i. INSTITUT NATIONAL DE LA STATISTIQUE ET DES ETUDES ECONOMIQUES
Série des Documents de Travail du CREST
(Centre de Recherche en Economie et Statistique)

n° 2002-42

Asymmetric Information in Insurance : General Testable Implications*

P.-A. CHIAPPORI[1] B. JULLIEN[2]
B. SALANIÉ[3] F. SALANIÉ[4]

---

* We thank participants at the European Group of Risk and Insurance Economists, Mathias Dewatripont, Victor Ginsburgh, Etienne Wasmer, Lucy White and especially David de Meza, Pierre Picard and David Webb for comments. Parts of this paper were written while Bernard Salanié was visiting the University of Chicago, which he thanks for its hospitality. Support from the NSF (grant#0096516) is gratefully acknowledged.
[1] University of Chicago, corresponding author. Email : pchiappo@midway.uchicago.edu
[2] University of Toulouse I (GREMAQ-CNRS and IDEI).
[3] CREST, CNRS URA 2200 and CEPR.
[4] University of Toulouse I (LEERNA-INRA)

Abstract

Several recent papers on empirical contract theory and insurance have tested for a positive correlation between coverage and ex post risk, as predicted by standard models of pure adverse selection or pure moral hazard. However, these models rely on strong and empirically implausible assumptions (such as one dimensionality, identical preferences, etc.). We provide a testable implication of asymmetric information that is valid in a very general set-up. We then show that the positive correlation property can in fact be extended to competitive insurance markets, and also to cases where risk aversion is public. We also provide an empirical test of our results on a new French dataset.

Résumé

Plusieurs articles récents sur l'économétrie de l'assurance ont testé l'existence d'une corrélation positive entre couverture et risque, qui est prédite par les modèles usuels d'antisélection pure et d'aléa moral. Ces modèles reposent toutefois sur des hypothèses très restrictives. Nous donnons ici une implication testable de l'information asymétrique qui est valable dans une classe de modèles très large. Nous montrons ensuite comment la propriété de corrélation positive s'étend quand le marché de l'assurance est concurrentiel ou quand l'aversion pour le risque est publique. Nous procédons ensuite à un test de nos résultats sur des données françaises.

1 Introduction

While the economics of insurance under asymmetric information dates back to the 1970s, only recently has there been extensive testing of its theoretical conclusions. A standard problem facing any empirical work on the topic is that the robustness of the testable predictions derived by existing theory is often unclear. Theoretical asymmetric information models typically use oversimplified frameworks, that can hardly be directly transposed to real life situations. To give but one example, Rothschild-Stiglitz's (1976) celebrated model of competition under adverse selection in insurance assumes that accident probabilities are exogenous (which rules out moral hazard), that only one level of loss is possible, and more strikingly that agents have identical preferences which are moreover perfectly known to the insurer. The theoretical justification of these restrictions is straightforward: analyzing a model of "pure", one-dimensional adverse selection is an indispensable first step. But their empirical relevance is dubious, to say the least. In "real life" insurance, moral hazard can hardly be discarded a priori (and interacts with adverse selection in a non-trivial way, as precaution depends on risk and preferences[1]); losses are continuous variables, often ranging from small amounts to hundreds of thousands of dollars; last but not least, preference heterogeneity is paramount and largely unobserved.

All this clearly suggests that an indispensable prerequisite for any empirical work is the *theoretical* derivation of robust predictions that can be taken to data. This is the first goal of the present paper. Specifically, we concentrate on a central property of asymmetric information models in insurance, on which recent empirical work has largely focussed[2]. The property states that under both moral hazard and adverse selection, one should observe of a *positive correlation* (conditional on observables) between risk and coverage: if different insurance contracts are actually sold to observationally identical agents, then the frequency of accidents among the subscribers of a contract should increase with the coverage it offers.[3] In the Rothschild-Stiglitz (1976) model, where riskiness is an exogenous and unobservable characteristic of agents, the correlation stems from the fact that "high risk" agents are ready

[1] See Chassagnon-Chiappori (1997)), de Meza-Webb (2001), Jullien-Salanié-Salanié (2001).

[2] See for instance Chiappori-Salanié (2000) and the references in Chiappori (2000).

[3] While this paper focusses on the insurance sector, the methodology developed here could be useful in other cases. For example, one of the first papers to test the Stiglitz-Weiss (1981) theory is Ausubel (1999), in the context of credit cards. Ausubel finds convincing evidence of adverse selection, through a similar test of correlation: customers who accept higher interest rates are more likely to default.

to pay more than "low risk" ones for additional coverage, and will therefore choose contracts with higher coverage. Under pure moral hazard, as in Arnott-Stiglitz (1988), an opposite causality generates the same correlation: an agent who, for any unspecified (and exogenous) reason, switches to a contract with greater coverage makes less effort and thus becomes riskier.

Popular as this prediction may be, its robustness is, in principle, not guaranteed; whether it would remain valid in the presence of moral hazard, heterogeneous preferences or multiple levels of losses has not (yet) been demonstrated.[4] The first part of our paper is devoted to a *theoretical* analysis of this issue. We show that the original intuition derived from Rothschild-Stiglitz extends to more general models, as already conjectured by Chiappori and Salanié (2000), although its scope and robustness varies with the type of competition at stake. Specifically, we extend the property in three directions. First, using a revealed preference argument, we derive a new property that is robust to any assumption on the nature of competition. Second, we consider the case of competitive markets, and show that asymmetric information (with any combination of adverse selection and moral hazard) indeed implies a positive correlation between risk and coverage, for suitably defined such notions. This result is a direct extension of Rothschild-Stiglitz's initial idea to a very general framework (entailing heterogeneous preferences, multiple level of losses, multidimensional adverse selection plus possibly moral hazard, and even non-expected utility). Third, we study the case of imperfect competition, and we underline the key role of the agent's risk-aversion. If it is public information, then some form of positive correlation is verified. In particular, with only one level of loss and expected utility, contracts with higher coverage must exhibit a larger frequency of accidents. Conversely, if risk-aversion is private information, the property does not necessarily hold: this was shown in Jullien-Salanié-Salanié (2001). The aversion to risk thus is a key parameter whose informational status drives the testable implications of simple models in the presence of market power.

