy, the insurer can enforce upon the owner a strategy under which the latter would always exercise the proper amount of care, thus wiping out the ill-effects of moral hazard. So, we must now turn to a precise definition of what will be meant by a "no-claims discounts" strategy for the insurer. Basically, an insurer offering no-claims discounts is saying this to his potential clients: "Before each renewal of your policy, we shall look at your long-run record with us. If this record looks reasonable, then your premium for renewal will be low. However, if your record should be unreasonable, then your premium will be high." A client's record with an insurance company is clearly the history of claims having been filed by this client in the past. This record may be thought of as "reasonable" if, *on average*, claims filed in the past have not significantly exceeded the level of claims that would be expected if it were known that the proper amount of care was indeed being exercised by the client. Formally, we propose the following definition:

A strategy $f \in S_1$ for the insurer, $f = (f^1, f^2, \ldots)$, will be referred to as a *no-claims discounts strategy* (or as an NCD strategy, for short) if it has the following form: There exist three numbers, c_*, π_*, and π^*, satisfying $0 \leqslant c_* \leqslant 1$ and $0 \leqslant \pi_* \leqslant \pi^* \leqslant 1$, such that

$$f^1 = \pi_*$$

and, for $t = 1, 2, \ldots,$

$$\begin{aligned} f^{t+1}(\theta_1^t) &= \pi_*, \qquad \text{if} \quad \sum_{s=1}^{t} \theta_D^s / N(t) < \bar{D}(c_*) + \alpha^{N(t)} \\ &= \pi^*, \qquad \text{otherwise,} \end{aligned}$$

where $\{\alpha^\tau\}$ is a sequence of nonnegative real numbers and where $N(t)$ is defined by $N(t) = \Sigma_{s=1}^{t} \theta_B^s$ for $t = 1, 2, \ldots$.

Recall that θ_B^s is 1 or 0 according as the asset has or has not been insured in period s and that θ_D^s is the size of the claim filed in period s (with $\theta_B^s = 0$ implying $\theta_D^s = 0$). Accordingly, the ratio $\sum_{s=1}^{t} \theta_D^s / N(t)$ is the average size of the claims filed by the owner of the asset in the first t periods. (Here, 0/0 must be interpreted as 0.) Under an NCD strategy, the insurer will, in period $t + 1$, charge a low premium π_* if the average size of the claims filed up to that period falls below the expected damage to the asset, given the level of care c_* (where c_* is interpreted as the level of care deemed proper by the insurer) plus some margin of error $\alpha^{N(t)}$. The sequence $\{\alpha^\tau\}$ is introduced here in order to allow the insurer to "go easy" on his client, by agreeing that if the average size of the claims filed in the past is only slightly above $\bar{D}(c_*)$, then this would still qualify as "reasonable" and the premium to be charged on this kind of record would be at the low rate π_*. Thus, an insurer who adopts an NCD strategy must announce four things for his strategy to be completely specified, namely, c_*, π_*, π^*, and the "forgiveness" sequence $\{\alpha^\tau\}$.

Now, at long last, we are ready to state

THEOREM 1. *Let* $c_0 \in [0, 1]$ *and* $\pi_0 \in [0, 1]$ *be such that*

(1) $\pi_0 - \bar{D}(c_0) \geqslant 0$ [4] *and* $u(-c_0 - \pi_0) \geqslant \bar{U}$.

Then, there exists a pair $(\tilde{f}, \tilde{g}) \in \tilde{S}$ *having the following properties*:

(2) $\tilde{f}$ *is an NCD strategy*;

(3) $H_1(\tilde{f}, \tilde{g}) = \pi_0 - \bar{D}(c_0)$, $H_2(\tilde{f}, \tilde{g}) = u(-c_0 - \pi_0)$;

