Optimal Insurance Contracts under Moral Hazard[*]

Álvaro Parra and Ralph A. Winter[†]

August 30, 2023

Abstract

This chapter surveys the theory of optimal insurance contracts under moral hazard. Moral hazard leads to insurance contracts that offer less than full coverage of losses. What *form* does the optimal insurance contract take in sharing risk between the insurer and the individual: a deductible, or co-insurance of some kind? What are the factors that influence the design of the contract? Posed in the most general way, the problem is identical to the hidden-action principal-agent problem. The insurance context provides structure that allows more specific implications for contract design. This chapter reviews the static models of optimal insurance under ex ante and ex post moral hazard, and the potential for repeated interactions and long-term contracts to mitigate moral hazard. We also discuss optimal insurance for liability risk, under which not just insurance but limited liability offers protection against large losses.

1 Introduction

The Concept of Moral Hazard

This chapter offers a synthesis of the economic theory of moral hazard in insurance, with a focus on the design of optimal insurance contracts. In this context, moral hazard refers to the impact of insurance coverage in distorting incentives. The topic divides naturally into *ex ante* moral hazard and *ex post* moral hazard. Ex ante, an individual facing the risk of an accident such as a home fire, a car accident or a theft, can generally take actions to reduce the risk. Without insurance, the costs and benefits of accident avoidance, or precaution, would be

[*]Chapter in forthcoming Handbook of Insurance, 3rd Edition, Georges Dionne, ed. We are grateful to two reviewers for very insightful comments.

[†]Sauder School of Business, UBC

internal to the individual. The incentives for avoidance would be optimal. With insurance, however, some of the accident costs are borne by the insurer. The insured individual, bearing all of the costs of accident avoidance but only some of the benefits will then under-invest in precaution. This is ex ante moral hazard. Ex post, once the event of a need for medical care (for example) has occurred, an individual will spend more resources on care the greater the proportion of those expenses covered by insurance. Insurance covering the replacement of lost or stolen items is also subject to ex post moral hazard.

An insurance contract may specify levels of precaution (the number of fire extinguishers, the frequency of inspection of equipment and so on). And it may constrain expenditures ex post. If insurance contracts were complete in the sense of specifying the individual's care in all dimensions and in all future contingencies prior to the accident, and expenditures in the event of need, then moral hazard would not be an issue. But insurance contracts do not specify precisely the precautions to be taken by the insured. An automobile insurance policy, for example, does not specify the attention and focus that a driver must dedicate to driving safely. A health insurance policy does not dictate the diet or exercise routine of the individual covered. The precaution decisions taken after an insurance contract is signed are inefficient because of the positive externality on the insurer entailed in greater precautionary effort. The optimal insurance contract must be designed, within the constraints of asymmetric information and enforceability, in anticipation of rational decisions on precaution and insured expenditures.

The term moral hazard originated in the insurance context that we study here. The domain of the term has expanded, however, to include virtually all contracts, well beyond the traditional context of insurance contracts. Labor contracts, for example, are designed with the knowledge that the effort, diligence, and enthusiasm of the employee cannot be specified completely in the contract and instead must be induced through incentives provided in the contract. The relationships between a homeowner and a contractor, a lawyer and a customer, partners in a joint venture, the editor of this volume and the authors of this chapter, are all subject to moral hazard. Even a marriage is subject to moral hazard in that costs are imposed on one marriage partner whenever the other one shirks.[1] Moral hazard in general refers to the distortions resulting from externalities among parties to an incomplete contract, on the decisions taken after the contract.

Moral hazard is distinguished from externalities in general by the existence of a contractual relationship between the decision maker and the party exposed to the externality. An insurance company and the insured individual have a contract; a polluting firm located up-

[1] In an ideal marriage, however, costs imposed on the spouse are internalized in an individual's own utility function. Love solves the moral hazard problem.

stream from a city taking water from a river generally has no contractual relationship with the users. Even this limit on the definition of moral hazard is tenuous, however. Moral hazard can encompass any externality if one adopts the broad theory of social contract (Hobbes 1651). Law and social norms can be interpreted together as a contract specifying the rights and obligations of individuals in a society. All individuals are in the social contract, and externalities are the consequence of incompleteness in the social contract. The ethical adjective *moral* in moral hazard is suggestive of this broader interpretation, but the phrase moral hazard has a much narrower origin in the insurance industry. Moral hazard, as describing the tendency for insurance to create incentives for individuals to be less careful in protecting themselves or property against risks, gained frequent usage in the late nineteenth and early twentieth centuries with the growth of private and social insurance in Europe and the United States (Dembe and Boden 2000).[2]

Kenneth Arrow (1963) pioneered the economic analysis of moral hazard. Pauly (1968) also offered an important early contribution. A large part of the microeconomics literature over the past twenty-five years has been devoted to the implications of incomplete contracts and incentives. The literature on the principal-agent theory, beginning with Holmström (1977), Holmström (1979), Shavell (1979a), Mirrlees (1975; 1976), is central in this movement. In returning to the original context in reviewing the implications of moral hazard for insurance contracts, we draw on the developments in this literature. In accepted terminology, it is the hidden-action version of the principal-agent model that provides the structure for optimal insurance contracts under moral hazard. Shavell (1979b) is an early link between the general contracting theory and insurance.

Moral hazard is distinguished from adverse selection, another form of asymmetric information, by the timing of the informational asymmetry. In moral hazard problems, the insurer and the insured are symmetrically informed at the time of contracting. The indi-

[2] It has been suggested that the etymology of the term moral hazard may involve a second historical use of the term *moral* (Dembe and Boden 2000). Daniel Bernoulli (1738; 1954) in his resolution of the St. Petersburg paradox first posed by Nicholas Bernoulli in 1714, applied a theory that he referred to as "the theory of moral value." *Moral* referred in this case to the subjective or psychological value placed on the gain in an individual's wealth. The *moral expectation* was distinguished from the *mathematical expectation*. Today we refer to Bernoulli's moral value as the Bernoulli utility or Von Neumann-Morgenstern utility of wealth. Bernoulli's use of the term "moral" as meaning subjective or personal value is consistent with the usage of moral in the eighteenth and nineteenth centuries as meaning in accordance with customs or the norms of human conduct, rather than ethical as in current English usage. As Dembe and Boden note at p.261, "The classical eighteenth-century mathematical analysis of subjective utility in risk-bearing situations can thus be considered as essentially value-neutral, despite being couched in the language of *moral* values and expectations." It is tempting for an economist, who considers maximizing behavior under incomplete contracts simply to be rational behavior, to trace the use of the term moral hazard in economics to the essentially value-neutral language of moral expectation in the eighteenth century. Dembe and Boden place considerable weight on this possibility, although Kenneth Arrow (1963) is quite explicit in citing the prior insurance literature, in introducing the moral hazard terminology to the economics literature.

vidual's precaution decisions taken after the contract are not observed by the insurer, or at least not by a court enforcing the contract. In asymmetric information models, in contrast, insured individuals have private information at the time of contracting.[3] The moral hazard / adverse selection distinction is cast in the contracting literature as the difference between hidden action (private information regarding actions of the insured individual, or agent) and hidden information (private information on the characteristics of the insured individual).

The Questions

The design of optimal insurance contracts in static environments: It has been well known since Professor Arrow's classic 1963 paper that the contractual response to moral hazard is to leave some of the risk uninsured, i.e., to leave some of the risk with the risk-averse individual rather than transferred entirely to the insurer. Leaving the individual with some share of the consequences of a marginal change in precaution improves the individual's incentives to take precautions. The optimal contract will balance the risk-sharing benefits of greater insurance with the incentive benefits of less insurance. Our first question is the central issue for the design of optimal insurance contracts is what *form* the risk-sharing takes. Will optimal insurance involve a contract in which the individual bears the cost of all losses, up to some limit? This is a *deductible*. Will the optimal contract involve full insurance of marginal losses up to some coverage limit? Or will it involve some continuous sharing of the marginal accident costs?

We begin with the most general model of an insurance contract in a static setting with ex ante moral hazard. This is essentially the general principal-agent model, applied to insurance. The insurance context imposes a structure on the general principal-agent contract that yields predictions, such as the following (Holmström 1979): a pure deductible insurance contract is optimal if precautionary efforts affect the probability of an accident but not the severity of the random losses conditional upon an accident.

The application of the general model to insurance begins with the simplest theory: individual effort affects the probability of a loss of known size. This approach focuses on ex ante moral hazard and leads to loss sharing between the insurer and the insured. Where individual effort affects the likelihood of a loss, but the loss size is random but *verifiable* ex post, the loss sharing takes the form of a deductible. We consider as well the case in which individual effort affects the likelihood of a loss but the loss size is *not verifiable* ex post — so that the contract must take the simple form of a specific payment in the event of a loss. The non-moral hazard benchmark in this case involves an exogenous probability of a loss.

[3] Models of insurance markets with asymmetric information are reviewed by Georges Dionne, Neil Doherty and Nathalie Fombaron in a different chapter in this handbook.

We should note that this non-moral hazard benchmark involves an insurance contract with a payment in the event of an accident that is larger than the expected loss (conditional upon an accident) for an individual with a utility that has the property of decreasing absolute risk aversion. Introducing moral hazard reduces the insurance payment in the optimal contract.

Having considered the structure in which effort affects only the probability of an accident and remaining within the ex ante moral hazard framework, we remain within the ex ante moral hazard content and consider the assumption that ex ante care affects the (random) *severity* of a loss, rather than the probability of loss. The loss is assumed to be verifiable. This yields a contract that is, in one respect, the opposite of a deductible, which, as we have discussed, is the solution to the endogenous probability case: losses up to some critical value are fully covered. Higher losses are partially covered. In a more general setting, we review conditions under which coinsurance with risk sharing between the insurer and the insured is optimal.

We then turn to the case of ex post moral hazard, motivated by its most important example, medical care insurance, in which the medical expenses in the event of an accident are chosen by the individual (and their doctor). Zeckhauser (1970) offers the basic model of ex post moral hazard. We review this setting taking a more modern revelation-principle or direct-mechanism approach to optimal contracting. Within the ex post moral hazard context, Ma and Riordan (2002) develop a model that captures very clearly the trade-off between insurance and the incentives to spend efficiently on ex post care. We review the analysis by Ma and Riordan of both demand side management (coinsurance, or incomplete insurance) and supply side management through the provision of incentives to physicians.

The design of optimal insurance contracts in dynamic environments: The move from static theories of insurance to dynamic (multi-period) theories lead us to two questions. First, when multi-period contracts are feasible, is there an incentive to enter a multi-period insurance contract, as opposed to relying on a sequence of short term contracts, to balance incentives and insurance? What are the characteristics of an optimal long term insurance contract? The answer under the simplest set of assumptions is that there is, surprisingly, no gain at all to multi-period contracting. Entering one-period contracts does just as well. But departures from this set of assumptions in a number of dimensions lead to a role for long term contracts as a response to moral hazard.

The second question within dynamic environments takes as an assumption that only short term contracts are available. In this case, the question is whether *repetition* of the market for short-term insurance contracts mitigates moral hazard through the threat of increases in future insurance premiums in response to accident claims? (Since we observe such increases, the question could be phrased as *how* future insurance markets enhance incentives.) We suggest through an initial, stylized model that a useful approach to this question is based on the application of Holmström's classic career-concerns model (Holmström 1999). With hidden action alone — the basis for moral hazard in static insurance models — repeated markets do nothing. Ironically, however, greater asymmetry in information — the addition of an assumption of hidden characteristics in addition to hidden information — leads to an explanation of how future insurance markets can enhance incentives. The intuition is that exerting additional effort today leads to a lower risk of an accident, which then increases the likelihood that the insurance market tomorrow will identify the individual (stochastically) as a better type. This lowers expected future premiums. The prospect of lower future premiums in response to the avoidance of an accident today improves incentives to exert effort today.

Liability Insurance: An important source of risk in modern society, especially in the U.S., is the risk under tort law of liability for damages from accidents. The economics of liability insurance involves three special considerations, compared to standard "first-party" insurance. First, limited liability offers an alternative means (compared to insurance alone) of protection against risk. How does limited liability interact with the incentives to purchase insurance? Second, in contrast to mainstream insurance theory, which involves only two agents — the insured and the insurer — liability insurance brings into play a third party affected by insurance purchases and incentives: the potential accident victim who stands to be compensated for accident damages. Third, in the case of liability insurance individual incentives to take care to avoid an accident are complex even before we introduce insurance. The incentives in the pre-insurance world involve externalities imposed by an individual whose decisions potentially cause an accident or affect the risk of an accident (such an individual is called a "tortfeasor") on the potential victim of the accident, and on other tortfeasors. Incentives are affected by the tort rule. For example, strict liability and negligence rules involve distinct effects on incentives for both care and engaging in the activity involving accident risk. The interaction of incentives, tort rules, limited liability and insurance — the economics of tort law — is a large and complex area. We offer a brief and relatively informal overview of the key insights in this area. We refer to Shavell (1987) for a comprehensive and illuminating treatment of the economics of tort law.

The conclusion to this chapter offers an outline of additional topics in the economics of moral hazard and insurance that would benefit from further research. These include the implications of multiple dimensions of care for moral hazard; moral hazard in insurance under the assumption of uniform pricing or non-exclusivity of insurance coverage; and moral hazard on the *supply* side of insurance markets — distortions in, for example, insurers' investment decisions. Finally, we offer a conjecture on the impact of corporate insurance on risk management including care decisions. We suggest that corporate insurance should often be interpreted as the purchase of risk-management services from insurers rather than simply protection against losses.