In the last part of the paper, we illustrate the theoretical analysis by testing the predictions it generates on real-life data. We first show that our revealed-preference result, when combined with a zero-profit condition, yields a prediction that can be tested from data on claims and reimbursement schedules only.[5] We test this prediction using data collected over the year

[4] Note, in particular, that in a general context, the frequency of accidents is only one indicator of riskiness, as the size of losses also matters; therefore the notion of positive correlation between risk and coverage is less straightforward.

[5] The 'revealed preference' prediction derived in our paper holds with positive profits as well, but testing it requires data on claims, reimbursement schedules and premia, which may be hard to obtain. The version we use fits better the data available to us.

1989 by a large French car insurer. We find that the estimated (generalized) correlation is close to and not significantly different from zero. This finding is consistent with the results of Chiappori-Salanié (1997, 2000), who however tested a much simpler and less robust variant of the theory developed in this paper.

Section 1 builds a general model of insurance under asymmetric information, that allows for non-expected utility preferences, multiple loss, adverse selection on risk and preferences, and moral hazard on risk. In Section 2, we apply a revealed preference argument to obtain a first testable implication, that relates the premium differential to expected indemnities. Section 3 analyses the stronger version of the correlation property; we show that it holds both when competition drives profits to zero and when risk aversion is public information. Section 4 tests the generalized positive correlation property derived in Section 3. Section 5 concludes.

2 The General Framework

Suppose that we observe a population of insurance policy holders, their insurance policies and their insurance claims. Typically an insurance contract specifies an indemnity $R(L) \geq 0$ for every possible claim level $L \geq 0$ and a premium P paid up-front. By definition, $R(0) = 0$ and we set $L = 0$ in the case of no claim. For each contract the indemnity function is fixed, but we allow the premium to vary with the characteristics of the insured that are observed by the insurer, which we denote X. In what follows, we assume that the econometrician observes X (from the insurer's files). Based on the data on observed claims and premia, the econometrician can estimate premia $P_i(X)$ and distributions of claims $F_i(L \mid X)$ conditional on X for each contract C_i. Our goal is to derive predictions that can be tested on such data.

For this, let us introduce a model allowing for both adverse selection and moral hazard. Consider a population of insurance policy holders that is indistinguishable for the insurers, which means that we control for characteristics X, and derive predictions valid for each value of X. From now on, we omit the variable X, although it should be clear that all results are conditioned on it. We thus denote P_i and $F_i(L)$ the premium and the empirical distribution of claims for contract C_i within the population of individuals with given characteristics X.

Each agent within this population faces the risk of an accident, equivalent to a monetary loss. Each agent can buy an insurance contract $C = (R(.), P)$. Note that the agent need not always report a loss, if it is associated with no indemnity. This is the case for instance when the loss is smaller than the deductible in the contract. For conciseness we identify claims and losses, but our predictions are valid for reported claims (see section 3.2). Each potential insured is characterized by a (possibly multidimensional) parameter θ, which is his private information. The parameter θ may affect the agent's preferences. Moreover an agent of type θ may secretly choose the distribution of claims F in some subset $\mathcal{F}^\theta$. The set $\mathcal{F}^\theta$ may be a singleton, as in pure adverse selection models, or include more than one choice, as when agents choose prevention efforts in moral hazard models. Within this very general setup, we make the following assumptions:

1. Each agent's preferences can be represented by a preference ordering over the final distribution of wealth, monotonic with respect to first order stochastic dominance.

2. Agents are risk averse in the sense that they are averse to mean-preserving spreads on wealth.

3. Risk-sharing: the net loss $L - R(L)$ is non-decreasing with L.

These assumptions are very weak. Models of insurance with risk-loving individuals do not seem to be very promising; and contracts for which $R(L)$ increase faster than L are almost systematically ruled out because of their perverse incentives properties.[6]

Under this form, it is clear that the class of models we consider encompasses most existing contributions, including the following works which all assume a Von Neumann Morgenstern utility function $u^\theta(W, F)$:

- *Pure adverse selection* (Rothschild-Stiglitz (1976) or Stiglitz (1977)): here $\mathcal{F}^\theta$ is a singleton. The Von Neumann-Morgenstern utility function u^θ does not depend on F, but it may depend on θ as in the multidimensional model of Lansberger-Meilijson (1999).

- *Moral hazard plus adverse selection on prevention cost* (Chassagnon-Chiappori (1997)): here $u^\theta(W, F) = v(W) - c^\theta(F)$, where v is common to all types of agents.

- *Moral hazard plus adverse selection on risk aversion.* In De Meza-Webb (2001), utility takes the form $u^\theta(W, F) = v^\theta(W) - c(F)$; in Jullien-Salanié-Salanié (2001), $u^\theta(W, F) = v^\theta(W - c(F))$. In both models, c is common to all types of agents, which differ only through their utility of wealth v^θ.

[6]Under such contracts, the agent gains in *worsening* the outcome of the accident (i.e., increasing the loss), a type of fraud that is extremely difficult to detect. This fact has been largely recognized, in particular, by the litterature on insurance fraud (see Picard 2000).

Lastly, it is important to stress what our results do *not* require. Although we allow for a general form of adverse selection (including multidimensional characteristics) plus possibly moral hazard, we do not impose any single-crossing condition. We do not restrict the number of types, nor their distribution. Neither do we assume expected utility maximization; our results hold in a non-expected utility framework as well.