(4) $(\tilde{f}, \tilde{g})$ *is enforceable.*

In order to interpret what is being asserted here, note first that the inequalities in Condition (1) are equivalent to the statement that if, in an *isolated* contract, the insurer were to pick the action π_0 and the owner were pick the action $(1, c_0)$, then the outcome would yield, for each party, a payoff at least as high as what the party could get on its own. That is, Condition (1) characterizes the *individually rational* outcomes for a single contract. (See Figs. 1 and 2, where these outcomes form the shaded areas.) Thus, what Theorem 1 states is that, in the long run, the insurer can enforce any individually rational outcome, by adopting a suitable NCD strategy. More precisely, given any individually rational outcome, the insurer has an NCD strategy $\tilde{f}$ which enforces a response $\tilde{g}$, such that the pair $(\tilde{f}, \tilde{g})$ is averageable and yields average payoffs tending to the payoffs in the given outcome. As might be expected, the strategy $\tilde{g}$—a best response for the owner to the insurer's NCD strategy $\tilde{f}$—consists of picking the pair $(1, c_0)$ in all periods: The owner always insures his asset and always exercises the proper amount of care.

Theorem 1 tells us, in particular, that all the outcomes which are fully Pareto optimal (and individually rational) are enforceable through NCD strategies. In other words, there exists a family of NCD strategies for the insurer which elicit a *socially optimal* level of care (see Section 2) from the owner of the asset in every period. These are the strategies which correspond, via Theorem 1, to the fully Pareto optimal outcomes for a single contract.

Now, the *converse* of Theorem 1 would read as follows: Let $(\tilde{f}, \tilde{g}) \in \tilde{S}$ be

[4] The theorem remains true also when this inequality is dropped.

an averageable pair of strategies having the following three properties: (i) $\tilde{f}$ is an NCD strategy; (ii) The pair $(\tilde{f}, \tilde{g})$ is enforceable; (iii) $H_1(\tilde{f}, \tilde{g}) \geqslant 0$. Then, the pair $(H_1(\tilde{f}, \tilde{g}), H_2(\tilde{f}, \tilde{g}))$ is an individually rational outcome for a single contract. Is this converse assertion true? The answer is yes, and it follows readily from the following more general proposition:

THEOREM 2. *Let $(\tilde{f}, \tilde{g}) \in \tilde{S}$ be an averageable pair of strategies. Then, the inequality*

$$H_2(\tilde{f}, \tilde{g}) \leqslant U^*(H_1(\tilde{f}, \tilde{g}))$$

holds, with U^* as defined in Section 2, that is,

$$U^*(H_1(f, g)) = \max_{0 \leqslant c, \pi \leqslant 1} u(-c - \pi) \quad \textit{subject to } \pi - \bar{D}(c) \geqslant H_1(\tilde{f}, \tilde{g}).$$

Theorem 2 has some straightforward implications. Suppose, for example, that the insurer is a monopolist whose objective is to maximize the long-run average of his own net earnings. Then, it follows from Theorem 2 that the best the insurer can do is find the pair (c_0, π_0) which solves the problem

$$\max_{0 \leqslant c, \pi \leqslant 1} \pi - \bar{D}(c) \qquad \text{subject to} \quad u(-c - \pi) \geqslant \bar{U},$$

and then pick an NCD strategy which enforces the single period outcome $(\pi_0 - \bar{D}(c_0), u(-c_0 - \pi_0))$ in the long run. On the other hand, suppose (as in the theory of optimal taxation) that the insurer is a public agency whose objective is to maximize the long-run average utility of the owner of the asset, given only that he, the insurer, shall not make a loss. (This is the framework being considered by Mirrlees [8].) Then, it follows from Theorem 2 that the best he can do is find the pair (c_0, π_0) which solves the problem

$$\max_{0 \leqslant c, \pi \leqslant 1} u(-c - \pi) \qquad \text{subject to} \quad \pi - \bar{D}(c) \geqslant 0,$$

and then pick an NCD strategy which enforces the single period outcome $(0, u(-c_0 - \pi_0))$ in the long run. In both cases, it follows from Theorem 1 that the insurer can do what is being asserted and it follows from Theorem 2 that he can do no better.

4. PROOFS

In order to prove Theorems 1 and 2, we shall have to use several assertions from probability theory. We turn first to a statement of these assertions.