2 Ex Ante Moral Hazard: a General Distribution of Losses

A general formulation of the optimal insurance problem under moral hazard is a simple adaptation of the standard principal agent model under hidden action (Holmström 1979; Bolton and Dewatripont 2005). A risk-averse individual faces a random loss x with a distribution that depends upon the effort, a, that the individual takes ex ante to avoid the loss. Let this distribution be $F(x;a)$ with continuous density $f(x;a)$ on support $[0,\overline{x}]$. Assume that increases in a reduce the random loss in the sense of first-order stochastic dominance: $\partial F(x;a)/\partial a \leq 0$ with the inequality strict for a positive measure of effort levels. The individual's utility over wealth w and effort is expressed as $u(w) - v(a)$, with $u' > 0, u'' < 0, v' > 0$ and $v'' > 0$. The individual's initial wealth is w.

The separability of the utility function in wealth and effort is one of two common formulations of preferences in principal-agent models. The second is the opposite: that costs of precaution are entirely pecuniary, with utility given by $u(w-a)$.[4]

Throughout this chapter we assume that an individual has access to a competitive insurance market that will provide any contract yielding non-negative expected profit. The essence of moral hazard is that the individual's effort is not contractible. The insurance contract specifies only an up-front premium, r, and the insurance coverage $I(x)$ that will be provided for each realization of the loss x. The insurance contract is exclusive in the sense that the individual enters only one contract. The individual's effort is determined after the contract by the individual acting in her own interest given the contract. We follow the standard approach in contract theory in writing the contract *as if* effort entered the contract, but subject the choice of contract to the *incentive compatibility constraint* that the choice of effort be the level that the individual will actually choose given the rest of the contract. The optimal contract is the solution to the following.

$$\max_{r,I(x),a} \int_0^{\overline{x}} u(w - r - x + I(x)) f(x;a) dx - v(a) \tag{1}$$

[4]Ma and Riordan (2002) adopt a general assumption on preferences that accommodates both non-pecuniary and pecuniary costs of effort.

subject to

$$a \in \arg\max_{e} \int_0^{\overline{x}} u(w - r - x + I(x)) f(x; e) dx - v(e) \tag{2}$$

$$\int_0^{\overline{x}} I(x) f(x; a) dx - r \leq 0 \tag{3}$$

The individual chooses the contract to maximize expected utility (1) subject to the incentive compatibility constraint (2) and the individual rationality or non-negative expected profit constraint (3). The notation "$a \in \arg$" allows for the possibility that there are multiple solutions to the agent's maximization problem. At this level of generality, not much can be said about the optimal insurance coverage, $I(x)$. Some insight can be gained into the optimality conditions by assuming that the first-order condition for the agent's incentive compatibility constraint is not only necessary but sufficient for the agent's optimum — in other words, that the second-order condition holds. While a common step in the analysis of principal-agent problems, the assumption, unfortunately, is ad hoc. The conditions that have been established to guarantee sufficiency of the "first-order approach" to the principal-agent model are strong. Rogerson (1985a) showed that the first order approach to principal-agent problems is valid under two additional assumptions: the *monotone likelihood ratio property* (MLRP) and the concavity of the distribution function in a, for any x.[5] The MLRP here is the following

$$\frac{d}{dx} \frac{f_a(x; a)}{f(x; a)} \leq 0 \tag{4}$$

The MLRP and the assumption of concavity of F assure that the agent's objective, $\int_0^{\overline{x}} u(w - r - x + I(x)) f(x; a) dx - v(a)$, is concave. The MLRP is a reasonable assumption, but as Jewitt (1988), Bolton and Dewatripont (2005) and others point out, the concavity condition (which is a convexity condition, in the standard principal-agent formulation) is quite restrictive.[6]

We follow convention in adopting the first-order approach in spite of its restrictiveness. Under this approach, we can replace the incentive compatibility constraint (2) with the first-

[5] The assumptions of concavity of the distribution function F and the MLRP for the random loss x correspond to the assumptions in a conventional principal-agent model of the *convexity* of the distribution function and an MLRP with the opposite inequality. Here, x is a loss; in the conventional principal-agent problem, x is profit.

[6] Jewitt (1988) provides sufficient conditions for the first-order approach beyond the restrictive convexity condition, and including the observation by the principal of multiple relevant statistics. Ábrahám et al. (2011) extend the sufficiency of the first-order approach to two period models when agents can perform hidden borrowing and lending. The first-order approach is sufficient when the model satisfies MLRP, Log-convexity of effort on output, and non-increasing absolute risk aversion utility.

order condition

$$\int_0^{\overline{x}} u(w - r - x + I(x))f_a(x;a)dx - v'(a) = 0 \tag{5}$$

The optimal insurance problem in the presence of ex ante moral hazard is thus the maximization of (1) subject to (3) and (5).

For a given insurance coverage function $I(x)$, the constraints (3) and (5) define a system of equations in a and r. We let the solution in (a, r) to this system be represented by the operators $a = A[I(x)]$ and $r = R[I(x)]$. (Appendix A contains a proof of the uniqueness of this solution in a neighborhood of the optimum.) We can substitute these operators into the objective function (1) to obtain an *indirect utility function* over the coverage function, $I(\cdot)$:

$$EV[I(x)] = \int u(w - R[I(x)] - x + I(x))f(x; A[I(x)])dx - v(A[I(x)]) \tag{6}$$

The optimal insurance problem is the choice of the function $I(\cdot)$ to maximize (6). We can obtain the optimality condition, using the standard approach to optimization on a space of functions, by considering a small deviation in the optimal coverage function $I^*(\cdot)$. Let us impose a small variation $\epsilon h(x)$ on top of the coverage policy $I(x)$, to that the new coverage policy is $I(x) + \epsilon h(x)$. Given $I(x)$ and $h(x)$, the indirect utility can be rewritten as the function of this small variation ϵ.

$$E\widetilde{V}(\epsilon) = \int u(w - \tilde{r}(\epsilon) - x + I(x) + \epsilon h(x))f(x; \tilde{a}(\epsilon))dx - v(\tilde{a}(\epsilon))$$

with $\tilde{r}(\epsilon)$ and $\tilde{a}(\epsilon)$ defined as the solutions in r and a to the following two equations

$$\int u(w - r - x + I(x) + \epsilon h(x))f_a(x;a)dx = v'(a) \tag{7}$$

$$r = \int (I(x) + \epsilon h(x))f(x, a)dx \tag{8}$$

Next we decompose the marginal effect of a small variation ϵ. Taking the derivative respect to ϵ, and evaluating at 0, we obtain

$$\begin{aligned} E\widetilde{V}'(0) = \int \Big[& u'(w - r - x + I(x))(-r' + h(x))f(x;a) \\ & + u(w - r - x + I(x))f_a(x;a)\tilde{a}'\Big]dx - v'(a)\tilde{a}' \end{aligned} \tag{9}$$

r' and a' can solved by differentiating (7) and (8).

$$\tilde{a}'(0) = \frac{\int u'h(x) f_a dx - \int u' f_a dx \cdot \int h(x) f dx}{v''(a) - \int u f_{aa} dx + \int u' f_a dx \cdot \int I(x) f_a dx} \tag{10}$$

$$\tilde{r}'(0) = \int I(x) f_a dx \cdot \tilde{a}' + \int h(x) f dx \tag{11}$$

Substituting (10) and (11) back to (9), we have

$$\begin{aligned} EV' = & \int u f_a dx \cdot \tilde{a}' \\ & - v'(a)\tilde{a}' \\ & - \int I(x) f_a dx \cdot \tilde{a}' \cdot \int u' f dx \\ & - \int h(x) f dx \cdot \int u' f dx \\ & + \int u' h(x) f dx \end{aligned} \tag{12}$$

The five terms in this expression reflect the marginal effect of a small variation of coverage, with ex ante moral hazard.[7] The terms represent:

(a) a change to the expected utility due to the shift of loss distribution;

(b) the disutility of marginal effort;

(c) the utility cost of the change in the premium due to the shift of loss distribution;

(d) the utility cost of a change of premium due to an increased level of coverage;

(e) the utility impact of a change in the level of coverage.

The changes in expected utility due to (a) and (b) are exactly offsetting, from the incentive compatibility first-order condition (5) — an envelope-theorem effect. This leaves only the last three terms to represent the marginal impact of a change in coverage on utility. The last two terms would appear in a complete contract, without moral hazard; these represent (at the fixed care level) the utility cost of the premium paid for marginal additional coverage, and the benefit of additional coverage. This leaves only the middle term of (12), the change in the premium due to the endogenous change in care level, as reflecting the moral hazard problem. This term is not only the utility cost of the change in ex ante premium reflecting the market's rational forecast of the shift in the lost distribution, it represents the expected value in utility terms of the externality that the individual imposes on the insurance company

[7] This decomposition parallels equation (7) in Shavell (1979), which considers the simpler insurance problem with only one possible value for the loss if an accident does occur.

in choosing the care level once the contract is entered into.

Starting with (12), we can establish three propositions in the general ex ante model. First, even in the presence of moral hazard, an insurance contract offers positive coverage. Starting from zero coverage, $I(x) = 0$ and letting $h(x) = x$, i.e. full coverage, the right hand side of (12) represents the marginal gain from moving ε towards full coverage starting at $\varepsilon = 0$. With these values for $I(x)$ and $h(x)$, the term (c) drops out. The first two terms continue to sum to zero, so that the right hand side of (12) equals $-\int xf dx \cdot \int u' f dx + \int xu' f dx = cov(u', x) > 0$. (The covariance is positive since a higher x leads to a lower wealth level, and u' is decreasing in wealth because of the concavity of u.) Zero coverage is therefore dominated by a marginal amount in coverage at all values of loss. Moral hazard never eliminates the value of insurance. Hence our first general principle:

In the general model, some insurance is optimal even in the presence of moral hazard.

The second proposition is that moral hazard will always lead to partial coverage. Full coverage is never optimal. To see this, evaluate the right hand side of (12) with $I(x) = x$ and any $h(x) < 0$; that is, consider a marginal *reduction* in coverage, starting from full coverage. With a reduction, rather than an increase, in coverage the right hand side of (12) changes sign. And with the *marginal* reduction in coverage, the last two terms drop out because at full coverage wealth is invariant to x and therefore u' is constant in these two equations. This leaves only the negative of term (c): $\int xf_a dx \cdot \tilde{a}' \cdot \int u' f dx = \partial [\int xf(x,a)dx]/\partial a \cdot \tilde{a}' \cdot \int u' f dx < 0$.[8] Thus, a marginal reduction in coverage starting from full coverage will yield a positive expected utility gain equal to the expected utility value of the gain in efficiency from reducing moral hazard. Full coverage is always dominated.

In the general model, full insurance is never optimal in the presence of moral hazard.

Define the coinsurance in an insurance policy $I(x)$ as $x - I(x)$, i.e. as the share of the loss that the individual must bear. The third proposition in the general ex ante moral hazard problem is about the monotonicity of coinsurance in the size of the loss. Intuitively, one might expect that coinsurance is increasing in x since this would give the individual incentive to avoid high losses, which involve the highest cost to the insurer. In general, however, coinsurance is not monotonic in an optimal insurance policy. Suppose, for example, that there are four possible realizations of loss: $1, 2, 3$ and 4; that with zero care the distribution

[8]To elaborate on the proof of this inequality, note that $\int f dx = 1 \Rightarrow \int f_a dx = 0 \Rightarrow E[f_a/f] = \int (f_a/f) f dx = 0$ so that $\int xf_a dx = \int x(f_a/f) f dx = E[x \cdot (f_a/f)] = cov(x, f_a/f) < 0$, by MLRP. Turning to the second term, $\tilde{a}' > 0$ since less coverage leads to more effort. The third term is positive since u' is positive.

of loss on the support $\{1, 2, 3, 4\}$, is $(.1, .4, .1, .4)$. An increase in care, we suppose would move the distribution closer to $(.4, .1, .4, .1)$. This is a first-order stochastic dominant shift downwards in the random loss, with greater care, so it is not unreasonable.[9] But it is easy to verify that the optimal insurance policy leaves the individual with more loss at the realizations 2 and 4 than it does at the realizations 1 and 3.

The key condition sufficient for monotonicity of coinsurance, and violated by the example, is the monotone likelihood ratio property, (4), which we have been assuming in our derivation. The third proposition on the general ex ante moral hazard problem, following from a standard result in the principal-agent model, is that optimal coinsurance is non-decreasing in the size of the loss under the MLRP.

The first-order condition for the optimal $I(x)$ can be found by solving the maximization of (1) subject to (5) and (3) via a point-wise Lagrangian, yields first-order conditions that can be reduced to the following:

$$u'(w - r - x + I(x)) \left[1 + \mu \, \frac{f_a(x; a)}{f(x, a)}\right] - \lambda = 0 \tag{13}$$

Without the incentive compatibility constraint, i.e., if $\mu = 0$, we would get the Arrow-Borch condition for first-best optimal insurance, that u' be constant across states (Arrow 1971; Borch 1962). This implies full insurance minus a constant: $I(x) = x - k$. *With* the incentive compatibility constraint, insurance is less than first-best. By the MLRP, $f_a(x; a)/f(x; a)$ is non-increasing in x. Since u' is strictly decreasing, (13) then implies that $x - I(x)$ is non-decreasing in x.

In the general model, under MLRP, the individual's dollar share of the loss is non-decreasing in the size of the loss.

The likelihood ratio enters because the optimal contract rewards to the extent that individual incentives matter ($\mu > 0$), the individual is "punished" more severely via reduced coverage in states for which a reduction likelihood is sensitive to increased effort. This encourages effort by deviating from full insurance in the most efficient way.