3 A First Testable Implication

To compare two contracts C_1 and C_2 proposed on the market, we rely on the following simple definition:

Definition 1 *Contract C_2 covers more than contract C_1 if $R_2(L) - R_1(L)$ is non-decreasing.*

Two typical examples can illustrate the definition. In the case of two straight deductible contracts, where $R_i(L) = \max\{L - d_i, 0\}$, C_2 covers more than C_1 if $d_2 \leq d_1$. Also, in the case of two events, $L \in \{0, \bar{L}\}$, the condition amounts to $R_2(\bar{L}) \geq R_1(\bar{L})$.

We first establish a simple but useful revealed preference property.

3.1 A revealed preference argument

Assume that, when faced with two contracts C_1 and C_2 (where C_2 covers more than C_1), two observationally identical individuals make different choices (i.e., one chooses the contract with the lower coverage, C_1, while the other opts for C_2). First notice that as $R_2(L)$ is larger than $R_1(L)$ for all L, the premium must be higher for contract C_2 ($P_2 - P_1 > 0$), for otherwise choosing C_1would violate first order stochastic dominance. Risk aversion then allows us to strengthen the bound on the premia differential:

Proposition 1 *Assume that agent θ prefers contract C_1 to C_2, and C_2 covers more than C_1. Let F_1^θ be the distribution of claims of agent θ under C_1. Then*

$$P_2 - P_1 \geq \int_0^{+\infty} R_2(L) dF_1^\theta(L) - \int_0^{+\infty} R_1(L) dF_1^\theta(L) \tag{1}$$

Proof: *see Appendix*

The result states that if an agent chooses one contract over another with better coverage, the increase in premium must be sufficient for the expected income of the agent to decrease at unchanged behavior. If this were not the case, a risk neutral agent would prefer C_2 to C_1, and a fortiori a risk-averse agent.

As this result only uses revealed preference, it is very general. For instance, it still holds if there is some compulsory insurance, as it only involves the comparison between two available contracts, conditional on the fact that the agent buys a contract. Also, it does not require perfect competition: the property holds under monopoly or oligopoly as well.

A first trivial application is when there are only two events $L \in \{0, \bar{L}\}$. Contracts involve a single level of indemnity, so that $R_i(L)$ takes value 0 or $R_i(\bar{L})$. In this case let p_i be the empirical probability of a claim under contract C_i. If contract 2 covers more than contract 1, then $R_2(\bar{L}) > R_1(\bar{L})$, and (4) obviously gives

$$P_2 - P_1 \geq p_1(R_2(\bar{L}) - R_1(\bar{L})) \tag{2}$$

The results extends as follows to the case of two contracts with straight deductibles $d_1 > d_2$, $R_i(L) = \max\{L - d_i, 0\}$. From the empirical data, we can obtain the probability p_i that a positive claim occurs under C_i and the expected claim e_i conditional on a claim occurring. We then obtain $P_2 - p_1(e_1 - d_2) \geq P_1 - p_1(e_1 - d_1)$.

Corollary 2 *Suppose that C_2 and C_1 are two straight deductible contracts, and C_2 covers more than C_1. Let p_i be the probability of a claim under C_i. Then*

$$P_2 - P_1 \geq p_1(d_1 - d_2). \tag{3}$$

3.2 Testing the Implication

To turn Proposition 1 into a feasible test, we need to integrate the inequality over the set of agents with identical observable characteristics who choose contract C_1. The empirical distribution of claims for contract C_i within this population is $F_i(L) = E\left\{F_i^\theta(L) \mid C_i\right\}$. Then the inequality becomes: if C_2 covers more than C_1, then

$$P_2 - P_1 \geq \int_0^{+\infty} R_2(L)dF_1(L) - \int_0^{+\infty} R_1(L)dF_1(L). \tag{4}$$

Notice first that the empirical distribution of claims depends on the contract in two ways. First the contract affects the level of risk chosen by each insured under moral hazard. Second it affects the distribution of the types θ who chose contract C_i. Notice also that the test requires to have an estimate of the premium that the individuals would have to pay for contract C_2, which depend on the observable characteristics X. Thus the insurer's information on the insured must be known by the econometrician. An exception occurs when the insurer cannot legally discriminate on the basis on some variables (sex, race), which can thus be omitted.

Finally, it is important to show that (4) holds in settings when L is observable only if the insured reports a claim. Indeed, under contract C_1, it is possible that the insured does not declare some accidents L knowing that $R_1(L) = 0$. Nevertheless, and assuming away any declaration costs, the insured could have declared such accidents; denote G_1 the distribution of claims in this case. Note that the insured gets the same payoff under (C_1, F_1) and under (C_1, G_1). Since by assumption he prefers C_1 to C_2, then he must prefer (C_1, G_1) to (C_2, G_1). Therefore (4) must hold at G_1:

$$P_2 - P_1 \geq \int_0^{+\infty} R_2(L)dG_1(L) - \int_0^{+\infty} R_1(L)dG_1(L).$$

Now the second term of the right-hand-side is the same if one replaces G_1 by F_1, since these weights only differ at points where $R_1(L) = R_1(0) = 0$. And in the first term of the right-hand-side, replacing G_1 by F_1 reduces the expected indemnities, since some claims with $R_2(L) \geq 0$ are not declared anymore. Therefore the inequality remains valid if one replaces G_1 by F_1, as announced.

4 The Positive Correlation Property

The result in Proposition 1 provides a test that doesn't rely on the market structure, but requires estimating the conditional premia. However, this test does not translate obviously into a correlation structure between risk and coverage. This is not surprising. In contrast with the previous results, the positive correlation property cannot be expected to hold independently of the market structure or the information structure. We develop below two contexts in which the property indeed holds. Once again, we omit the observable variables X, although it should be clear that all results are conditioned on it.