ASSERTION A (Law of the Iterated Logarithm). *Let $\{X^t\}$ be a sequence of independent identically distributed random variables with finite means μ and finite variances σ^2. Then, for every $\lambda > 1$, almost surely,*

$$\limsup_T \left[\left| \mu - \frac{1}{T} \sum_{t=1}^{T} X^t \right| \Big/ (2\lambda\sigma^2 \log\log T/T)^{1/2} \right] < 1.$$

Proof. See, e.g., Breiman [2, p. 291].

ASSERTION B. *Let K be a real number. Let $\{X^t\}$ be a sequence of random variables. Let x^t denote a value of X^t and, for every $x^1,..., x^T$, let $(1/T)\sum_{t=1}^{T} x^t$ be denoted $\bar{x}^T$. If $\bar{x}^T > K$, $E(x^{T+1} \mid x^1,..., x^T) \leqslant K$. Then, almost surely,* $\limsup \bar{x}^T \leqslant K$.

Proof. See Blackwell [1, Theorem 1].

ASSERTION C (Strong Law of Large Numbers). *Let $\{X^t\}$ be a uniformly bounded sequence of random variables. Then, almost surely,*

$$\frac{1}{T} \sum_{t=1}^{T} [X^t - E(X^t \mid X^1,..., X^{t-1})] \to 0.$$

Proof. See, e.g., Loeve [7].

ASSERTION D. *Let X be a random variable bounded by some real number B. Let A be an event with* $\Pr(A) = 1 - \varepsilon$. *Then*

$$|E(X) - E(X \mid A)| \leqslant 2B\varepsilon/(1 - \varepsilon).$$

Proof. This result is immediate.

ASSERTION E (Consequence of Egoroff's Theorem). *Let $\{\mathbf{X}^t\}$ be a sequence of random vectors with values in $\mathbb{R}^n$. Let $A \subset \mathbb{R}^n$ be a closed set with the property that $\rho(\mathbf{X}^t, A) \to 0$ a.s., where ρ denotes Euclidean distance. Then, for every $\delta > 0$, there exists an event E with* $\Pr(E) \geqslant 1 - \delta$ *such that $\rho(\mathbf{X}^t, A) \to 0$ uniformly on E.*

Proof. See, e.g., Halmos [4, p. 88].

ASSERTION F. *Let $A \subset \mathbb{R}^n$ be a convex compact set, and let $\{\mathbf{X}^t\}$ be a sequence of random vectors with values in $\mathbb{R}^n$ such that $E(\mathbf{X}^t \mid \mathbf{X}^1,..., \mathbf{X}^{t-1}) \in A$ for all t. Let $(1/T)\sum_{t=1}^{T} \mathbf{X}^t$ be denoted $\bar{\mathbf{X}}^T$. Then $\rho(\bar{\mathbf{X}}^t, A) \to 0$ a.s., where ρ is Euclidean distance.*

Proof. This follows from Assertion C, together with the convexity of A.

ASSERTION G. *Let $\{X^t\}$ be a sequence of bounded random variables such that, for all t and for all values $x^1,..., x^{t-1}$ taken on by $X^1,..., X^{t-1}$, $E(X^t \mid x^1,..., x^{t-1}) \leqslant \bar{U}$. Let $\{\delta^t\}$ be a sequence of random variables such that $\delta^t = 0$ or 1 and δ^t is measurable in the sigma-field generated by $X^1,..., X^{t-1}$. Then*

$$\text{Prob}\left\{\limsup \left(\sum_{t=1}^{T} \delta^t X^t\right) \Big/ \left(\sum_{t=1}^{T} \delta^t\right) \leqslant \bar{U} \text{ or } \sum_{t=1}^{\infty} \delta^t < \infty\right\} = 1.$$

Proof. See Freedman [3].

We turn now to the proofs of the two theorems stated in Section 3.