3 Ex Ante Moral Hazard in Special Cases

More specific predictions about the form of an optimal insurance policy follow from additional assumptions on the distribution of losses. A natural structure on insurance losses is a *two-stage compound lottery*: an accident occurs or not, and then conditional upon an accident

[9] This type of distribution can easily result from an exogenous uncertainty that has a bi-modal distribution.

nature draws from a random distribution of losses. Care can affect either stage of the lottery. We follow Ehrlich and Becker (1972) in distinguishing between care taken to reduce the probability of an accident, and care to reduce the (random) size of the loss contingent upon an accident. In Ehrlich and Becker's terminology, the former is *self-protection.* The latter these authors call self-insurance, but we will use the term *loss reduction*, because self-insurance has a different meaning in the insurance literature. Both types of care are important in various settings. Expenditures on fire sprinklers reduce the size of a loss, but not the probability of a fire. Expenditures on burglar alarms or security systems reduce the probability of a theft, whereas the decision not to buy expensive silverware reduces the loss if there is a theft. In the important case of earthquake insurance, *all* precaution is loss-reducing. Driving an automobile more slowly and carefully and avoiding driving entirely if road conditions are bad reduces both the probability of an accident and the costs of an accident should it occur.

We consider, in turn, the implications of moral hazard on these two types of care, reviewing first Holmström's key result on the form of optimal insurance contracts under self-protection. I then exploit the two-stage structure for an additional question: the optimal insurance contract when an insurer can observe the event of an accident, but cannot observe the size of the loss.

3.1 Self-Protection and Moral Hazard: the optimality of deductibles

Self-protection refers to the case where an individual can take effort, a, to affect the probability, $p(a)$, of an accident, but not $F(x)$, the distribution of losses conditional upon there being an accident. Many risk situations fit this description. A driver may be constrained to drive at a particular speed on the freeway but be careless to some degree in his driving, or in how long he drives while tired. In this situation, the probability of an accident is affected by care, but the random severity of the accident if it does occur may depend very little, if at all, on care.

Holmström (1979) shows that under self-protection, the optimal insurance policy is a deductible, d, with full coverage above the deductible. In other words, the optimal $I^*(x)$ satisfies $I^*(x) = \max(0, x-d)$ for some deductible, d. Holmström takes a first-order approach to this problem, without assuming explicitly a set of assumptions under which the first-order approach is valid. A basis for the first-order approach is relatively straightforward in this special case, however. We can *define* the level of care as the extent to which the probability of an accident is reduced below the probability, p_0, that the accident would occur with zero care. (There is no loss in generality in adopting this definition.) The probability of an accident is then linear in a, $p = p_0 - a$. An assumption that the disutility of care, $v(a)$, is convex with Inada conditions $v'(0) = 0$ and $v'(a) \to \infty$ as $a \to \overline{a}$ for some $\overline{a}$, is then enough to ensure that the agent's problem, $\max_a [1-(p_0-a)]u(w-r)+(p_0-a)\int_0^{\overline{x}} u(w-r-x+I(x))f(x)dx - v(a)$ is concave, and that the agent's first-order condition is both necessary and sufficient for the incentive compatibility condition.[10]

The optimal insurance contract under these conditions solves the following problem

$$\max_{r,I(x),a} \quad [1-(p_0-a)]u(w-r)+(p_0-a)\int_0^{\overline{x}} u(w-r-x+I(x))f(x)dx - v(a)$$

subject to the IC and zero-profit conditions:

$$u(w-r)-\int_0^{\overline{x}} u(w-r-x+I(x))f(x)dx - v'(a) = 0$$

$$r-(p_0-a)\int_0^{\overline{x}} I(x)f(x)dx = 0$$

as well as a non-negativity constraint

$$I(x) \geq 0$$

reflecting the assumption that an individual cannot be compelled to report an accident.

Ignoring, for the moment, the non-negativity constraint, the first-order condition on the choice of $I(x)$ implies that

$$u'(w-r-x+I(x)) = \lambda \Big/ \left[1-\frac{\mu}{(p_0-a)}\right]$$

which is independent of x. This can be achieved only if the individual bears the same net loss, $x-I(x)$, in all realizations of x. Care does not affect the distribution of losses conditional upon the event of an accident, so there is no reason to have the individual bear risk on the individual conditional upon that event. With moral hazard on the probability, $x-I(x)$ is positive in order to elicit care. Incorporating the non-negativity constraint then leads directly to the optimality of a pure deductible policy, $I(x)=\max(x-d,0)$ for some deductible d.

Where effort affects the probability of an accident but not the distribution of losses conditional upon an accident, the optimal insurance policy is a deductible with full coverage of marginal losses above the deductible.

Holmström shows that an insurance contract more general than a pure deductible, including a deductible but with possibly coinsurance for losses above the deductible, is optimal

[10]The same set of assumptions can be used to justify the first-order approach in Shavell (1979). Shavell adopts an assumption that the costs of care are pecuniary (i.e. a reduction in wealth) rather than purely non-pecuniary, as we assume here.

when the effort affects not only the probability of an accident but the losses conditional upon an accident as well; all that is required for a deductible is that there be an atom at 0 in the distribution of losses. The co-insurance of losses above 0 provides incentives for an individual to exert effort, as per the principal-agent model that we outlined.[11]

3.2 Loss reduction and moral hazard

The opposite assumption to the previous case isolates the impact of optimal contracting of effort that affects the distribution of losses conditional upon an accident, but not the probability of the accident itself. Insurance to replace household furnishings in the event of an earthquake illustrates this case. A homeowner cannot affect the chance of an earthquake, but can take measures to reduce the contingent losses. Investing in seismic protection reduces losses and even purchasing less expensive home furnishings reduces losses.

Let us take the first-order approach to this incentive contract design problem. Assume that an individual faces with exogenous probability p a loss that is distributed with distribution $F(x;a)$. An increase in effort shifts $F(x;a)$ downwards in the sense of first-order stochastic dominance. The individual in this case maximizes $[1-p]u(w-r) + p\int_0^{\overline{x}} u(w - r - x + I(x))f(x;a)dx - v(a)$ subject to ICC and zero-profit constraints that are the obvious modifications of (2) and (3). (We ignore, for the moment, any bounds on $I(x)$.) It is straightforward to verify that the first-order condition on $I(x)$ solving the optimal contracting problem yields (13) as in the general model. Using this plus the first-order condition on r yields

$$u'(w-r) = \int u'(w-r-x+I(x))f(x;a)dx \tag{14}$$

The deviation from perfect insurance in any moral hazard problem serves only to generate incentives in the most efficient way. Equation (14) reveals that in this case the Arrow-Borch condition holds across the event of a loss: the marginal utility of wealth conditional upon no accident equals the conditional expectation of marginal utility given an accident. This must be an optimality condition because if the condition were violated there would be an insurance benefit to transferring a dollar between the event of an accident and the event of no accident — and the transfer would involve no cost in terms of incentive distortion because

[11] An alternative theory supporting the optimality of deductibles in insurance contracts is costly state verification (Townsend 1979; Gale and Hellwig 1985). This is a theory that endogenizes the extent of asymmetry in information, rather than taking it as given as in the basic principal-agent approach. Insurers cannot always costlessly observe the loss that an individual has incurred. If the loss (the "state") can be verified only at a cost, then the optimal insurance policy will call for coverage only when the claimed loss exceeds a specific level. In other words, a deductible is optimal when the state can be verified only at a cost. The theory involves essentially a re-interpretation of the Townsend and Gale-Hellwig corporate-finance models in terms of insurance contracts.

p is exogenous. The condition (13) implies that the amount of risk borne by the individual, $x - I(x)$, is non-decreasing in x, so as to give the individual incentive to reduce the stochastic size of the loss by exerting more care. The monotonicity of $I(x)$ plus the equality of the expected marginal utility of wealth in the events of accident and no-accident implies that $I(x) > x$ for low x. If we then add a constraint $I(x) \leq x$ on the contract (insurance payout cannot exceed loss because the individual has the ability to cause a loss), the optimal coverage is $I(x) \leq x$, full coverage, for all losses below a critical value $\widehat{x}$ (Rees and Wambach (2008)). The condition that care affects the size of the loss but not the probability is the opposite of the set of assumptions giving rise to a deductible, and the nature of the optimal contract is the opposite. Low losses are fully covered, instead of not covered at all.

Where effort affects the loss conditional upon an accident but not the probability of an accident, the optimal insurance policy involves full coverage of losses below a critical value.

We do not in reality observe insurance contracts with full coverage of low losses. The implication that should be drawn from the model of loss reduction and moral hazard, however, is not the stark prediction of full coverage of low losses but rather the relative inefficiency of deductibles, the opposite kind of contract, in the presence of moral hazard on the size of damages. Providing zero coverage in the event of a small loss, via deductibles, distorts the incentive to mitigate losses when the losses are random.

3.3 Optimal Insurance with an Unobservable Loss

The conventional approach to the optimal insurance problem with or without moral hazard, on which we have focused to this point, assumes that the care decision, a, on the part of the insured individual cannot be observed, but that the size of the loss, x, is observed perfectly. The theory allows for hidden action, but not hidden information. A natural assumption in many cases, however, is that the size of the loss is both random and not verifiable by the insurer ex post. Consider, for example, insurance offered in many major premium credit cards against a delay of baggage arrival during air travel.[12] The provider of baggage insurance has no information on the cost to the individual of the baggage delay. This cost may be zero or, for the business traveler who has to replace a suit for a meeting, substantial. The only feasible insurance policy when the size of the loss is unverifiable is a fixed payment, I, in the event of a loss.

We consider here the implications of this assumption for both the no-moral-case and the case of ex ante moral hazard.

[12] The standard insurance policy against baggage delay (as of 2023) allows the insured individual to claim up to 500 dollars to purchase clothing if baggage is delayed by more than 6 hours.

No Moral Hazard: If the individual cannot affect the distribution of loss then there is no moral hazard problem, apart from the constraint that an individual would not report the true loss, x. The individual facing a random loss x with exogenous probability p chooses an optimal insurance contract (r^*, I^*) to solve the following problem

$$\max_{r,I} \qquad (1-p)u(w-r) + pE[u(w-r-x+I)] \tag{15}$$

subject to

$$r = pI$$

The first-order condition for this problem is the familiar Arrow-Borch condition that the expected marginal utility in the event of an accident must equal the marginal utility in the event of no accident:

$$E[u'(w-r-x+I)] - u'(w-r) = 0 \tag{16}$$

If the utility function satisfies decreasing absolute risk aversion, $d[-u''(w)/u'(w)]/dw < 0$, then $u''' > 0$, i.e. u' is convex. The first-order condition evaluated at $I = E[x]$ is then positive:

$$E\left[u'\left(w-r-x+E[x]\right)\right] - u'(w-r) > u'(w-r-E[x]+E[x]) - u'(w-r) = 0$$

where the inequality follows from Jensen's and the convexity of u'. Thus the optimal insurance coverage is greater than the expected loss.

> When the loss conditional upon an accident is random and cannot be verified by the insurer ex post, then the optimal insurance contract for an individual whose utility satisfies decreasing absolute risk aversion is an insurance payment in the event of an accident that is greater than the conditional expectation of the loss.

Moral Hazard on Probability of an Accident: Suppose that the probability of an accident equals $p_0 - a$, where a is the individual's effort in avoiding an accident. The individual disutility of effort is, as earlier, a convex function, $v(a)$, satisfying Inada conditions. In this case, the optimal insurance policy solves the following problem:

$$\max_{r,I,a} \qquad [1-(p-a)]u(w-r) + (p-a)E[u(w-r-x+I)] - v(a) \tag{17}$$

subject to

$$a \in \arg\max_{e} [1-(p-e)]u(w-r) + (p-e)E[u(w-r-x+I)] - v(e)$$
$$r = (p-a)I$$

Substituting the individual's first-order condition for the incentive compatibility constraint into (17), and solving for the first-order condition for this maximization problem yields the following equation.[13]

$$E[u'(w-r-x+I)] - u'(w-r) = \lambda \frac{E[u'(w-r-x+I)][1-(p-a)] + u'(w-r)(p-a)}{(p-a)[1-(p-a)]}$$

Comparing this first-order condition with the non-moral hazard optimal insurance condition (16), we see that the expected marginal utility with accident is higher than the one without accident, which implies the coverage I with moral hazard is lower than the coverage when there is no moral hazard. The partial insurance coverage again enhances the incentive for the individual to take care.

Ex Ante Moral Hazard on the Size of the Loss: The distinction between *ex ante* moral hazard on the size of the prospective loss in the event of an accident, and *ex post* moral hazard in the event of an accident is subtle. Ex ante moral hazard on the loss size is illustrated by the baggage delay insurance example. In this example, it is reasonable to assume that the probability of an accident is exogenous. The traveler does not cause the baggage delay. The care that an individual would take to avoid a high cost of delay, however, is the outcome of a decision on the part of the individual. The individual would take into account the need for specific items in baggage and would pack the essential items in their carry-on baggage. Ex post moral hazard, on the other hand, is illustrated by the distortion in incentives for an individual deciding on medical care once an accident or illness has occurred. We consider ex ante moral hazard on losses here and ex post moral hazard in the next subsection.

If the loss on the part of the individual is observable, the inability of the insurer to observe care creates a moral hazard problem, as we showed in the most general, principal-agent model of insurance. The insurance policy covers higher losses with greater insurance payouts in this

[13] With the substitution of the agent's first-order condition for the ICC, the Lagrangian is

$$\begin{aligned}\mathcal{L} = {} & [1-(p-a)]u(w-(p-a)I) + (p-a)E[u(w-(p-a)I-x+I)] - v(a) \\ & + \lambda\{u(w-(p-a)I) - E[u(w-(p-a)I-x+I)] - v'(a)\}\end{aligned}$$

case, but does not cover the full marginal loss (under standard conditions). The co-payment by the insured individual provides an incentive for effort to reduce losses.