4.1 Competitive Environment

As is well known, the mere definition of a competitive equilibrium under asymmetric information is a difficult task, on which it is fair to say that no general agreement has been reached. For the moment, we only make a mild assumption, namely that competition, whatever its particular form, leads to zero profits. Technically, let $\pi(C_i)$ be the profit the insurer makes on contract C_i. Then in the absence of loading or taxation, but allowing for a cost per contract K, we can write

$$\pi(C_i) = P_i - \int_0^{+\infty} R_i(L)dF_i(L) - K.$$

We thus assume the following:

Zero profit assumption : $\pi(C_i) = 0$ *for every contract that is traded.*

The zero profit assumption holds in the Rothschild-Stiglitz model and in fact in most theories of competitive equilibrium that have been proposed in the literature. An exception is the model of cross subsidies of Miyazaki, to which we will come back later. Of course, it needs not hold in non-competitive models such as Stiglitz (1977) or Jullien-Salanié-Salanié (2001).

Under the zero profit assumption, empirical riskiness and coverage are related as follows:

Proposition 3 *Assume that the zero profit assumption holds. If two contracts C_1 and C_2 are bought in equilibrium, and C_2 covers more than C_1, then*

$$\int_0^{\infty} R_2(L)dF_2(L) \geq \int_0^{\infty} R_2(L)dF_1(L). \tag{5}$$

If, moreover, the contracts are different and agents are strictly risk averse, then the inequality has to be strict.

Proof: From Proposition 1 we have for each θ that chose C_1 :

$$P_2 - \int_0^{+\infty} R_2(L)dF_1^{\theta}(L) \geq P_1 - \int_0^{+\infty} R_1(L)dF_1^{\theta}(L).$$

Taking the expectation conditional on the choice of C_1, we obtain:

$$P_2 - \int_0^{+\infty} R_2(L)dF_1(L) \geq P_1 - \int_0^{+\infty} R_1(L)dF_1(L).$$

The zero profit assumption gives us

$$P_2 - \int_0^{+\infty} R_2(L) dF_2(L) = P_1 - \int_0^{+\infty} R_1(L) dF_1(L)$$

Substracting these two equations immediately yields the result. Finally, to see why the inequality must be strict when contracts are different, assume not, then from the argument above (1) holds as an equality. But then all risk averse agents will strictly prefer contract C_2, which contradicts the assumption that both contracts are sold in equilibrium. ■

The results state that the empirical risk is larger for the contract with the higher coverage, in the sense that the average indemnity would be smaller with the distribution of claims of the other contract. Two remarks are in order at that stage:

1. The result only requires that profit doesn't increase with coverage, $\pi(C_1) \geq \pi(C_2)$. Thus, the zero profit condition is not necessary provided that the less profitable contract covers more. This is precisely the case for cross-subsidies a la Miyazaki, where the losses made on the full coverage contract (chosen by high risk agents) are subsidized by the profits stemming from the alternative, partial coverage contract (that attracts low risk agents).

2. While the inequality (5) holds true whatever the type of asymmetric information at stake, it boils down to an equality when the asymmetry is not related to risk, in the sense that the distribution of risk is identical accross contracts ($F_1 = F_2$). This is trivially the case in the absence of asymmetry, and also in models of pure adverse selection on preferences (or risk aversion): under fair pricing, the only equilibrium entails pooling. Of course, the interesting part of the Proposition is that whenever some information asymmetry on risk is involved (whether adverse selection, moral hazard or both), then one expects a menu of *different* contracts to be offered at equilibrium, and a *strict* inequality must hold.

The general insight can be summarized as follows. First assume that competition leads to actuarially fair contracts and yet our result does not hold: at least two contracts C_1 and C_2 are sold at equilibrium, and C_1 covers less than C_2 but has ex post riskier buyers. Since C_1 has higher ex post risk, its "unit price" (i.e., the ratio of premium to coverage) will be larger. But this leads to a contradiction, as under fair pricing, rational agents will never choose a contract entailing less coverage at a higher unit price.

Testing Proposition 3 only requires observing the insurers' observables X, two contracts, one of which has higher coverage, and being able to estimate the conditional distributions of claims. In particular it doesn't require to know the premia under the two contracts.

It is easy to derive consequences of this property. First note that a contract with full insurance, if available, must generate larger expected claims than any other contract. Second, in the case of straight deductibles, we obtain:

$$p_2 e_2 - p_1 e_1 \geq (p_2 - p_1) d_2. \tag{6}$$

Thus if contract C_2 leads to a higher probability of a claim, it must also generate larger expected claims.

Of particular theoretical interest is the case in which contracts specify a fixed level of reimbursement for any accident. Then the empirical riskiness must be positively correlated with the coverage, which is the test performed in Chiappori-Salanié (2000):

Corollary 4 *Assume that the zero profit assumption holds and that $L \in \{0, \bar{L}\}$. If two contracts C_1 and C_2 are bought in equilibrium, and C_2 covers more than C_1, then $p_1 \leq p_2$.*

It is easily seen that this corollary also holds if a constant administrative cost of processing a claim is allowed in the definition of profits, provided that contracts C_1 and C_2 have nonzero coverage. However, de Meza-Webb (2001) indeed offer a model in which agents choose between insurance and no insurance. Then costs per claim are only incurred for insured agents; and this changes the computation of the actuarial premium which allowed us to derive Proposition 3.[7] More generally, one may argue that a contract with higher coverage is also more comprehensive[8], so that costs per claim may be higher. Under general contracts and costs $c_i(L)$ which may differ across contracts, the result in Proposition 3 becomes

$$\int R_2(L)[dF_2(L) - dF_1(L)] \geq \int c_1(L)dF_1(L) - \int c_2(L)dF_2(L)$$

[7] This point is due to Koufopoulos (2001).

[8] Consider for example automobile insurance, for which the basic contract only covers damages to third parties. Extending the coverage to the damages incurred by the insuree requires that the insurer devote resources to estimating these damages (we thank David de Meza and David Webb for this remark).

and whether the left-hand-side remains positive now becomes an empirical question. Clearly more information is needed on costs per claim to provide a fully convincing test of this inequality.

Similar phenomena occur if one takes into account experience rating, taxation of indemnities or premia, or a loading factor. In the case of experience rating, the occurrence of an accident causes an increase in future premia, which can be approximated by a reduction in the indemnity $R_i(L)$; Proposition 1 and 3 then change accordingly. Similarly, any taxation modifies the computation of actuarial premia, and Proposition 3 must be restated.[9] In all these cases, a test of our predictions is still possible, provided some assumptions are made on these newly introduced parameters.