Proof of Theorem 1. For convenience, let $u(-c_0 - \pi_0)$ be denoted U_0. To specify the insurer's NCD strategy $\tilde{f}$ select π_*, π^*, c_*, and a sequence $\{\alpha^\tau\}_{\tau=1}^{\infty}$ as follows: $\pi_* = \pi_0$, $\pi^* = 1$, $c_* = c_0$, and

$$\alpha^\tau = \sqrt{(2\lambda\sigma^2 \log \log \tau)/\tau},$$

where $\lambda > 1$ and $\sigma^2 = \text{Var}\, D(c_0)$. Let $\tilde{g}$ be given by $\tilde{g}^t \equiv (1, c_0)$ for all t, i.e., in every period, the owner insures the asset and exercises care at the level c_0. We must show that the pair $(\tilde{f}, \tilde{g})$ satisfies Conditions (3) and (4) in the statement of Theorem 1. To see that Condition (3) is satisfied, note that $\{D(c_0, \omega^t)\}$ is a sequence of independent identically distributed random variables. By Assertion A, we have, a.s., that for all but finitely many T

$$\frac{1}{T} \sum_{t=1}^{T} D(c_0, \omega^t) - \bar{D}(c_0) < \alpha^T.$$

Since $\theta_B^t \equiv 1$, we have $\sum_{t=1}^{T} \theta_B^t \equiv T$ and so, almost surely, it is true for all but finitely many values of T that

$$\left(\sum_{t=1}^{T} \theta_D^t\right) \Big/ \left(\sum_{t=1}^{T} \theta_B^t\right) < \bar{D}(c_0) + \alpha^T.$$

Therefore, by the construction of $\tilde{f}$, we have, almost surely, that $\tilde{f}^T = \pi_*$, except for finitely many values of T. Now, using Assertion C, we get

$$\frac{1}{T} \sum_{t=1}^{T} h_1^t(\tilde{f}, \tilde{g}) \to \pi_* - \bar{D}(c_*) = \pi_0 - \bar{D}(c_0) \qquad \text{a.s.}$$

and

$$\frac{1}{T} \sum_{t=1}^{T} h_2^t(\tilde{f}, \tilde{g}) \to u(-c_* - \pi_*) = U_0 \qquad \text{a.s.}$$

Next, we must show that Condition (4) is satisfied, i.e., that $\tilde{g}$ is a best response to $\tilde{f}$. Assume, contrary to Condition (4), that there exists a $g \in S_2$ such that

$$\Pr(\bar{U}^T > U_0 + \varepsilon_0 \text{ for infinitely many values of } T) > \varepsilon_0,$$

where $\bar{U}^T = (1/T) \sum_{t=1}^{T} h_2^t(\tilde{f}, g)$. By Assertion G, we may assume, without loss of generality, that $g_B^t \equiv 1$ for all t. Let $\bar{d}^T$ be the average claim in the first T periods, when the owner plays g. That is,

$$\bar{d}^T = \frac{1}{T} \sum_{t=1}^{T} \theta_D^t(\tilde{f}, g).$$

Define two real-valued functions d_* and d^* as follows:

$$d_*(v) = \bar{D}(\pi_0 - u^{-1}(v)) \qquad \text{and} \qquad d^*(v) = \bar{D}(1 - u^{-1}(v)).$$

Here $d_*(v)$ and $d^*(v)$ are the levels of expected damage to the asset which give the owner a utility level v when the insurance premium is π_0 or 1 ($=\pi^*$), respectively. The domains of d_* and d^* are closed intervals and their common range is contained in $[0, 1]$. The assumptions on u and $\bar{D}$ imply that both d_* and d^* are increasing convex functions. Furthermore, the graph of d_* must lie to the right of the graph of d^*. (That is, $(d_*)^{-1}(y) > (d^*)^{-1}(y)$ for all y in the common range of d_* and d^*.) Moreover, the choice of $\pi^* = 1$ implies that $(d^*)^{-1}(y) \leqslant \bar{U}$ for all y in the range of d^*. All these properties are summarized in Fig. 3.