Where the size of the loss is not observed, so that the insurer has even less information, the moral hazard problem on losses disappears. The insurance policy is restricted to paying a lump sum in the event of an accident, as we have discussed. This leaves the individual with the full share of loss at the margin, and therefore the full benefit of care. The optimal policy would be identical to (15) with the distribution of losses given by $F(x; a^*)$ where a^* is the first-best level of care. The non-observability of care by the insurer is costly to the individual, of course, in that it constrains the class of insurance contracts that can be written. Ironically, the extra limitation on what the insurer can observe eliminates the incentive distortion.

Suppose that the insured individual's ex ante effort affects the size of the loss conditional upon an accident. Then if the loss is observed by the insurer, the optimal policy includes partial coverage of marginal losses, under standard conditions. If the loss is not verifiable, then the individual's ex ante effort to reduce losses is first-best. The moral hazard problem disappears.

4 Ex Post Moral Hazard

The ex post moral hazard problem arises when an individual's expenditures on reducing the damages from an accident are covered by insurance, and the insurer cannot identify exactly the efficient expenditure ex post. The most important example of ex post moral hazard problem is in medical insurance, which has from the beginning been a focus of the literature on moral hazard (Arrow 1963; Pauly 1968; Zeckhauser 1970; Ma and McGuire 1997; Ma and Riordan 2002). Insurers cannot identify the exact state of health of an individual, and must instead rely on the decision by the individual and her doctor as to the level of care and expenditure.

The conventional view is that the insured individual, capturing the full benefit of marginal expenditure on medical care but bearing less than the full cost, will spend excessively relative to the first-best. We begin by outlining the simplest model supporting this intuition. This model draws on Zeckhauser (1970). Zeckhauser's approach is prescient in recognizing the ex post moral hazard problem as one of hidden information rather than hidden action. We reformulate the Zeckhauser model with a continuum of states rather than a finite number, and make use of the revelation principle in the reformulation. The revelation principle (Myerson 1979) implies in this context that in designing an insurance contract in a model with hidden information, one cannot do better than adopting a direct mechanism (i.e. a mechanism in which the individual reports her type) that is incentive-compatible (i.e., subject to the constraint that individuals have the incentive to report the truth). The revelation principle had not been developed at the time of Zeckhauser's contribution. See Myerson (1979) as well as Dasgupta et al. (1979), and Gibbard (1973).

The moral hazard problem in medicare is more complex than the simple model suggests in at least two respects. Additional medical care at one point in time, especially preventative care, may reduce expenditures to be made later and thus benefiting the insurer. The individual bears only part of the benefits as well as part of the costs of preventative medical care. That is, medical care has some elements of ex post moral hazard and some elements of ex ante moral hazard, in that preventative care lessens ex post damages. Second, as analyzed in Ma and Riordan (2002), managed care systems such as HMO's can mitigate the potentially severe incentive problem in medical care. We offer in section 4.2 a brief outline of the Ma and Riordan model.[14]

4.1 Basic Model of Ex Post Moral Hazard in Medical Care:

Consider an individual with an uncertain need for medical care, i.e., uncertain preferences over medical care and all other commodities. Let x be the individual's expenditure on medical care, y the expenditure on all other commodities, and θ be the uncertain state of the world. State $\theta = 0$ refers to perfect health, and an increase in θ is interpreted as worsening health. The distribution of θ is smooth. The patient's utility function is $u(x, y; \theta)$, which satisfies $u_x \geq 0$, $u_{xx} \leq 0$, $u_y \geq 0$, $u_{yy} \leq 0$, $u_x(x, y, \theta) = 0$ for all x exceeding some finite $\hat{x}(\theta)$ for every θ; and

$$\frac{\partial}{\partial \theta}(\frac{u_x}{u_y}) > 0 \tag{18}$$

The condition (18) states that marginal rate of substitution between health care and expenditure on all other goods is increasing due to increased severity of illness. The fact that u_x reaches 0 at finite x means that there is a limit to the marginal value of health care even at a zero price. The extra month spent in hospital for a stubbed toe carries negative utility. The individual's initial wealth is w.

The individual has the opportunity to purchase insurance prior to the realization of θ. The insurance policy $[r, I(x)]$ has a premium r, and provides coverage $I(x)$ when expenditure on health is x. Ex post, having entered an insurance policy $[r, I(x)]$, the individual chooses health care expenditure x and other expenditure y to maximize $u(x, y; \theta)$ subject to the budget constraint $x + y - I(x) \leq w - r$. This maximization problem forms the incentive compatibility constraint, in the choice of an optimal insurance contract. The optimal contract

[14] An alternative model is offered in Blomqvist (1977).

solves

$$\max_{r,I(\cdot)} \int u\left(x(\theta), y(\theta), \theta\right) f(\theta) d\theta$$

subject to the incentive-compatibility and zero-expected-profit conditions:

$$\forall\theta \qquad (x(\theta), y(\theta)) = \arg\max_{\tilde{x},\tilde{y}} u(\tilde{x}, \tilde{y}; \theta) \qquad \text{subject to} \qquad \tilde{x} + \tilde{y} - I(x) \leq w - r \tag{19}$$

$$r - \int I(x) f(\theta) d\theta = 0$$

The revelation principle allows us to reformulate the problem as the choice of an expenditure plan contingent on health:

$$\max_{x(\theta),y(\theta)} \int u(x(\theta), y(\theta); \theta) f(\theta) d\theta \tag{20}$$

subject to

$$\forall\theta \qquad \theta \in \arg\max_{\hat{\theta}} u(x(\hat{\theta}), y(\hat{\theta}); \theta) \tag{21}$$

$$\int (x(\theta) + y(\theta)) f(\theta) d\theta \leq w \tag{22}$$

Taking a first-order approach, we assume that the incentive compatibility constraint (21) can be replaced by the first order condition of the agent's maximization problem. Under the assumption that the optimal solution $x(\theta), y(\theta)$ varies smoothly with θ, the first-order condition is

$$u_x(x(\theta), y(\theta); \theta) x'(\theta) + u_y(x(\theta), y(\theta); \theta) y'(\theta) = 0$$

That is, the marginal rate of substitution equals the ratio of the rates of change of x and y with θ: $u_x/u_y = -y'(\theta)/x'(\theta)$.

First-best benchmark: When the state of health is observable to the insurance company, the incentive compatibility constraint can be dropped. In this case, the maximization of (20) subject to (22) yields

$$u_x = u_y = \lambda \qquad \forall\theta$$

This condition means that the patient is fully insured under the first-best contract: the marginal utility of income is equal across states and, from the patient's ex post maximization problem, equal to the marginal utility on each class of expenditures.

At this level of generality, the pattern of first-best insurance can be almost anything. Consider, for example, an individual with a preference for only two activities: helicopter skiing and reading library books. The individual should purchase *negative* insurance against the event of a broken ankle that would preclude skiing: this event carries a much lower marginal utility of wealth so that transferring wealth out of the event ex ante raises expected utility. Optimal insurance at a fair premium involves equating the expected *marginal* utility across events, not compensating the individual for lost utility.

Optimal Contract: Adopting a first-order approach yields a Lagrangian for the optimal contracting problem given by

$$\mathcal{L} = \int [u - \mu(\theta)(u_x x' + u_y y') + \lambda(w - x - y)] f(\theta) d\theta$$

Let $G(x, y, x', y') = [u - \mu(\theta)(u_x x' + u_y y') + \lambda(w - x - y)] f(\theta)$. Adopting a calculus-of-variations approach, we have that the optimal expenditure plan satisfies

$$\frac{\partial G}{\partial x} = \frac{d}{d\theta}(\frac{\partial G}{\partial x'})$$
$$\frac{\partial G}{\partial y} = \frac{d}{d\theta}(\frac{\partial G}{\partial y'})$$

The first-order conditions are then

$$(u_x - \lambda + \mu u_{x\theta}) f = u_x(\mu f)'$$
$$(u_y - \lambda + \mu u_{y\theta}) f = u_y(\mu f)'$$

Take the ratio of these first order conditions, we have

$$\frac{u_x - \lambda + \mu u_{x\theta}}{u_y - \lambda + \mu u_{y\theta}} = \frac{u_x}{u_y}$$

which implies

$$\lambda(u_y - u_x) = \mu u_y^2 \frac{\partial}{\partial \theta}(\frac{u_x}{u_y}) \geq 0 \tag{23}$$

if the multiplier is positive.

To generate predictions, we must impose enough structure to set aside unusual preferences such as those of the helicopter skier. We can do this by imposing a structure of additive utility in x and y: $u(x, y, \theta) = v(x, \theta) + b(y)$, with v and b satisfying conditions that yield our assumed restrictions on derivatives of u. (In particular, for every θ there is some $\hat{x}(\theta)$ where $v(x, \theta)$ reaches a maximum.) With the additive utility function (which implies a positive multiplier), we have from (23)

$$\lambda(b_y - v_x) = \mu b_y v_{x\theta} \geq 0$$

which implies that $b_y \geq v_x$

This gives us the first property of the optimal contract: it involves greater than first-best expenditure on y relative to x. This is possible only if the slope of the optimal insurance coverage, $I'(x)$ is positive, because of the incentive compatibility condition (19) in the first formulation of the problem. As in the ex ante moral hazard problem, ex post moral hazard does not eliminate the gains from trade in insurance markets. The optimal contract provides a positive amount of insurance.

The second property, however, is that the optimal insurance policy provides less than full coverage of medical expenditures: $I'(x) < 1$. Just as in the ex ante moral hazard problem, some exposure to the risk of needing medical expenditures is left with the individual. This follows simply from the requirement at the optimum that $v_x(x(\theta);\theta) > 0$. A positive marginal utility of medical care means that at each θ the individual is spending less than $\hat{x}(\theta)$, which is the expenditure under full insurance. From (19) this is possible only if $I'(x) < 1$, showing that insurance, as in Zeckhauser (1970), is incomplete.

In a general setting of insurance with ex post moral hazard on medical expenditures, the optimal insurance policy provides positive insurance but less than full insurance of medical expenditures.

4.2 Ma-Riordan (2002)

The design of optimal insurance contracts in an environment with ex post moral hazard must balance the benefits of greater insurance in terms of superior allocation of risk-bearing, against the incentive distortion that greater insurance unavoidably brings. This is the same trade-off that must be struck in any principal-agent model with a risk-averse agent. Ma and Riordan (2002) provide a particularly elegant formulation of this trade-off in the context of ex post moral hazard. In their model, an individual faces a probability λ of becoming ill, with $0 < \lambda < 1$. The illness varies by its severity l with a density $f(l)$ and distribution $F(l)$. The insured individual learns of the benefits of treatment, i.e. the severity of the disease, after the realization of the illness and severity, but neither the illness nor the severity can be contracted upon by the insurer. The consumer's preferences are presented (in the "utility loss" version of the Ma-Riordan model) by $U(y) - bl$, where y is the income available for expenditure on goods other than medical care. The treatment of the illness is a fixed amount C, and eliminates the disease with certainty.[15]

[15] The full model in Ma and Riordan allows for a pecuniary loss al as well as the utility loss.

A first-best contract, in the Ma-Riordan model, has the consumer paying a fixed premium, P, and receiving treatment whenever the benefits of treatment l is above a particular threshold L. The first-best contract maximizes expected utility subject to a zero-profit constraint:

$$\max_{P,L} \quad (1-\lambda)U(Y-P)+\lambda\left[\int_0^L [U(Y-P)-bl]f(l)dl+[(1-F(L)]U(Y-P)\right]$$

subject to

$$P \geq \lambda[1-F(L)]C$$

The solution involves a threshold L^* that equates the benefits of treatment of a disease of severity L^* with the cost of treatment: $bL^* = U'(Y-P)C$.

When the insurer cannot observe l, the first-best contract cannot be struck. Instead, the contract calls for a co-payment, D, on the part of the patient, and leaves the treatment decision up to the patient. The optimal contract (P, D) maximizes expected utility subject to an incentive compatibility constraint, that the patient chooses treatment when the severity of illness exceeds a threshold L that is chosen rationally given the contract, as well as a non-negative profit constraint:

$$\max_{P,D,L} \quad (1-\lambda)U(Y-P)+\lambda\left[\int_0^L [U(Y-P)-bl]f(l)dl+[(1-F(L)]U(Y-P-D)\right]$$

subject to

$$U(Y-P)-bL = U(Y-P-D)$$
$$P \geq \lambda[1-F(L)](C-D)$$

Ma and Riordan characterize the optimal co-payment D^* as balancing the expected utility cost of a marginally higher co-payment against the corresponding benefits of a lower premium. In the case where the probability of an illness is small (i.e., letting λ approach zero) the trade-off yields a relatively simple expression:

$$\frac{C-D}{D} = [\frac{U'(Y-P-D)}{U'(Y-P)}-1] \;/\; \{\frac{f(L)}{[1-F(L)]}\frac{DU'(Y-D)}{b}\}$$

The term in the numerator of the right hand side is the insurance benefit of a marginally lower co-payment, as given by the "Arrow-Borch distortion" of the given co-payment. This distortion depends on the concavity of the utility function on the domain between $Y-D$ and Y. The incentive distortion of a marginally higher co-payment depends on the consumer's elasticity of expected demand for treatment with respect to the co-payment (at a constant premium). This elasticity is the expression in the denominator of the right hand side. This expression illustrates the trade-off between providing insurance and controlling moral hazard. If the consumer is highly averse to income risk and therefore has a high Arrow-Borch distortion, then the insurance company bears a high fraction of the treatment cost in order to better insure the patient. On the other hand, if the demand for treatment is highly price-elastic, then the potential incentive distortion is high and the consumer bears a substantial co-payment in order to curtail excessive demand. Optimal cost sharing balances these two considerations.