4.2 Expected Utility with Public Risk Aversion

While the previous section was dealing with competitive environments, we now allow for market power and imperfect competition. This generalization comes at a cost. In order to keep the correlation property, we need to assume that (i) the agent has a Von Neumann-Morgenstern utility function $u^\theta(W, F)$, and (ii) observationally identical agents exhibit the same risk aversion[10], the latter being thus independent of the distribution F. Under this assumption the utility function is determined up to an affine transformation:

There exists a function $v(W)$ such that, for any θ, one can write

$$u^\theta(W, F) = a^\theta(F)v(W) - c^\theta(F)$$

with $a^\theta(F) > 0$.

The class of models satisfying this assumption, although restrictive, includes the standard models of pure adverse selection à la Stiglitz (1977) and pure moral hazard à la Arnott-Stiglitz (1988), as well as more complex frameworks.

Let two contracts C_1 and C_2 be bought in equilibrium by some individuals within the population at stake. For $i = 1, 2$, denote $w_i(L) = v(-L + R_i(L) - P_i)$ the utility under contract i after a loss L. Then a simple revealed preference argument (see the Appendix) implies that

Proposition 5 *Under public aversion, for any two contracts,*

$$\int_0^{+\infty} (w_2(L) - w_1(L))(dF_2(L) - dF_1(L)) \geq 0. \tag{7}$$

[9] See the empirical application in this paper, where we also allow for a loading factor.

[10] This is equivalent to assuming that risk aversion is publicly observable, hence is included in the observables X.

This looks like a positive correlation property, but it involves the utility function v, which is unknown to the econometrician. The function $w_2 - w_1$ is increasing in the range where it is positive, but may not be in the negative range. However combined with (4), it yields some interesting conclusions in cases of interest. Its implication is clearest for the case of two events, accident and no accident:

Corollary 6 *Assume that risk aversion is public and that $L \in \{0, \bar{L}\}$. If two contracts C_1 and C_2 are bought in equilibrium, and C_2 covers more than C_1, then $p_1 \leq p_2$.*

Proof: From (2), we must have $P_2 > P_1$. For C_2 to be bought it must be that

$$R_2 - P_2 \geq R_1 - P_1$$

From (7) :

$$(p_2 - p_1)\left(v(-\bar{L} + R_2 - P_2) - v(-\bar{L} + R_1 - P_1) + v(-P_1) - v(-P_2)\right) \geq 0;$$

hence $p_2 \geq p_1$.■

This result was already known in the Rothschild-Stiglitz case. Our contribution here is to highlight the key role played by the assumption of identical risk-aversion. In particular, once agents have chosen their preventive efforts they can be ordered according to their riskiness; and then the assumption guarantees that agents which are *ex-post* riskier indeed prefer contracts with higher coverage.

Finally, the assumption of identical risk-aversion is necessary for the result to hold. The underlying intuition is simple, and can be described in the polar case of an insurance monopoly. Start with the benchmark situation where agents have identical risks, but different risk aversion. Then in the optimal monopoly contract, partial coverage is used to screen agents according to their risk aversion, exploiting the fact that more risk averse individuals are willing to pay more for additional coverage; typically, more risk averse agents are fully covered, while less risk averse clients reveal their type (and benefit from a lower premium) by accepting partial coverage. Now, introduce an infinitesimal difference is risk that is fully correlated with risk aversion; specifically, the more risk averse agents have a (slightly) smaller accident probability. The optimal contract will still offer more coverage for the more risk averse individuals, at a higher price, despite the fact that the aggregate risk for that population is (slightly) smaller - a pattern that creates a *negative* correlation between risk and coverage.[11]

[11] Of course, the situation just described is somewhat specific, because it relies on a strong, *exogenous* correlation between risk and coverage. A more interesting context is

5 An Empirical Test

Tests of the positive correlation between risk and coverage on insurance contracts have provided mixed results. Most papers on automobile insurance (see, e.g., Chiappori-Salanié (1997, 2000)) cannot reject the no-correlation null: there in fact appears to be no correlation between the coverage of a contract and the ex post riskiness of its subscribers. Puelz-Snow (1994) was an early exception; but Dionne-Gouriéroux-Vanasse (2001) attributes their result to the spurious effect of a linear specification. Cawley and Philipson (1999) find no evidence of a positive correlation in their study of life insurance contracts. On the other hand, the market for annuities seems to be plagued by adverse selection problems, as documented by Brugiavini (1993) and more recently Finkelstein-Poterba (2000); Bach (1998) reaches similar conclusions in her study of mortgage-related unemployment insurance contracts.

Since all of these papers rest on a simplified analysis of the insurance market, it is interesting to see whether the more general predictions we obtained in this paper fare better when taken to the data. Note that the maintained assumptions of the theory are different for each of our results. Proposition 1 only relies on a revealed preference argument, while Proposition 3 adds a zero-profit condition and Corollary 6 assumes that risk-aversion is public and losses are 0-1. Ideally, we would start by testing Proposition 1. Unfortunately, this relies on data on premia as well as claims and reimbursement schedules. So far we have not been able to obtain data of consistent quality in these three dimensions. The data used by Chiappori-Salanié was very good on contracts, but not on the size of claims. On the other hand, we do have data recorded by a large French car insurer in December 1989 that are very good on reimbursement schedules (in fact straight deductibles) and the size of claims, but unfortunately not on premia. This will allow us to test Proposition 3, since it does not require data on premia.

This dataset covers the two years between October 1, 1987 and September 30, 1989. We only kept those individuals that are observed for at least a full year, and we normalized the data so that our figures correspond to exactly one year. The dataset we use in the application comprises 69,892 policies. About half (34,288) of these policies entail comprehensive coverage. We focus on these policies because they are straight deductible contracts, with varying

studied by Jullien-Salanié-Salanié (2001), who consider a model where risk-aversion is the agent's private information and agents secretly choose some prevention effort (moral hazard). Then the correlation between preferences and realized risk is endogeneized; the authors show that a monopoly may optimally propose two contracts that may involve a violation of the positive correlation property.