It follows from Assertion D that we can select $\hat{\varepsilon}$, $0 < \hat{\varepsilon} < \varepsilon_0$, to be so small that for every c and for every t, $E[-c - D(c, w^t) \mid \hat{E}] \leqslant \bar{U} + (\varepsilon_0/4)$. Let A be the convex hull of the graphs of d_* and d^*. It follows from Assertion F that $\rho[(\bar{U}^T, \bar{d}^T), A] \to 0$, a.s., where ρ denotes Euclidean distance. From Asser-

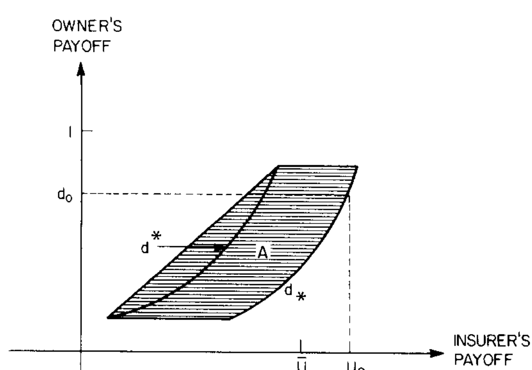

FIGURE 3.

tion E we now obtain that for every $\hat{\varepsilon} > 0$ there exists an event $\hat{E}$ with $\Pr(\hat{E}) > 1 - \hat{\varepsilon}$ and there exists T such that for every $t \geqslant T$, $\rho[(\bar{U}^t, \bar{d}^t), A]$ is arbitrarily small. Since d_* is a convex function, we can get this distance to be small enough to assure that the inequality

$$\bar{U}^t - (d_*)^{-1}(\bar{d}^t) < \varepsilon_0/2$$

is satisfied for $t > T$.

To establish the desired contradiction, it remains to be shown that

$$\limsup \bar{U}^t \leqslant U_0 + \varepsilon_0, \qquad \text{a.s. on } \hat{E}.$$

Note that $\alpha^t \to 0$ and therefore there exists a T_0 such that, for every $t \geqslant T_0$, we have

$$d_0 + \alpha^t < d_*(U_0 + \tfrac{1}{4}\varepsilon_0),$$

where $d_0 = d_*(U_0)$. To complete the proof, we now make use of Assertion B. Assume $\bar{U}^t > U_0 + (\varepsilon_0/4)$, and $t > \max\{T_0, T\}$. From the choice of T and T_0, $\bar{d}^t > d_*(U_0 + (\varepsilon_0/4)) > d_0 + \alpha^t$. Thus, under strategy $\tilde{f}$, a high premium $\pi^* = 1$ is being charged. Therefore $E[U^{t+1} \mid U^1,..., U^t, \hat{E}] \leqslant \max_c E[-c - D(c, w^t) | \hat{E}] \leqslant \bar{U} + (\varepsilon_0/4) \leqslant U_0 + (3\varepsilon_0/4)$. Thus, the premises of Assertion B are satisfied and we may also conlcude that in $\hat{E}$, whose probability exceeds $1 - \hat{\varepsilon}$, $\limsup \bar{U}^t \leqslant U_0 + (3\varepsilon_0/4)$, contradicting the original assumption about the strategy g. Therefore, recalling the choice of $\hat{\varepsilon}$, we conclude that $E[U^{t+1} | (\bar{U}^t, \bar{d}^t), \hat{E}] < U_0 + (\varepsilon_0/4)$. Thus, the premises of Assertion B are satisfied and we may also conclude that in $\hat{E}$, whose probability exceeds $1 - \hat{\varepsilon}$, the distance $\rho[(\bar{U}^t, \bar{d}^t), C] \to 0$ a.s. Therefore, $\limsup \bar{U}^t \leqslant U_0 + (3\varepsilon_0/4)$, contradicting the original assumption about the strategy g.