4.3 Ex post moral hazard with managed care

The health care insurance market has responded to the ex post moral hazard problem with supply-side managed care, especially Health Maintenance Organizations (HMO's). An HMO is an organization that provides managed care for health insurance contracts in the United States. The HMO serves as a liaison between insurers and health care providers (hospitals, doctors, etc.). Unlike traditional indemnity insurance, an HMO covers only care rendered by those doctors and other professionals who have agreed to treat patients in accordance with the HMO's guidelines and restrictions in exchange for a steady stream of customers.

The HMO strengthens the role of a health care provider as a gatekeeper for treatment. Even prior to HMO's physicians played this role.[16] Ma and Riordan (2002) and Ellis and McGuire (1990) develop models of treatment decision as a collective decision that maximizes a weighted sum of the benefits of treatment to the physician and the patient, in which physicians are induced by contracts with insurers to be sensitive to costs. In the Ma-Riordan model, a physician's payment involves a fee, S, for a diagnosis but also a bonus, B, if the diagnosed patient does not receive treatment. The collective decision as to treatment between the doctor and the patient provides a threshold L that maximizes the physician's expected payment minus the patient's loss from illness (multiplied by a weight representing patient bargaining power). The constraints on the solution are that the expected costs to the insurance company are covered and the physician must find the relationship profitable. The managed care relationship, when the patient's bargaining power is high, expands the set of feasible treatments relative to the Ma-Riordan model outlined above. The optimal contract

[16] As Arrow (1963, pp. 960) explained

> By certifying to the necessity of given treatment or the lack thereof, the physician acts as a controlling agent on behalf of the insurance companies. Needless to say, it is a far from perfect check; the physicians themselves are not under any control and it may be convenient for them or pleasing to their patients to prescribe more expensive medication,private nurses, more frequent treatments, and other marginal variations of care.

contains both a positive co-payment by the consumer and a positive bonus for not treating on the part of the physician. As Ma and Riordan express it, the optimal arrangement involves both demand-side and supply-side management. In some cases in the Ma-Riordan model, the first-best treatment is possible. The point is not that first-best efficiency is possible in the real world, but rather that supply-side management is vital in contractual responses to ex post moral hazard.

5 Dynamics of Insurance Contracts under Moral Hazard

We have, to this point, analyzed moral hazard within static models. Introducing a dynamic setting leads us to two important questions. First, insurance contracts can in principle be long term. Long term insurance contracts would involve experience-based premiums, i.e. premiums that depend on an individual's accident history. Are long term contracts optimal because they allow for greater incentives? Second, where insurance contracts are short term market dynamics clearly matter as an empirical fact. Experience in past insurance contracts can affect the optimal contracting at a particular date. Automobile insurance policies, for example, often involve discounts for drivers with safe driving records. Does a sequence of short term contracts induce greater effort on the part of an individual in order to attract lower premiums in the future?

5.1 Long-term Insurance Contracts

The economics of long term insurance contracting under moral hazard are complex. The literature shows that optimal insurance coverage in long term contracting under moral hazard cannot be divorced from an individual's decisions on saving or borrowing (e.g, Rogerson 1985b). At a general level, this is not surprising. Consumption smoothing over states of the world through insurance is clearly linked to consumption smoothing over time, through savings and borrowing. If an individual had no access to capital markets at all, and therefore had to consume earnings each period, income smoothing over time would be eliminated. The optimal sequence of short term contracts would be unaffected by dynamics in the sense that finite repetition of the problem would have no impact on the optimal contract. On the other hand, if an insurer has better access to capital markets than individuals, long term contracts play a role of smoothing consumption over states and over time simultaneously. The optimal long term contracting problem then confounds optimal insurance with optimal savings through the insurance company.

In order to focus on the pure insurance motives for contracting, as opposed to savings through insurance, we initially set aside the differential access to capital markets by assuming that the individual and the insurer can borrow or invest at the same interest rate. Even from a pure insurance perspective, however, long term contracts might appear to offer greater contractual flexibility in responding to moral hazard. Experience-based premiums, for example, might appear to be an additional instrument that the contract can rely upon to elicit stronger incentives. If the event of an accident leads not only to losses to the individual because of partial coverage or coinsurance but also to higher future premiums, then the incentives to avoid accidents would seem to be stronger. The flexibility offered by long term contracts in responding to moral hazard with a richer set of contractual parameters would seem to be of value.

This intuition is false. In the simplest setting, the benchmark setting outlined below, long term contracts contribute nothing to the resolution of moral hazard. The ability to raise future premiums in response to accident occurrence adds nothing to incentives in the optimal contract that cannot be achieved with partial insurance. More precisely, any allocation of wealth across states that can be achieved by long term contracts in a competitive market can also be achieved by a sequence of short term contracts. Long term contacts and experienced-based premiums must be explained by other features of insurance markets such as hidden information. Long-term contracts cannot be explained by moral hazard alone.

5.1.1 Benchmark Setting

The benchmark in a moral hazard setting is one in which long term contracts offer no advantage over a sequence of short term insurance contracts. This setting requires that individual savings (or consumption) decisions be observed by insurers. If savings were not observed, then the individual's future wealth would not be directly observed by an insurer — and, as we shall discuss, savings decisions would possibly be random (a mixed strategy). Because risk aversion depends on wealth in general, the result would be that a hidden information problem is generated endogenously in the second period. The exception to this would be preferences satisfying constant absolute risk aversion, since under these preferences changes in wealth do not affect the preferences over insurance contracts (Chiappori et al. 1994; Park 2004). We set aside this problem by assuming that saving is observable. We continue to assume, however, that saving is not contractible — focusing attention on conventional insurance contracts that specify simply premiums and payouts.

We set out the benchmark of the equivalence between a long term contract and a sequence of short term contracts in a two-period model within the simple setting. Consider an individual facing the risk of a loss L (an "accident") in each of two periods. The probability of the loss in either period if no care is taken is p_0, and we can measure the care that the individual takes to avoid the accident in either period as the reduction in this probability. Denoting care in period t by a_t, the probability of an accident in period t is $p_t = p_0 - a_t$. The individual's initial wealth at the beginning of period 1 is w. (To save on notation, the individual receives no income in either period.) We treat care as involving a pecuniary cost, for simplicity, although nothing depends upon this assumption: the cost of care a in any period is function $g(a)$ that is increasing and convex with $g'(0) = 0$ and $g'(a)$ unbounded on the interval $[0, p_0]$. The individual can save or borrow at an interest rate r. Finally, the individual faces a competitive insurance market that is willing to offer any insurance policy (we consider one-period contracts and then two-period contracts) that returns zero expected present value of profit, with second-period profits discounted at the same rate r. That is, in the benchmark case, the individual has access to the same capital market as the insurance market.

Short term contracts in benchmark setting: When the individual has access only to short term insurance contracts, the timing of the game is as follows. In each period t, the individual first chooses an insurance contract $[R_t, I_t]$, which requires an immediate payment of R_t in return for insurance coverage of I_t in the event of an accident. The individual then chooses care a_t; nature chooses between an accident or no accident with probabilities $(p_0 - a_t)$ and $1 - (p_0 - a_t)$, and the individual's net loss in the period is then $-R_t - g(a_t) - L + I_t$. The initial wealth minus the first-period loss can be saved at an interest rate r for consumption in the second period.

In this simple model, the individual's choice over insurance contracts will give rise to a consumption allocation, which is a vector of consumption in the six time-events: $\mathbf{c} = \{c_n^1, c_a^1; c_{nn}^2, c_{an}^2, c_{na}^2, c_{aa}^2\} \in R_+^6$. (The superscript on each element of this vector denotes the time period; the subscript denotes the history of accidents.) The individual's preferences over consumption are represented by the utility $u(c^1) + \delta u(c^2)$, where δ is a discount factor.

A consumption allocation $\mathbf{c}$ is *feasible under short term contracting* if there exists a contracting plan, $\{(R_1, I_1, a_1), (R_a, I_a, a_a), (R_n, I_n, a_n)\}$ (here subscripts a and n denote plans for second period contracts given the history of loss realization in period 1) and a savings plan (s_a, s_n) satisfying four sets of constraints:

(1) The plan implements $\mathbf{c}$ in the sense that if the plan is followed, the resulting consumption allocation is $\mathbf{c}$.

(2) Participation constraints, or non-negative expected profit constraints.

(3) Incentive compatibility constraints on a_n and a_a, and sequential rationality of a_1, s_a and s_n.

(4) Sequential rationality in the choice of the second period contracts in the plan.

Long term contract in benchmark setting: Under a long term contract $\{(R_1, I_1, a_1), (R_a, I_a, a_a), (R_n, I_n, a_n)\}$ all of the contractual payments (premiums and payouts) are committed to at the beginning of the first period. A consumption plan c is feasible under a long term contract if an analogous set of constraints is satisfied — but now the non-negative profit constraint is on the present-value of payments and receipts embedded in the long run contract.

Since the long term contract involves a weaker set of constraints than the short term contract, it follows that any consumption allocation feasible under short term contracts is feasible under a long term contract. The long term contract can simply duplicate a plan of short term insurance contracts. The converse is also true in this setting: it is straightforward to show that the optimal consumption pattern under a long term contract can be implemented with a sequence of short term contracts and a sequentially optimal savings plan.[17] Rey and Salanié (1990) demonstrate this.[18] Intuitively, a long term contract can offer no additional gains in expected discounted utility because insurers cannot offer a transfer of wealth across time at better terms than those available to the individual. The result of irrelevance of multi-period contracts in the simplest insurance context, in which wealth is observed the beginning of the second period, is an example of the more general irrelevance of long term contracts in principal-agent models under particular conditions. (See for example Malcolmson and Spinnewyn 1988; Salanié (2005) reviews these conditions.)

Where long term insurance contracts are feasible in a dynamic setting, under the simplest benchmark set of assumptions, long term contracts offer no advantage over a sequence of short term contracts. Short-term contracting is an equilibrium.

5.1.2 Four departures from the Benchmark

Against this benchmark, we can set out several factors — departures from the assumptions of the benchmark — that give rise to gains from long term contracting. The first departure is

[17] For brevity, we omit the detailed development of the model and the proof of this statement. The intuition is clear: the optimality conditions for a long term contract can be reduced to two sets of conditions: (1) a Borch condition on the optimal smoothing of consumption across states in each time period; and (2) an optimal smoothing condition over time on the realized consumption in period 1 and the conditionally expected consumption in period 2. A sequence of short run insurance contracts meets the first condition. The second condition is met by the individual's optimal saving decision when the individual faces the same interest rate as the insurer.

[18] Rey and Salanié state "When the agent has access to a credit market (which presupposes some commitment on the agent's side) and this access can be contracted upon, a short-term contract can also be described as a 'spot' contract that specifies both consumption and savings or income. Thus in this case 'spot contracting' is sufficient for long-term optimality".

to allow for *a difference between the interest rate available to consumers and the interest rate available to insurers.* If insurers have superior access to capital markets, then a long term contract is of course preferred, since the contract can provide both gains from insurance and gains from the better interest rate. Rogerson (1985b) analyzes the repeated moral hazard problem under the assumption that the principal has access to capital markets, with an objective of the present value of profits, but the agent has no access to capital markets. The optimal contract, in this case, involves spreading the impact of a shock to wealth across future periods. The outcome for the agent thus depends on the outcomes in past periods, a property referred to by Rogerson as the *memory effect.* In our context of insurance, the memory effect would result in higher premiums as a consequence of losses in past periods. That is, the prediction is a long term contract with experience-based premiums. But the role of experience-based premiums is solely to spread the impact of a loss (which must be borne partly by the insured individual in a model with moral hazard) across time. The experience-based premium is not motivated by the enhancement of incentives. Rey and Salanié (1990) show that when renegotiation is allowed and there is no adverse selection at the time of recontracting, an optimal long term contract with full commitment can be implemented via series of two-period contract. Their result encompasses our context of moral hazard if technologies and preferences are time separable.

Moving from theory to reality, the superior access to capital markets by insurers is an important basis for long term contracts in life insurance. Whole life insurance, commonly described as combining insurance and savings, exists in part because of a tax arbitrage: individual savings or investment is taxed whereas the tax rate faced by life insurance companies is very low. Life insurance is recognized as a legitimate tax shelter.[19]

The second departure from our benchmark is to allow for *hidden information at the outset of the multi-period game.* We have, throughout this chapter, assumed away hidden information about individuals' risk types in order to focus on hidden action. That is, we have focused on moral hazard issues rather than adverse selection. A substantial literature, reviewed in the chapter by Dionne, Fombaron, and Doherty in this Handbook, analyzes the contractual response to hidden information on risk types. Dynamics are a key part of hidden information insurance models. An individual's choice of insurance in a sequence of short term contracts, for example, will be influenced by the inference that future insurers draw from the contract choice about the individual's risk type. Commitments made in long term contracts can mitigate the resulting distortions.

[19]Note that the tax-shelter aspects of life insurance are somewhat constrained, at least in the U.S., to prevent the avoidance of inheritance tax. In flexible-premium policies, large deposits of premium could cause the contract to be considered a "modified endowment contract" by the Internal Revenue Service (IRS), which would involve a tax liability, negating many of the tax advantages associated with life insurance.