Table 1: Deductibles for Comprehensive Policies

Deductible (francs) %	Number of policies
0	11,218
500	7,247
1,000	14,044
2,000	1,779

deductible levels. As shown in Table 1, the most common deductible levels are 0 franc and 1,000 francs[12]. We will use the 1,000 francs contract as our contract C_1 and the 0 franc contract as our C_2.

We start from Corollary 2, which tells us that we should have

$$P_2 - P_1 \geq p_1(d_1 - d_2)$$

where here $d_1 = 1,000$ and $d_2 = 0$. A first difficulty is to define claims, as p_i is the probability of a claim under contract C_i. For every claim in France, insurance experts assign responsibilities to the policyholders involved. A policyholder that is deemed not to be responsible is fully reimbursed; a responsible policyholder is only reimbursed if (s)he has a comprehensive policy, and then only up to the deductible. Thus we should clearly focus on claims in which the policyholder was responsible and incurred some damage; there are 4,529 such claims in the data[13], so that the average p (over all contracts) is about 0.065.

Remember that our predictions are conditional on all variables X that are observed by the insurer (and hopefully by the econometrician). There are a large number of such variables in the data. As in Chiappori-Salanié (2000), our approach is to define "cells" of policyholders with identical values of those X variables that prior studies have identified as the most relevant. We choose six 0-1 X variables:

- whether the policyholder has the best experience rating (a 50% bonus)
- whether (s)he is a man
- whether his/her car is relatively powerful

[12] One franc was about 16 cents in 1989.

[13] Only 268 policyholders have more than one responsible claim; we neglect these multiple occurrences in the following.

- whether his/her car is relatively expensive
- whether the driver is young or old
- whether the car is driven in an urban area.

This defines $2^6 = 64$ cells. Each of them holds about 1,000 policyholders on average, but some are much smaller; thus we drop from the analysis the 16 cells that contain fewer than 200 insurees.

Within each cell X, we first test whether $p_2(X) \geq p_1(X)$, as done by Chiappori-Salanié. We first estimate $p_1(X)$, $p_2(X)$ and their estimated standard errors, then we compute a Student statistic by dividing $p_2(X) - p_1(X)$ by its standard error. This yields a collection of $48 = 64 - 16$ numbers $t_P(X)$. Under the null hypothesis that $p_2(X) = p_1(X)$ for all X, these numbers should be distributed as a $N(0,1)$ normal distribution. Figure 1 shows the estimated nonparametric density of the $t_P(X)$ (weighted by cell sizes), along with the $N(0,1)$ density. Contrary to Chiappori-Salanié, we find clear evidence here that $p_2(X) > p_1(X)$, as the distribution of the Students lies to the right of the normal curve.[14]

Of course, this is not completely conclusive, since we have only derived this theoretical prediction (in corollaries 4 and 6) when there is only one loss level. Our dataset gives us a rather detailed breakdown of the costs of each claim for the insurer, for each type of guarantee. What matters to us is the total cost of the claim, which is easily reconstructed from the data. Figure 2 plots the estimated nonparametric density of the costs for all responsible claims. Clearly, it is very dispersed: some claims are very costly and some very cheap[15]. Thus the corollaries may not apply, and it is of some interest to test the generalized positive correlation property, as given in Proposition 3. It is indeed possible that while the contract with the 1,000 francs deductible has more claims, these are less costly than under the contract with no deductible.

In order to test this property, we need to state a zero-profit condition. This should take into account both the existence of a loading factor λ, that reflects processing costs and a normal rate of profit, and the regulatory tax rate t on premia. This tax rate is 18%; and we take the loading factor to be 25%, a figure often quoted by insurance economists. Our results of course are somewhat sensitive to this choice, but not dramatically so. Denoting

[14] A Kolmogorov-Smirnov test gives a value of 0,381, largely above the 1% critical value for 48 observations.

[15] There are indeed a few negative costs. This is due to payment rules among insurers: small claims are settled by the insurer of each policyholder involved, and lump-sum payments are sometimes exchanged.