Proof of Theorem 2. Let a set $\Delta \subset \mathbb{R}^2$ be defined as follows:

$$\Delta = \{(x, y) \mid x \leqslant \pi - \bar{D}(c), y \leqslant u(-c - \pi) \text{ for some } c, \pi \in [0, 1]\}.$$

By our assumptions on D and u, Δ is a closed convex set. Furthermore, it follows from assumption (ii) on D (see Section 2) that $(0, \bar{U}) \in \Delta$. Let $(\tilde{f}, \tilde{g}) \in \tilde{S}$. For every t and for every state of information θ^{t-1}, we have

$$(E(h_1^t \mid \theta^{t-1}), E(h_2^t \mid \theta^{t-1})) \in \Delta,$$

which follows from the observation that the random variables $h_1^t \mid \theta^{t-1}$ and $h_2^t \mid \theta^{t-1}$ are generated either by a pair (c, π) with insurance purchased in period t or by the owner deciding, in period t, to decline the insurance offer, which leads to the payoff pair $(0, \hat{U})$ with $\hat{U} \leqslant \bar{U}$. Therefore, we may write

$$E[(h_1^t, h_2^t) \mid h_1^1,..., h_1^{t-1}, h_2^1,..., h_2^{t-1}] \in \Delta.$$

By Assertion F,

$$[H_1(\tilde{f}, \tilde{g}), H_2(\tilde{f}, \tilde{g})] \in \Delta,$$

so, by the definition of Δ there exists $c \in [0, 1]$ such that

$$\begin{aligned} H_2(\tilde{f}, \tilde{g}) &\leqslant u(-c - \bar{D}(c) - H_1(\tilde{f}, \tilde{g})) \\ &\leqslant u(-\hat{c} - \bar{D}(\hat{c}) - H_1(\tilde{f}, \tilde{g})) \end{aligned}$$

where $\hat{c}$ is some socially optimal level of care. This completes the proof.

5. Discussion

Now let us review the structure of the insurer's NCD strategy in Theorem 1. Consider the pair (π_0, c_0) appearing in the statement of the Theorem. If the insurer could somehow convince the owner of the asset to exercise care at a level c_0, then he would be happy to sell him insurance at a price π_0, and the resulting outcome would be to the benefit of both parties. In order to convince the owner of the asset to set care at c_0, the insurer announces that the premium for insurance in any given period may depend on the owner's record of claims up to that period. Theorem 1 shows that a relatively simple rule of this kind will indeed convince the owner of the asset to exercise the proper care c_0. The insurer picks two premium levels, high and low, with the low level set at π_0 and the high level, say, π^*, set to satisfy $u(-\pi^*) \leqslant \bar{U}$. (For the owner of the asset, π^* is a penalty premium, because he can attain the utility level $\bar{U}$ even without recourse to insurance.) In any given period, the insurer now charges a low premium π_0 if the owner's average past claims have *not* been "excessive." Otherwise, he charges the high premium π^*. The insurer's problem is to define exactly what will be meant by a past claims record being excessive. If the definition of "excessive" should be too strict, then the owner of the asset would end up paying a penalty premium quite often, even when he exercises proper care. On the other hand, if the definition of "excessive" should be too lax, then the owner of the asset would be able to set the level of care below c_0 sufficiently often (thereby obtaining an average utility level in excess of U_0) while at the same time keeping his claims record below the level that causes insurance premia to go up. For example, suppose that the average size of the owner's past claims is to be considered "excessive" if it falls above $\bar{D}(c_0)$. Then, with probability 1, the owner of the asset will be charged a penalty premium infinitely often, even if he consistantly exercises the proper care c_0. On the other hand, suppose that the average size of past claims is to be considered "excessive" only if it falls above $\bar{D}(c_0) + \varepsilon$ for some $\varepsilon > 0$. Then, the owner of the asset will be in a position to take advantage of this leniency and lower care to some level c satisfying $\bar{D}(c) < \bar{D}(c_0) + \varepsilon$ without the penalty premium affecting his long-run average utility. Thus, what the insurer must do is find an appropriate statistical instrument that would make it possible for him simultaneously to reduce the likelihood of two types of errors: One type consists of an erroneous detection of a (nonexistent) deviation from the proper care level c_0 and the other type consists of a failure to detect an actual (and prolonged) deviation from that proper care level. Theorem 1 provides such an instrument. In the tth contract between insurer and owner the penalty premium π^* will be charged if, and only if, the record of past claims filed shows an average size of claims exceeding $\bar{D}(c_0) + \alpha^t$, where $\{\alpha^t\}$ is a sequence of positive real numbers converging to 0 sufficiently slowly. Indeed, any sequence $\{\alpha^t\}$ converging to 0 and having the property that $\alpha^t \geqslant (2\lambda\sigma^2 \log\log t/t)^{1/2}$ for all t and for some $\lambda > 1$ will do. On the one hand, it follows from $\alpha^t \to 0$ that the owner of the asset cannot profitably deviate from the proper care level c_0. On the other hand, the slow rate of convergence of α^t provides a guarantee that, if the owner of the asset *should* consistently exercise the proper level of care c_0, then, with probability one, he will be charged the penalty premium π^* only finitely many times. (Note that if we take $\alpha^t = (2\lambda\sigma^2 \log\log t/t)^{1/2}$ but set $\lambda = 1$, then with positive probability, a properly careful owner will find himself penalized infinitely many times.) Informationally, the insurer's strategy is very simple: He uses only two premium levels, π_0 and π^*, and he must keep a running account only of the average size of the claims filed by the owner in those past periods where he (the owner) chose to buy insurance.