The third departure from the benchmark model again involves the interaction of savings and insurance. We assumed in our benchmark model that under the strategy of short term contracting the insurers had full information in every period on the agent's wealth. Only the agent's effort in each period was not observable. If, however, *the individual's savings decision is not observable by the insurance company* then we have to consider two possibilities: (1) the agent's savings under the optimal long term contract is non-random (i.e. a pure strategy) or (2) the agent's savings and therefore second-period wealth is random, even conditioning upon the first period accident outcome. Chiappori et al. (1994) offer a surprising result: *when renegotiation is allowed, then any optimal contract associated with non-random savings decisions cannot implement any effort level above the minimum effort level.*[20] Any higher effort level in the second period requires that a second period incentive compatibility constraint on a_i^2 (in the notation of our model above) be binding. This means that the agent's utility from the second period effort level implemented by the contract must be the same as the utility level that the agent could achieve under the contract with another effort level a'. But with savings unobservable, an ex ante incentive compatibility constraint imposes a constraint on the agent's savings decision, s_i, at the end of the first period. A savings decision optimal conditional upon effort level a' planned for the second period would increase utility for the agent beyond what is determined by the contract. In other words, the ex post incentive compatibility constraint on a_i^2 is incompatible with the ex ante incentive compatibility constraint on s_i. This result requires only that optimal savings depend upon the effort planned for the second period. If the costs of effort are pecuniary (as in our model) and the agent's utility exhibits constant absolute risk aversion this condition will fail. Apart from the case of CARA utility, in short, the optimal contract will involve random savings when savings cannot be observed. Hidden information arises endogenously in the second period choice of insurance coverage.

Our final departure from the benchmark assumptions is to allow for the possibility that *effort persists over time.* Examples of persistent effort include maintenance (or other safety-investment decisions), with effectiveness that depends on the consistency of repairs. Fernandes and Phelan (2000) show that the optimal contract in this context can be quite complex using recursive methods. In a binary effort model, Mukoyama and Sahin (2005) and Kwon (2006) show that persistence decreases the need for loss sharing to incentivize the agent. In particular, when persistence is high, it is sufficient to incentivize the agent only in the last period. In earlier periods, the prospect of a large future reward induces the agent to exert

[20]When renegotiation is ruled out (and, therefore, renegotiation proofness not imposed), DeMarzo and Sannikov (2006) and He (2012) show that the optimal long-term contract can induce effort greater than the minimum effort. Chiappori et al (1994) argue, however, that the optimal long-term contract is not renegotiation proof.

effort. In the absence of loss-sharing benefits, persistence can lead to full insurance, as the insurer is better equipped to deal with risks than the risk-averse agent. In contrast, Jarque (2010) shows that the persistence problem can be reduced to a related repeated moral hazard model without persistence when the agent is risk averse with linear disutility of effort, and output depends on the sum of depreciated past efforts. She shows that under these assumptions a simple solution exists and that full insurance is never optimal.

> Gains to long-term insurance contracting follow from each of four departures from the simple benchmark set of assumptions: differential access to capital markets between the insured individual and the insurer; hidden information at the outset of contracting; non-observability by the insurer of the individual's saving decision; and persistence over time of the impact of effort in reducing random accident losses.

We note, in closing our discussion of long term insurance contracts, that these contracts are very rare in reality. Insurers are almost always free to change their premiums on an annual basis, if not more frequently. The various theories dealing with the gains from long-term contracting do not confront this reality. One disadvantage of long term contracting is that in a market where other insurers are offering short term contracts, and where individuals have some private information on their own risks, a long term contract is likely to attract high-risk individuals because these individuals face a greater risk of future premium increases if they insure through a series of short-term contracts. The explanation of the optimal length of insurance contracts remains an open question.

5.2 Dynamics with Short-Term Contracts

Suppose only short term contracts are available. This is almost always the case in reality: firms are free to adjust their premiums annually. We take the constraint of short term contracts as exogenous rather than explaining it, and ask whether the *repetition* of the market for short-term insurance contracts mitigates moral hazard. One might expect that the threat of increases in future insurance premiums in response to accident claims. But under hidden action alone, with known types, this conjecture is false. In any model with a finite number of periods, the equilibrium contract in each period is identical to the equilibrium contract with only one period.

We suggest through an initial, stylized model that a useful approach to this question is based on the application of Holmström's classic career-concerns model (Holmström 1999). With hidden action alone — the basis for moral hazard in static insurance models — repeated markets do nothing. Ironically, however, greater asymmetry in information — the addition of an assumption of hidden characteristics in addition to hidden information — leads to an explanation of how future insurance markets can enhance incentives. The intuition is that exerting additional effort today leads to a lower risk of an accident, which then increases the likelihood that insurance markets tomorrow will identify the individual (stochastically) as a better type. This lowers expected future premiums.

The following is a very stylized model (essentially an example) of the incentive effects of future premiums on current care. We adopt a two-period model in which an individual faces the risk of the loss of 1 dollar in each period. We assume in this simple model that full insurance is mandated as a regulatory requirement rather than as a consequence of risk aversion. The individual's welfare is given by an initial wealth, w, minus the sum of premiums over the two periods. Individuals are of type, θ, which has a smooth distribution, F with density f. The individual's type is unknown both to the individual and a competitive insurance market, which provides insurance at an actuarially fair premium in each period. The individual can decrease the probability of an accident in period t by exerting effort a_t, which comes at a cost $c(a_t)$. The cost function c is unbounded, convex and satisfies $c'(0) = 0$. A loss is assumed to occur if $\theta + a_t < k$ and the probability of a loss in period t is $Pr[\theta + a_t < k] = F(k - a_t)$.

The individual exerts $a_2 = 0$ in the second period, since effort has no impact on the second period premium; we denote first-period effort simply as a. The competitive insurance market cannot observe a, but has an expectation a^* as to this effort level. The market knows the distribution of types, and in the second period uses a^* and the accident history from the first period (loss or not) to formulate a posterior distribution over θ. The second period premium is equal to the probability of a loss based on this posterior distribution. We denote the first period premium by ρ and the second period premium, given the history of no loss or loss, respectively as (ρ_{20}, ρ_{21}). An equilibrium is a vector $(a^*, \rho_1^*, \rho_{20}^*, \rho_{21}^*)$ satisfying two conditions: the premiums are actuarially fair given the expectation a^* and Bayesian updating in period 2 given the accident loss history; and the agent's optimal effort, taking the premiums $(\rho_1^*, \rho_{20}^*, \rho_{21}^*)$ as given, is $a = a^*$. In other words, the market's expectation a^* leads to a set of premiums that induce an optimal effort equal to a^*. The expectations are self-realizing, or rational.

The following expressions are easily demonstrated. The posterior densities are

$$\Pr[\theta | a^*, 0] = \frac{f(\theta)}{1 - F(k - a^*)}$$

for $\theta \geq k - a^*$and zero otherwise; and

$$\Pr[\theta | a^*, 1] = \frac{f(\theta)}{F(k - a^*)}$$

for $\theta < k - a^*$ and zero otherwise. The premiums satisfy

$$\rho_1 = \Pr[\theta < k - a^*] = F(k - a^*)$$

$$\begin{aligned}\rho_{20} &= \Pr[\theta < k|0] \\ &= \int_{k-a^*}^{k} f(\theta|0) = \int_{k-a^*}^{k} \frac{f(\theta)}{1 - F(k - a^*)} \\ &= \frac{F(k) - F(k - a^*)}{1 - F(k - a^*)}\end{aligned}$$

$$\begin{aligned}\rho_{21} &= Pr[\theta < k|1] \\ &= \int_{-\infty}^{k-a^*} f(\theta|1) = \int_{-\infty}^{k-a^*} \frac{f(\theta)}{F(k - a^*)} \\ &= \frac{F(k - a^*)}{F(k - a^*)} = 1\end{aligned}$$

The history of an accident with a non-negative effort in the first period ensures an accident in the second period with an effort equal to zero; this is why $\rho_{21} = 1$.[21]

The agent chooses a to maximize $w - \rho_1 - c(a) - [(1 - F(k - a))\rho_{20} + f(k - a)\rho_{21}]$.[22] This yields the following first-order condition, when simplified:

$$c'(a) = f(k - a)[\rho_{21} - \rho_{20}]$$

The marginal cost of effort is equal to the marginal benefit of a decrease in the future premium. Substituting in the values for the premiums in terms of a^*, we have an equilibrium when the following condition is satisfied at $a = a^*$:

$$\frac{c'(a)}{f(k - a)} = \frac{1 - F(k)}{1 - F(k - a^*)} \tag{24}$$

This equation defines a mapping from a^* to a, a fixed point of which is an equilibrium. The left-hand side of (24) equals 0 at $a = 0$ and is unbounded. The right-hand side is positive and bounded. A fixed point therefore exists by the intermediate value theorem. All of the equilibria have $a > 0$, demonstrating the point that the threat of future premium increases

[21] This is a consequence of our assumption that the event of an accident depends only on a and θ. If we added a random ϵ (as in Holmström 1999), assuming that an accident occurs when $\theta + a_t + \epsilon < k$, then the model would be "smoother" and the ρ_{21} would be non-trivial. We make our assumption to keep the algebra simple.

[22] Note that this objective involves no expectation, because under the exogenous legal requirement of full insurance the individual faces no uncertainty. Of course without this requirement, we would have to formulate the objective as maximizing expected utility.

incentivizes the agent to take effort to avoid an accident.[23]

In a canonical dynamic model in which only short term contracts are feasible, then: with hidden action but observable characteristics, the repetition of insurance markets has no impact on insurance contracting or incentives. Accidents have no impact on future premiums. With both hidden action and hidden characteristics, the repetition of markets changes the equilibrium in a way that reduces moral hazard. The prospect of future premium increases in response to accidents enhances incentives to exert effort.

6 Insurance in the presence of limited liability

The final topic we address within the models of moral hazard is the interaction between limited liability protection and the insurance decision. Suppose, for example, a U.S. individual has a very small net worth, small prospects of an increase in future wealth, and (like all of us) faces the risk of a major disease. The individual has the choice of buying catastrophe medical insurance, or forgoing the insurance and simply planning in the event of a major disease declaring bankruptcy and having the medical costs covered by Medicaid. Let's set aside for simplicity the lower quality of health care for those individuals under Medicaid. Which option will a rational individual choose? Clearly, if wealth is low enough, then relying on limited liability rather than paying a premium for insurance is the rational choice.

Limited liability means a difference between *optimal* insurance contracts for an individual facing a competitive insurance market, and socially *efficient* insurance insurance contracts. This is because if an accident (loss) occurs and is not covered by insurance, the loss will be incurred by another party or parties. In the case of a medical risk discussed above, taxpayers cover the uninsured expenses. In the case of an individual or corporation financed in part by debt, the debt holders bear the costs of bankruptcy (and if the corporation could commit to full insurance, the cost of raising debt would be lower). In the case of an individual assigned legal responsibility for accident costs incurred by another party, whom we call the victim, it is the victim who bears the uninsured costs.[24]

The simplest model of optimal insurance contract in the presence of limited liability is the following. An individual with utility u and net worth w, faces a known loss x with probability p. For convenience, let $u(0) = 0$. We assume, to start with, that there is no moral hazard. If the individual does not purchase insurance, the payoffs in the event of an accident and not

[23]We have adopted, along with other strong assumptions, the restriction to two periods. Following Holmström (1999), we conjecture that with multiple periods, the effort induced may even be greater than the first-best effort as the agent is rewarded for greater effort by an improved posterior in all future periods.

[24]We analyze this case in more depth, in the next section of this paper on liability insurance.

are $\max\{0, w-x\}$ and w. The individual's final wealth position cannot be negative because of limited liability.

If the individual purchases insurance in the amount I, then the individual's payoff in the event of an accident is $\max\{0, w-x+I-pI\}$ where pI represents the fair insurance premium that is subtracted from wealth. In the event of no accident, then the payoff is $w-pI$. It is clear a priori that if an individual buys insurance, it will be in an amount in the range $I \in [(x-w)/(1-p), x]$. (A lower amount of insurance would leave the individual's wealth 0 in the event of an accident, making the expenditure on insurance worthless.) The decision to buy insurance at all is the choice between accepting the following maximization problem:

$$\max_I \quad p \cdot u(w-x+I-pI) + (1-p) \cdot u(w-pI) \qquad \text{[insurance]}$$

or the expected utility with no insurance:

$$p \cdot \max\{u(w-x), 0\} + (1-p) \cdot u(w) \qquad \text{[no insurance]}$$

The following characterizes the optimal insurance decision for various values of the loss, and is easily proven: Ceteris paribus, the optimal insurance is full insurance ($I = x$) for a loss in the range $x \leq \overline{x}$ and no insurance for $x > \overline{x}$ for some $\overline{x}$.[25]

The impact of limited-liability protection in this simple model: it replaces full insurance completely if the potential loss is large enough.

Next consider the case where x is random and observable ex post. In this case, an insurance policy is a function $I(x)$ with a fair premium $P = pE[I(x)]$. The optimal insurance policy is the solution to

$$\max_{I(x)} \quad pE\left[u\left(w-x-pE\left[I(x)\right]+I(x)\right)\right] + (1-p)u\left(w-pE\left[I(x)\right]\right)$$

Note that $I(x) = 0$ for a range of x is of course possible. Solving the first-order condition point-wise in the choice of I at each x shows that similar to the previous case, we have full insurance, $I(x) = x$, up to some value $\overline{x}$ and $I(x) = 0$ beyond that. (Formally, this is a solution almost everywhere.) Committing to zero insurance for high losses saves more on premiums than insurance protection is worth, given limited liability protection, i.e. the availability of declaring bankruptcy as an option.