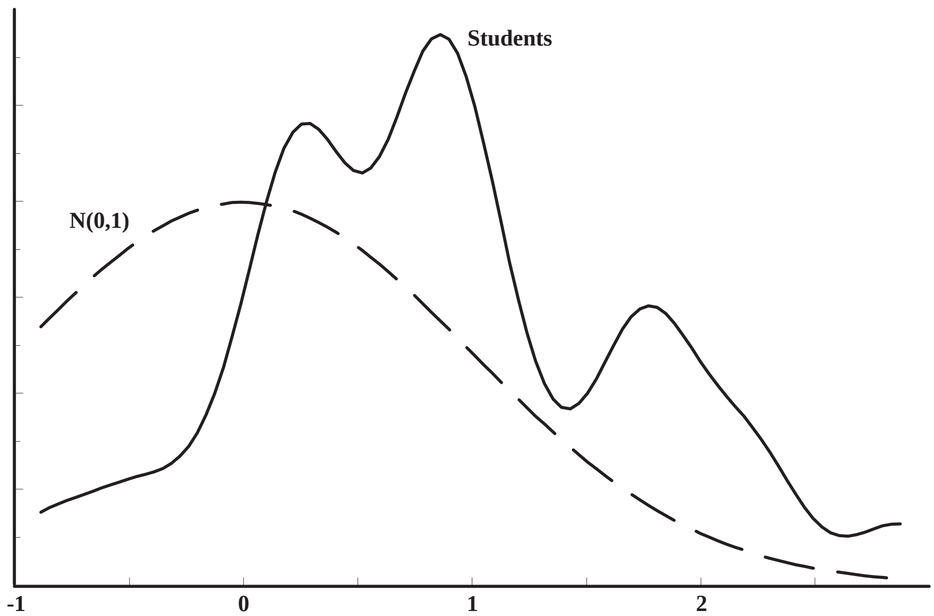

Figure 1: Studentized Estimates of $p_2 - p_1$

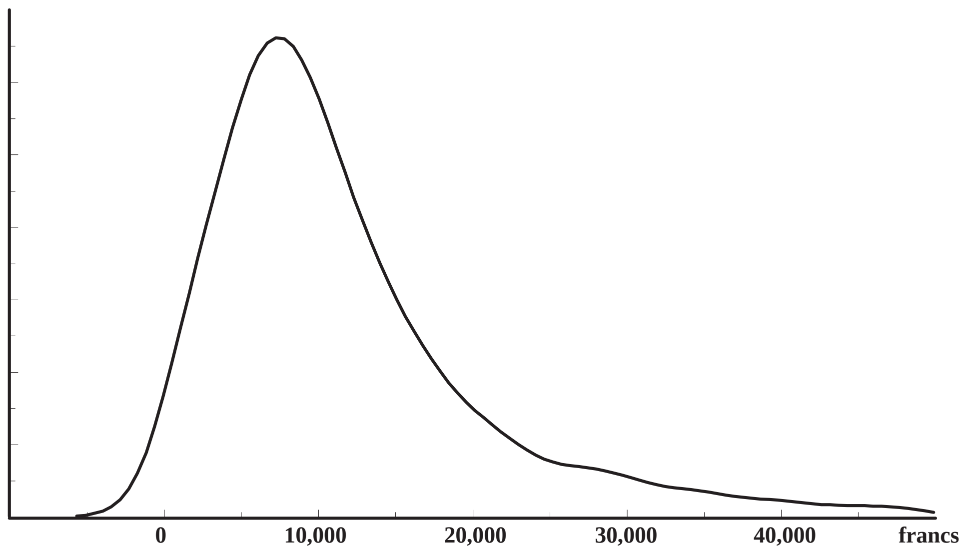

Figure 2: The Distribution of Costs of Claims

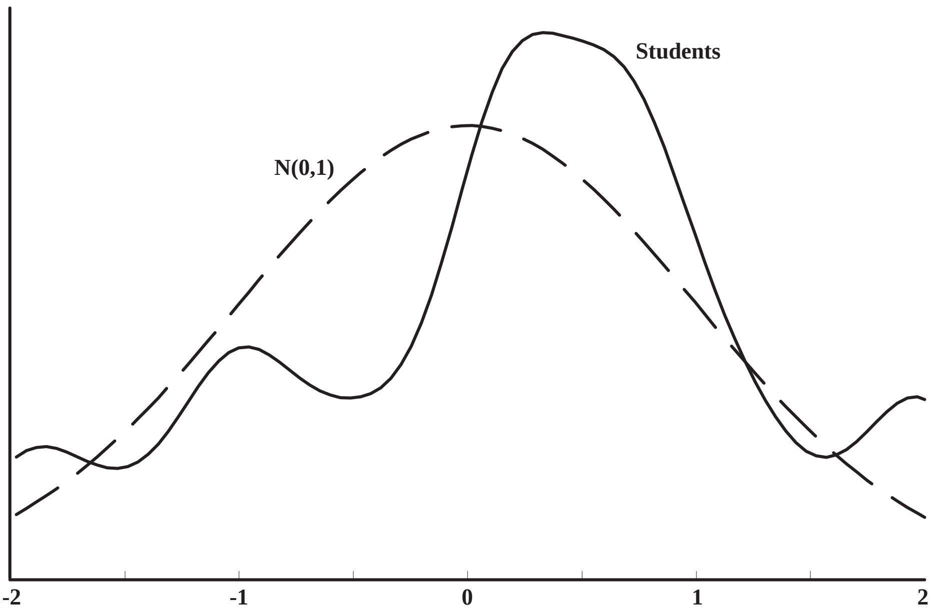

Figure 3: Studentized Estimates of the Generalized Positive Correlation

$\tilde{R}_i(X)$ to be the average total costs of claims under contract i for cell X, we therefore replace premia with

$$P_i(X) = (1+t)(1+\lambda)\tilde{R}_i(X)$$

in the inequality of Corollary 2. Thus we first estimate the quantity

$$(1+t)(1+\lambda)(\tilde{R}_2(X) - \tilde{R}_1(X)) - p_1(X)(d_1 - d_2)$$

then we standardize it and we obtain a new collection of Student statistics $t_G(X)$. Figure 3 shows the estimated nonparametric density of the $t_G(X)$ (again weighted by cell sizes), along with the $N(0,1)$ density. Now there is little evidence that the empirical distribution is to the right of $N(0,1)$, or indeed that it is asymmetric.[16] This confirms the findings of most of the literature.

[16] Indeed, a Kolmogorov-Smirnov test gives a value of 0,178, hence does not reject the null at 5%.

6 Conclusion

A first lesson stemming from this note is that in an asymmetric information context, a positive correlation between coverage and risk properly restated seems to be a natural and robust consequence of the competitive assumption. In that sense, our paper provides (somewhat *a posteriori*) a theoretical foundation for many existing empirical papers, although it points to the fact that the comparison of risk is not unambiguous and that a proper measure of risk must be used. Proposition 3 is characteristic of a competitive setting. Note nevertheless that one can weaken the zero profit assumption: the proof of Proposition 3 also goes through if we only assume that contracts with a greater coverage make (weakly) lower profits. This is for instance the case in equilibrium in the Miyazaki model of cross-subsidies; thus Proposition 3 also holds in that model. However, Proposition 3 must be restated with proportional loading or taxation, experience rating, or administrative costs of processing a claim.