The existence of an insurance strategy having the desirable properties discussed above follows from a theorem in probability theory known as the Law of the Iterated Logarithm. This theorem has been used recently also in two other studies, and this seems to be a good place to comment briefly on the relationships between these studies and the present essay. In the first of these studies, Rubinstein [11] considers the question of how the legal system ought to treat an individual who is found to have committed an offense of the kind that may, with some probability, be committed inadvertently, even by the most scrupulously law-abiding citizens. (Failing to report an item in one's income tax return may be a case in point.) If the legal system is to act leniently, then individuals may find it advantageous to commit the offense willfully. On the other hand, if the legal system acts very stringently, then it may often find itself penalizing honest individuals who are the victims of random circumstances. Rubinstein shows that the legal system can successfully avoid these two predicaments by adopting a strategy of punishing an offender in period t if, and only if, the number of times where the offense had been committed in the past exceeds the value $t(p + \alpha^t)$, where p is the probability of the offence being committed accidentally and

$\{\alpha^t\}$ is a sequence of "safety margins" obtained from the Law of the Iterated Logarithm. The main structural difference between the framework of Rubinstein's paper and the framework being considered here is that, in the former, a willfully committed offense is immediately detected, so an individual cannot plan to commit the offense willfully in the hope of being mistaken for an honest law-abiding citizen. In our present framework, negligent behavior on the part of the owner of the insurable asset can never be detected as such. Thus, it appears that the owner of the asset might be in a position to gain from undetected negligent behavior. The long proof of part (4) in Theorem 1 shows that this is not possible. Another difference is that, in the framework being considered here, the owner of the asset can choose to refrain altogether from buying insurance in any given period. Assertion G in the previous Section is required in order to show that the owner cannot use this option to advantage.

Radner [10] also uses the Law of the Iterated Logarithm as a means for detecting deviant behavior that is not directly observable. Radner is concerned with what he calls "cooperative agreements" which lead to long-run payoffs that dominate the payoffs associated with a single-period Nash equilibrium. He shows that for every $\varepsilon > 0$ there exists a pair of strategies under which, if the single game is repeated sufficiently many times, then deviant behavior cannot yield a long-run payoff which exceeds the payoff from the cooperative agreement by more than ε. There are several differences between Radner's framework and the framework being considered here but, in our view, the main difference lies in Radner's use of a so-called "trigger strategy" (also known in the folklore as a "grim" strategy) for the principal: Punishment, once it occurs, continues relentlessly, until the game ends. In an infinitely repeated version of Radner's model, punishment would have to be inflicted until the deviating agent's payoff is driven down virtually to the inferior level obtained in the single-period Nash equilibrium. In the present essay, we have concentrated on what we have called "no-claims discounts" strategies which, we feel, are closer to observed behavior than trigger strategies. In our framework, a deviating individual is punished only as long as his record appears unreasonable, and punishment ceases as soon as this record is restored to within reasonable bounds. The strategies being considered here have a certain "perfectness" property which is not shared by Radner's trigger strategies: For all t, and for all conceivable histories up to period t, the insurer's NCD strategy will elicit the proper care level from the owner from period t onwards.