[25] Assume that for a given value of x insurance is purchased. The first order condition for the insurance problem is then $-p(1-p)u'(w-x+I-pI) + p(1-p)u'(w-pI) = 0$ which is easily shown to be satisfied only when $I = x$. Next, the difference between the expected utility under full insurance and the expected utility with no insurance is easily shown to be decreasing in x, negative for sufficiently large x and positive for x positive but sufficiently small.

In sum, limited liability means that optimal insurance is less than full insurance. This choice is inefficient for two reasons. First, as is familiar from corporate finance, in the case where bankruptcy involves a transfer of assets to creditors (as it usually does) the the individual's assets, worth w to the individual, are likely to be worth less to the creditor assigned the assets. Physical assets are almost always complementary with the human capital of the pre-bankruptcy owners and of course the human capital cannot be transferred. An inefficiency is thus incurred ex post whenever limited liability is invoked, which it will be when a large loss occurs and liability insurance is zero. Second, the lack of complete insurance puts the "victim", or debt holder or other claimant to the individual's assets at risk. Since that party is almost always more risk averse than the insurer, the choice of incomplete insurance leads to inefficient risk-bearing. These inefficiencies arise from the failure of the insured to take into account the externalities imposed on those who are affected by the insurance decision. In the case of a corporation, the distortion is well-known as the debt-overhang problem (Myers 1977). Decisions involving medical insurance ignore the impact on taxpayers if bankruptcy involves the adoption of Medicare. Thus moral hazard arises even in the purchase of insurance, before considerations of care or effort to avoid an accident.

The introduction of care decisions, the traditional focus of moral hazard, brings a new dimension of efficiency considerations. With limited liability alone, it is clear that the individual has inadequate incentives to undertake effort to avoid an accident. The individual does not take into account the welfare of victims or other individuals affected by bankruptcy. The addition of insurance may make matters worse if the insurer cannot observe the care decisions — and in fact this problem may push the individual into a region where no utility-improving insurance coverage is possible. (As we have seen, this is a possibility even disregarding any inefficiency in care.)

It is possible, however, that insurance *adds* to the efficiency of care decisions. This is in contrast to the standard moral hazard problem in which insurance is the source of the distortion in care. Two main possibilities are behind the possible efficiency of care decisions. First, the insurer may be able to observe care decisions, at least to some extent. Insurance companies, after all, are the experts in managing their risks. In this case, minimum care levels would be part of the insurance contract instead of being decided upon the individual. Second, as we showed in the previous section of this chapter, incentives for care are generated by the repetition of insurance markets and the threat of future premium increases. The incentives provided by the repetition of markets may enhance incentives beyond the level without insurance. As we argue in the next section of this paper, *mandatory* insurance may therefore be a remedy to market inefficiencies arising from the limited liability distortions.

A third possibility, in theory at least, is that long-term insurance contracts may provide improved incentives relative to the incentives distorted by limited liability. Biais et al. (2010) study this possibility. They show that it is possible to incentivize the agent in a long-term contract with premiums that vary over time depending on the history of accidents. The optimal contract may even involve rebates. The authors study a principal-agent model in which a risk-neutral firm protected with limited liability must exert (unobservable) effort to reduce the likelihood of large but infrequent losses. The size of the project undertaken by the agent is dynamic and can be affected by moral hazard and investment. If a long enough time passes without losses, the contract prescribes increasing the project's size and the possibility of the agent receiving payments for good performance. If an accident occurs, payments are suspended, and investment stops. Downsizing occurs in this model when bad performance accumulates.

7 Liability Insurance, Incentives and the Economics of Tort Law

Insurance against legal liability is critical for any homeowner, car owner or business, especially in the U.S. Liability insurance involves a host of issues beyond standard first-party insurance, issues not covered in our review of moral hazard and insurance contracts to this point.

Liability insurance is similar to first-party insurance under limited liability in that instead of just two parties as in the standard insurance model, the insured and the insurer, liability insurance involves three parties. These are: the insured individual is the injurer (or "tort-feasor") who faces the risk of legal liability for the costs of an accident; the insurer; and the victim who incurs damages from the accident and is potentially compensated by the injurer under the law. This means that incentives may be distorted not just because of a positive externality from the insured's care decision on the insurer (as in the standard moral hazard problem) but also because of externalities imposed on the victim or potential victim. And in the case of a standard "joint-and-several liability rule, in which damages not covered by a tortfeasor accrue as liabilities to other tortfeasors, externalities in a tortfeasor's care decision are imposed on other tortfeasors as well. And not only the insured's care decision but also the basic decision to purchase liability insurance in the first place may be distorted.

The set of economic questions that arise in the theory of liability risk and insurance is large, and includes the following:

(1) **The insurance purchase:** How is the optimal liability insurance decision affected by limited liability? Are there conditions, for example, under which the injurer may rely entirely on limited liability for protection against risk, forgoing liability insurance

altogether?

(2) **The capitalization decision:** limited liability protection is stronger the smaller the size of assets relative to the size of the potential accident cost. Has the insured the incentive to under-capitalize in order to strengthen limited liability protection?

(3) **The efficiency of liability insurance:** Liability insurance was at one time attacked as against the public interest on the grounds that it removed the deterrent against negligence.[26] Is liability always efficient, or can the prohibition of liability insurance be justified in theory?

(4) **Mandatory liability insurance:** At the other extreme, are there circumstances under which mandatory liability insurance is an appropriate public policy, because individual incentives to purchase adequate liability insurance are distorted?

(5) **The care or effort decision:** Is the injurer's incentive to take effort to avoid an accident distorted by limited liability, and is this distortion exacerbated or lessened by liability insurance?

(6) **Efficiency of liability rules, with and without liability insurance:** Does the availability of liability insurance affect the socially optimal liability rule, specifying the conditions under which an individual will be found liable for the costs to others of an accident.

We offer a brief sketch of several of these issues; a single chapter on moral hazard and insurance cannot provide an analysis of each question in depth. As background, it is helpful to outline the legal rules that dictate when an injurer must compensate the victim of an accident. A rule of *negligence* means that the producer of a product or an individual engaged in a risky activity can be successfully sued for damages by a victim if the producer was negligent, i.e. engaged in inadequate care, in causing the accident. Even optimal care in general results in some accidents, and a negligence lawsuit would fail if the plaintiff could not prove negligence on the part of the defendant. Under a rule of *strict liability*, on the other hand, a producer is strictly liable for damages caused by its product (or a defendant is strictly liable for damages caused by its activity) whether there is negligence or not. Under negligence, the focus is on the defendant's conduct; under strict liability, the focus is on whether the defendant's product or activity caused the plaintiff's damages. Both rules are used. In product liability, for example, negligence was in force until the early 1960's, whereas strict liability has been the rule since then. In terms of economic theory, both rules have

[26]In 1941, McNeely wrote "it is not surprising that the most recent branch of the insurance family, liability insurance, should ... have been attacked as unlawful and anti-social. The principal ground of the attack has been that the protection afforded the insured by the policy removes the financial deterrent against negligent and criminal acts" McNeely (1941), p.26.

advantages. The efficiency of rules can assessed according to whether they elicit the efficient level of *care* in avoiding an accident as the efficient level of *activity* that gives rise to accident risk (a distinction due to Shavell (1980). A car driver, for example, may exhibit efficient care in avoiding accidents, but engage in too much activity (driving too many miles or too often when roads are icy and dangerous).

Perhaps the most fundamental principle of the economics of tort law is that in the case of a single injurer, the strict liability rule leads (under ideal conditions) to the *internalization* of the externalities that the defendant's behavior imposes on victims. If damages are pecuniary, then under strict liability the damages are simply transferred from the victim to the defendant. The defendant then bears all costs and benefits from its decision on care. The defendant's decisions on both care and activity level are then optimal. But there is a catch. If causation involves more than one defendant, strict liability cannot internalize the accident costs fully to any one defendant. The dilemma is that there may be many actors involved in causing accident risk but only one set of accident damages to assign (this is familiar in mechanism design as the "budget balancing" problem). Each defendant, bearing only a share of accident damages, will not internalize the full costs of decisions taken, with the result of inadequate care and excessive activity. A negligence rule can take care of the efficiency of care decisions even with multiple defendants because each defendant is held to standards of adequate care, but — unlike strict liability — cannot elicit efficient activity levels even when there is only one defendant. Meeting the negligence standard does nothing to ensure efficient activity. Moving to considerations of risk, strict liability becomes more advantageous when injurers are less risk averse (or can obtain liability insurance) since strict liability imposes risks on injurers. The negligence rule becomes more attractive when injurers are more risk averse than victims since the negligence rule imposes risk on victims (Shavell 2007).[27]

The impact of limited liability protection on the optimal purchase of liability insurance is identical to that covered in section 6 of this paper. That is, the impact of limited liability constraint on insurance purchase is the same for liability and first-party insurance. The limited liability constraint can be more prominent in the case of liability insurance, however, as some firms engage in activity that clearly involves the risk of large liability risks. (examples are firms owning sites with asbestos contamination and firms involved in hazardous waste and exposed to other environmental risks). Limited liability results in a lower optimal purchase of liability insurance because the protection it offers is a substitute for the protection offered by liability insurance.

The capitalization decision in the area of liability risks is critical because it means that

[27]This one paragraph summary of tort rules is of course just a glimpse into the economics of the law of torts. An excellent overview of the area is provided by Shavell (1987).

the protection offered by limited liability is strong for a firm with smaller equity relative to potential risks. But the other side of the coin is that bankruptcy can ensure that victims get very little or nothing in the event of an accident. In *Walkovsky v Carlton* 18 NY2d 414, taxicabs in a taxi company were individually incorporated and intentionally undercapitalized. This would not be a problem if each potential taxi customer decided whether to take a cab based on its capitalization, but of course that is far from reality and in any case would not take care of accident-risk externalities for pedestrians or other drivers. The solution in situations like this can be mandatory liability insurance. Insurance guarantees that victims are compensated. And insurance generates efficiency in care decisions and (we think more importantly) in activity decisions such as whether to drive with icy roads. (The care versus activity distinction can be important not just for tort law but also for insurance incentives.) The incentives are generated by the sensitivity of future premiums to the accident history, a link that we analyzed in section 5.2.

Mandatory liability works as a regulatory remedy in situations where the liability insurance is for car drivers. But for large risks such as environmental risks, the dilemma is that liability insurance markets are simply unreliable. Liability insurance in these cases is subject to periodic "tight" markets in which the insurance is very expensive or even non-available (Winter 1988; 1991). The tort law for some accidents, such as those involving asbestos, protects the compensation to potential victims by assigning liability to a large set of firms in the supply chain for the risky product (including even the insurance companies). In theory, as we have argued, liability insurance can reduce the care incentives for the individual; it is possible that full liability insurance can harm care incentives to the extent that overall efficiency falls relative to that flowing from limited liability alone. But empirically this may not be an important effect.

The final question we have posed is the impact of liability insurance on the optimal liability rule? Since the negligence rule leaves the risk of accident with the victim and strict liability leaves the risk with the injurer, we could say that the relative availability of insurance (first-party insurance for the victim; liability insurance for the injurer) influences the optimal liability rule. First-party insurance is, we suggest, the more reliably available. Accident risks facing an individual come from a wide range of sources and are predictable with actuarial data. The risk product liability for a single producer is not as predictable and the supply of liability insurance as a result less reliable. This favors a negligence rule.

8 Conclusion

This chapter outlines the theory of optimal insurance contracts under the condition that the individual's effort to avoid accidents cannot be contracted (ex ante moral hazard) as well as the case when an individual's expenditures on insured items such as medical care cannot be contracted (ex post moral hazard). At a general level, the design of an optimal insurance contract is an exercise in solving the principal-agent problem. But the context of insurance provides enough structure to allow predictions on the form of optimal insurance contracts. The main predictions of the ex ante model are easily summarized. Moral hazard reduces the insurance coverage that an optimal contract offers, although moral hazard does not eliminate the gains from insurance altogether. Some coverage remains optimal. The reduction in coverage can take the form of a deductible or coinsurance, and under standard assumptions the amount of losses left uninsured is non-decreasing in the size of the loss.

The same general principles extend to the characterization of optimal coverage under ex post moral hazard. The greater the level of risk aversion, the greater the insurance coverage of marginal expenditures in the ex post moral hazard model. And the higher the elasticity of demand for items covered by insurance, the lower the extent of insurance coverage.

The introduction of dynamics to the moral hazard model highlights the interplay of the spreading of risks across states via insurance with the spreading of income over time with capital markets. Where individual savings are observable (and therefore wealth levels are observable at the time that insurance contracts are struck) — and if insurers have no better access to capital markets (a higher interest rate) than do consumers — long term contracts offer no gains compared to the adoption of short term insurance contracts. Moral hazard alone is not associated with the efficiency of long term contracts with experience-based premiums. Long term contracts are more efficient than a sequence of short term insurance contracts if the insurer has superior access to credit. Such long term contracts allow the spreading of losses over time, not just across states. Hidden information, set aside in this chapter, is another basis for long term contracts, as well as for experience-based premiums. Contracts in which savings decisions are not observable generate endogenous hidden information, even when the characteristics of individuals are common knowledge at the outset of the contract: individual savings decisions are a mixed strategy in equilibrium with the consequence of uncertain preferences in later periods. Constant absolute risk aversion utility is the exception, since wealth does not affect preferences over risk for these preferences.

The single most important question in insurance dynamics, we suggest, is whether the repetition of insurance markets can provide incentives. Insurance premiums are often based on accident history in reality. We offer a model, with a mixture of hidden characteristics and hidden information, that can capture this mechanism. We offer an outline of the economic issues presented by limited liability and potential losses that are large relative to the net worth of the insured individual. These include incentives not just for care, the focus of mainstream moral hazard analysis, but also for the purchase of insurance itself. The interaction of limited liability protection and insurance protection is complex. In some circumstances, the purchase of insurance may enhance incentives for accident avoidance. Finally, we offered an outline of the economics of liability insurance, which presents a rich set of issues because of its interaction with both the tort system and limited liability protection.