This last comment makes it clear that NCD strategies are by no means the only ones which make it possible for the insurer to enforce a proper care level c_0 in return for a lower premium for insurance π_0. Indeed, we are indebted to R. J. Aumann and A. Neyman for pointing out to us that, in order to establish the existence of *some* strategy that enforces a proper level of care, it is not even necessary to appeal to the Law of the Iterated Logarithm. A more elementary tool—the Borel–Cantelli Lemma—would suffice for that purpose. To see this, let $\{p^t\}$ be a sequence of real numbers satisfying $0 \leqslant p^t \leqslant 1$ for all t, and assume that the series $\sum p^t$ converges. Now suppose that the owner of the insurable asset has already suffered through $t-1$ epochs of punishment. Then, what the insurer must do in any given period is to calculate the probability of the owner's record being as it really is, or worse, on the assumption that he (the owner) has been exercising proper care c_0 ever since the last time he had been punished. If this probability is less than p^t, then a new epoch of punishment begins (i.e., the owner is charged a high premium π^*) and this punishment continues relentlessly until it is known with certainty that the owner's long-run utility has been reduced to a level $\bar{U} + \varepsilon^t$, where $\varepsilon^t \to 0$. It can be shown, using the Borel–Cantelli Lemma, that this insurance strategy enforces a proper level of care as a best response. Here again, what is lost is the NCD structure of the insurer's strategy.

Finally, it ought to be mentioned that, with minor modifications, the present framework can be recast to fall within the realm of repeated games. Using the terminology of the theory of repeated games, our Theorems 1 and 2 amount to the assertion that the set of pairs of payoffs which are individually rational in the single game coincides with the set of equilibrium pairs of payoffs in the repeated game. In this sense, our theorems constitute a new version of what has come to be called the Folk Theorem of repeated games (see Hart [5]) for a class of games with imperfect information, where, in any given period, a player must make his move without full knowledge of the moves made by the other player in the preceding periods.

References

1. D. Blackwell, An analog of the minimax theorem for vector payoffs, *Pacific J. Math.* **6** (1956), 1–8.
2. L. Breiman, "Probability," Addison–Wesley, Reading, Mass., 1968.
3. D. Freedman, Another note on the Borel–Cantelli lemma and the strong law, *Ann. Probab.* **1** (1973), 910–925.
4. P. R. Halmos, "Measure Theory," Van Nostrand, Princeton, N.J., 1950.
5. S. Hart, Lecture Notes: Special topics in game theory, Stanford University, Stanford, Calif., 1979.
6. B. Hølmström, Moral hazard and observability, *Bell J. Econ.* **10** (1979), 74–91.
7. M. Loeve, "Probability Theory," Volume II, 4th ed., Springer-Verlag, Berlin, 1978.
8. J. A. Mirrlees, The theory of moral hazard and unobservable behavior, Mimeo, Nuffield College, Oxford, 1975.
9. M. Pauly, Overinsurance and public provision of insurance: The roles of moral hazard and adverse selection, *Quart. J. Econ.* **68** (1974), 44–62.

10. R. Radner, Monitoring cooperative agreements in a repeated principal–agent relationship, *Econometrica* **49** (1981), 1127–1148.
11. A. Rubinstein, An optimal conviction policy for offenses that may have been committed by accident, *in* "Applied Game Theory" (S. J. Brams, A. Schotter, and G. Schwodiauer, Eds.), Physica-Verlag, Würzburg, 1979, 406–413.
12. S. Shavell, On moral hazard and insurance, *Quart. J. Econ.* **93** (1979), 541–562.