A number of important extensions to the theory of moral hazard have not been covered here.

Multi-dimensional Effort: Investment in care to avoid accidents is in reality multi-dimensional. An individual insuring household belongings against theft can take care in locking the doors and windows, in leaving lights on when she is away, in buying deadbolt locks for the doors, in purchasing an alarm system, and so on. Some dimensions of care can be observed and contracted for more easily than others. The implications for optimal insurance contracts of this fact are worth exploring in depth, as an application of the Holmström-Milgrom (1991) model of multi-task agency.

Non-exclusive insurance: Arnott and Stiglitz (1988a; 1988b), Bisin and Guaitoli (2004) and Attar and Chassagnon (2009) have examined moral hazard under the assumption that the price of insurance is uniform in the amount of insurance obtained from a competitive market. That is, each supplier of insurance has no control over the amount of coverage an individual purchases from other insurers. Since in reality insurers can and do contractually restrict payments when coverage is obtained through other policies, this is better interpreted as fraud rather than conventional moral hazard. Optimal insurance under a non-exclusivity restraint and insurance fraud more generally are important areas beyond the treatment of moral hazard that we have offered.

Moral Hazard on the Supply Side of Insurance Markets: Moral hazard as we have noted is a problem of contracts in general, not just on the demand side of insurance contracts. Indeed, as the events of the financial crisis of 2008 revealed, the more important moral hazard problem in the insurance market is often on the *supply* side of insurance contracts. An insurance contract is a financial contract under which an insurer accepts the financial liability of future insurance payments in exchange for a premium payment today. When the insurer has limited liability, the insurer may have the incentive to invest excessively in risky activities because the insurer does not bear the full downside risk of those activities: in the event of bankruptcy, insured individuals do not receive payment specified in the insurance contracts, or if they do receive coverage, it may be from a government insurance guaranty fund. In either case, distortionary incentives for the insurer to enter risky activities result from the downside protection of limited liability. This is exactly parallel to the risk-shifting problem in the presence of debt, recognized in finance since Jensen and Meckling (1976) and Myers (1977). Insurance regulation has long recognized the incentive distortion in asset allocation decisions on the part of insurers, in constraints of the amounts that insurers can invest in risky assets. The actions of AIG, the largest insurer in the U.S., specifically the issuance of credit default swaps were central to the financial crisis. The decision of AIG to issue large amounts of these risky liabilities shows that the incentive distortion caused by supplier moral hazard is as important in decisions related to the liability side of the balance sheet as in decisions related to the asset side.

Optimal Corporate Insurance: The purchase of insurance by corporations has long been regarded as a puzzle (e.g. Main 1982). Why would corporations, which can easily diversify risks, be risk averse? The standard model of corporate finance predicts that non-systematic risk, risks unrelated to movements in the overall stock market, should be irrelevant to a corporation. Yet corporations purchase insurance against firm-specific losses such as fires or liability for accident damages. We suggest that corporation insurance has been regarded as a puzzle because scholars have mistakenly (but understandably) assumed that the incentive for corporations to purchase insurance reveals risk aversion. Risk aversion is of course the driving assumption in standard models of insurance. Corporate insurance may instead be motivated by specialization in expertise in managing risks of various sorts: assessing and optimal mitigation of risks. Insurance companies invest in expertise about both the assessment of fire insurance and the various means of reducing that risk, for example. That is their business. Corporations invest in expertise in running their own businesses and investment by each corporation in expertise of managing the risks of fire would involve duplication of investment in risk management across corporations. Corporate insurance then represents the assignment of particular risks to enterprises that specialize in the management of those risks.

Appendix

This appendix establishes uniqueness of a solution in a and r to the two constraints on the optimal insurance contract in the ex ante model, the zero profit condition (3), and the agent's first-order condition:

$$\int I(x)f(x,a[I(x)])dx - r[I(x)] = 0 \tag{25}$$

$$\int u(w - r[I(x)] - x + I(x))f_a(x;a[I(x)])dx = v'(a[I(x)]) \tag{26}$$

Proposition. The solutions in a and r, to (25) and (26), $a[I(x)]$ and $r[I(x)]$, are unique within a neighborhood of the optimal $I^*(x)$.

Proof. We proceed by showing that the solutions are unique at the optimal $I^*(x)$; the proof for a neighborhood about the optimal $I^*(x)$ follows directly. Suppose $I^*(x)$ is the optimal coverage policy. Given $I^*(x)$, rearrange the two equations, and define $M(r,a)$ and $N(r,a)$ as

$$M(r,a) = \int u(w - r - x + I^*(x))f_a(x;a)dx - v'(a)$$
$$N(r,a) = r - \int I^*(x)f(x,a)dx$$

Each equation, $M(r,a) = 0$ and $N(r,a) = 0$, defines a curve in (r,a) space. Taking total derivative for both $M(r,a)$ and $N(r,a)$, we can get the slope of both curves in (r,a) space.

$$r'_M(a) = \frac{\int uf_{aa}dx - v_{aa}}{\int u'f_a dx}$$
$$r'_N(a) = \int I^*(x)f_a dx$$

Note that r'_N is the marginal change of coverage payment. This term is non-positive for optimal contract; otherwise, holding other elements of the contract constant, the insurance company could ask the agent to reduce effort by a small amount, which would reduce the expected coverage paid by insurance company. This change would have only a second order effect on agent's utility since effort is chosen optimally, but a first order positive effect on the insurance company's profit. This contradicts the optimality of $I^*(x)$, showing that r'_M is non-positive.

The numerator of r'_M, $\int uf_{aa}dx - v_{aa}$, is the second order condition for the agent's maximization problem, and hence is negative. The denominator is negative for optimal contract. We have shown $x - I(x)$ is a non-decreasing function of x. Therefore, $u'(w - r - x - I(x))$

is a non-decreasing function of x, too. Since $\int f_a dx = \int \frac{f_a}{f} f dx = E[\frac{f_a}{f}] = 0$, and $\frac{f_a}{f}$ is non-increasing by the MLRP. Therefore,

$$\int u'(w - r - d(x)) f_a dx = E\left[u' \frac{f_a}{f}\right] = cov\left(u', \frac{f_a}{f}\right) < 0$$

Therefore, $r'_N \leq 0$, and $r'_M > 0$. Therefore at most one intersection of two curves $M(r, a) = 0$ and $N(r, a) = 0$ exists. $\square$

References.

Ábrahám Á, Koehne S, Pavoni, N (2011) On the first-order approach in principal–agent models with hidden borrowing and lending. Journal of Economic Theory 146(4): 1331–1361

Arnott R, Stiglitz JE (1988a) The basic analytics of moral hazard. Scandinavian Journal of Economics 90(3): 383–413

Arnott R, Stiglitz JE (1988b) Randomization with asymmetric information. The RAND Journal of Economics 19(3): 344–362

Arrow KJ (1963) Uncertainty and the welfare economics of medical care. The American Economic Review 53(5): 941–973

Arrow KJ (1971) Essays in the theory of risk-bearing. Amsterdam-London: North-Holland

Attar A, Chassagnon A (2009) On moral hazard and non-exclusive contracts. Journal of Mathematical Economics 45: 511-525

Bernoulli D (1738) Specimen theoriae novae de mensura sortis. Commentarii Academiae Scientiarum Imperialis Petropolitanae 5: 175-–192

Bernoulli D (1954) Exposition of a new theory on the measurement of risk. Translated by Sommer, L. Econometrica 22(1): 23-–36

Biais B, Mariotti T, Rochet, JC, Villeneuve S (2010) Large risks, limited liability, and dynamic moral hazard. Econometrica 78(1): 73–118

Bisin A, Guaitoli D (2004) Moral hazard with non-exclusive contracts. The Rand Journal of Economics 35(2):306–328

Blomqvist A (1977) Optimal non-linear health insurance. Journal of Health Insurance 16(3): 303–321

Bolton P, Dewatripont M (2005) Contract Theory. M.I.T. Press, Cambridge, MA

Borch K (1962) Equilibrium in a reinsurance market. Econometrica 30(3): 424–444

Chiappori PA, Macho I, Rey P, Salanié (1994) Repeated moral hazard: The role of memory, commitment and the access to credit markets. European Economic Review 38(8): 1527–1553

Dasgupta P, Hammond P, Maskin E (1979) The implementation of social choice rules: Some general results on incentive compatibility. Review of Economic Studies 46(2): 185–216

DeMarzo PM, Sannikov, Y (2006) Optimal security design and dynamic capital structure in a continuous-time agency model. The Journal of Finance 61(6): 2681–2724

Dembe AE, Boden LI (2000) Moral hazard: A question of morality?. New Solutions 10(3): 257–279

Eeckhoudt L, Gollier C, Schlesinger H (2005) Economic and financial decisions under risk. Princeton University Press.

Ehrlich I, Becker G (1972) Market insurance, self insurance and self protection. Journal of Political Economy 80(4): 623–648

Ellis RP, McGuire TG (1990) Optimal payment systems for health services. Journal of Health Economics 9(4): 375–396

Fernandes A, Phelan C (2000) A recursive formulation for repeated agency with history dependence. Journal of Economic Theory 91(2): 223–247

Gale D, Hellwig M (1985) Incentive-compatible debt contracts: The one-period problem. The Review of Economic Studies 52(4): 647–663

Gibbard A (1973) Manipulation of voting schemes: A general result. Econometrica 41(4): 587–601

He Z (2012) Dynamic compensation contracts with private savings. The Review of Financial Studies 25(5): 1494–1549

Hobbes T (1651), Leviathan

Holmström B (1977) On incentives and control in organizations. Ph.D. Dissertation, Stanford University

Holmström B (1979) Moral hazard and observability. The Bell Journal of Economics 10(1): 74–91

Holmström B (1999) Managerial incentive problems: A dynamic perspective. The Review of Economic Studies 66(1): 169–182

Holmström B, Milgrom P (1991) Multi-task principal-agent analyses: Linear contracts, asset ownership and job design. Journal of Law, Economics and Organization 7: 24–52

Jarque A (2010). Repeated moral hazard with effort persistence. Journal of Economic Theory, 145(6), 2412–2423

Jensen MC, Meckling WH (1976) Theory of the firm: Managerial behavior, agency costs, and ownership structure. Journal of Financial Economics 3(4): 305–360

Jewitt I (1988) Justifying the first-order approach to principal-agent problems. Econometrica 56(5): 1177–1190

Kwon I (2006) Incentives, wages, and promotions: Theory and evidence. The RAND Journal of Economics 37(1): 100–120

Ma CTA, McGuire TG (1997) Optimal health insurance and provider payment. The American Economic Review 87(4): 685–704

Ma CTA, Riordan M (2002) Health insurance, moral hazard, and managed care. Journal of Economics & Management Strategy 11(1): 81–107

Main B (1982) The firm's insurance decision. Some questions raised by the capital asset pricing model. Managerial and Decision Economics 3(1): 7–15

Malcomson J, Spinnewyn F (1988) The multiperiod principal agent problem. The Review of Economic Studies 55(3): 391–408

McNeely MC (1941) Illegality as a factor in liability insurance. Columbia Law Review 41(1): 26–60

Mirrlees JA (1999) The theory of moral hazard and unobservable behaviour-Part I. The Review of Economic Studies 66(1): 3–21

Mirrlees JA (1976) The optimal structure of incentives and authority within an organization. The Bell Journal of Economics 7(1): 105–131

Mukoyama T, Sahin A (2005) Repeated moral hazard with persistence. Economic Theory 25(4): 831–854

Myers S (1977) Determinants of corporate borrowing. Journal of Financial Economics 5(2): 147–75

Myerson RB (1979) Incentive compatibility and the bargaining problem. Econometrica 47(1): 61–73

Park IU (2004) Moral hazard contracting and private credit markets. Econometrica 72(3): 701–746

Pauly M (1968) The economics of moral hazard: Comment. The American Economic Review 58(3): 531–537

Rees R, Wambach A (2008) The microeconomics of insurance. Foundations and Trends in Microeconomics 4(1–2): 1–163

Rey P, Salanié B (1990) Long-term, short-term and renegotiation: On the value of commitment in contracting. Econometrica 58(3) 597–619

Rogerson W (1985a) The first-order approach to principal-agent problems. Econometrica 53(6): 1357–1368

Rogerson W (1985b). Repeated moral hazard. Econometrica 53(1): 69–76

Salanié B (2005) The economics of contracts: A primer (second edition), M.I.T. press

Shavell S (1979a) Risk-sharing and incentives in the principal and agent relationship. The Bell Journal of Economics 10(1): 55–73

Shavell S (1979b) On moral hazard and insurance. The Quarterly Journal of Economics 93(4): 541–562

Shavell S (1980) Strict liability versus negligence. The Journal of Legal Studies 9(1): 1–25

Shavell S (1987) Economic analysis of accident law. Harvard University Press

Shavell S (2007) Liability for accidents. Handbook of Law and Economics, Vol. 1, Edited by AM Polinsky and S Shavell

Townsend R (1979) Optimal contracts and competitive markets with costly state verification. Journal of Economic Theory 21(2): 265–293

Winter RA (1988) The dynamics of competitive insurance markets. Journal of Financial

Intermediation 3(4): 379–415

Winter RA (1991) The liability insurance market. Journal of Economic Perspectives 5(3): 115–136

Zeckhauser R (1970) Medical insurance: A case study of the tradeoff between risk spreading and appropriate incentives. Journal of Economic Theory 2(1): 10–26