Under imperfect competition, the zero profit assumption typically does not hold, and the correlation need not be positive. Indeed, the insurance companies extract rent from the policyholders, and optimal rent extraction may be such that more profit is extracted on contracts entailing more coverage. However, if risk aversion is public, which encompasses many frameworks (e.g. Rothschild-Stiglitz (1976) or Chassagnon-Chiappori (1997)) then at least with a single claim, the positive correlation property also holds. Notice however that public risk aversion is not a natural assumption in the context of insurance, as it eliminates any unobserved heterogeneity on a key determinant of the demand for insurance. Risk aversion clearly affects both the choice of an insurance policy and the precautionary attitude. Moreover it is an intrinsic property of preferences that cannot easily be observed by insurers.

Empirically, most data sets on automobile insurance (including the one studied in this paper) do not reject the null of zero correlation. The simplest explanation is probably the absence of significant asymmetric information in automobile insurance, although more complex stories can be evoked (see de Meza-Webb 2001 and Chiappori-Salanié 2000 for a detailed discussion). Finding a significant, *negative* correlation would raise a more serious challenge to the theory. Our paper, together with previous findings by Jullien-Salanié-Salanié (2001), suggests that the explanation should be grounded into market power and adverse selection on risk aversion. In fact, the theoretical results in this paper strongly suggest that there is a crying need for such models. An alternative is to turn the asymmetric information model on its head, by assuming that the insurer actually knows more than the insured.

This is done by Villeneuve (2000) within an otherwise standard hidden information model; he indeed finds that the correlation may be reversed, at least in a principal-agent framework. The competitive case however is more tricky, since competition tends in general (but not always) to result in full revelation.

Finally, since the positive correlation property has not fared well in empirical tests, it may be of interest to test predictions that rely on fewer assumptions. Proposition 1 provides one such prediction, as it only relies on a revealed preference argument and does not impose any particular market structure. In fact, it even holds in the Villeneuve model of an informed insurer, provided the conditioning variables X include all of the insurer's information. By the same token, a rejection of Proposition 1 would represent a rather strong challenge to the theory.

Bibliography

Arnott, R. and J. Stiglitz (1988), "The Basic Analytics of Moral Hazard", *Scandinavian Journal of Economics*, 90, 383-413.

Ausubel L. (1999), "Adverse Selection in the Credit Card Market", mimeo, University of Maryland.

Bach, K. (1998), *Negativauslese und Tarifdifferenzierung im Versicherungssektor*, DUV, Schesslitz.

Bennardo, A. and P.-A. Chiappori (2000), "Bertrand and Walras Equilibria under Moral Hazard", mimeo, University of Chicago, forthcoming in the *Journal of Political Economy*.

Brugiavini, A. (1993), "Uncertainty Resolution and the Timing of Annuity Purchases", *Journal of Public Economics*, 50, 31-62.

Cawley, J. and T. Philipson (1999), "An Empirical Examination of Information Barriers to Trade in Insurance", *American Economic Review*, 89, 827-846.

Chassagnon, A. and P.-A. Chiappori (1997), "Insurance Under Moral Hazard and Adverse Selection: The Case of Perfect Competition", mimeo, DELTA.

Chiappori, P.A. (2000), "Econometric Models of Insurance under Asymmetric Information", in G. Dionne, ed., *Handbook of Insurance*, North Holland, 365-94.

Chiappori, P.-A. and B. Salanié (1997), "Empirical Contract Theory: The Case of Insurance Data", *European Economic Review*, 41, 943-950.

Chiappori, P.-A. and B. Salanié (2000), "Testing for Asymmetric Information in Insurance Markets", *Journal of Political Economy*, 108, 56-78.

de Meza, D. and D. Webb (2001), "Advantageous Selection in Insurance Markets", *Rand Journal of Economics*, 32(2), 249-262.

Dionne, G., C. Gourieroux and C. Vanasse, (2001), "Testing for Evidence of Adverse Selection in the Automobile Insurance Market: A Comment", *Journal of Political Economy*, 109(2), April, pages 444-53.

Finkelstein, A. and J. Poterba (2000), "Adverse Selection in Insurance Markets: Policyholder Evidence from the UK Annuity Market", NBER DP 8045.

Jullien, B., B. Salanié and F. Salanié (2001), "Screening Risk-averse Agents Under Moral Hazard", mimeo IDEI Toulouse.

Koufopoulos, K. (2001), "A Note on 'Asymmetric Information in Insurance: a Robust Characterization' ", mimeo, London School of Economics.

Landsberger, M. and I. Meilijson (1996), “Extraction of Surplus under Adverse Selection: The Case of Insurance Markets”, *Journal of Economic Theory*, 69, 234-239.

Landsberger, M. and I. Meilijson (1999), “A General Model of Insurance under Adverse Selection”, *Economic Theory*, 14, 331-352.

Pauly M.V. (1974), “Overinsurance and Public Provision of Insurance: The Role of Moral Hazard and Adverse Selection”, *Quarterly Journal of Economics,* 88, 44-62.

Picard, P. (2000), “Economic Analysis of Insurance Fraud”, in G. Dionne, ed., *Handbook of Insurance*, North Holland, 315-62.

Puelz, R. and A. Snow (1994), “Evidence on Adverse Selection: Equilibrium Signalling and Cross-Subsidization in the Insurance Market”, *Journal of Political Economy*, 102, 236-57.

Rothschild, M. and J. Stiglitz (1976), “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information”, *Quarterly Journal of Economics*, 90, 629–649.

Shavell S. (1979), “On Moral Hazard and Insurance”, *Quarterly Journal of Economics*, 93, 541-562.

Stiglitz, J. (1977), “Monopoly, Nonlinear Pricing, and Imperfect Information: The Insurance Market”, *Review of Economic Studies*, 44, 407-430.

Stiglitz J. and A. Weiss (1981), ”Credit Rationing in Markets with Imperfect Information”, *American Economic Review,* 71, 393-410.

Villeneuve, B. (2000), “The Consequences for a Monopolistic Insurer of Evaluating Risk Better than Customers: The Adverse Selection Hypothesis Reversed ”, *Geneva Papers on Risk and Insurance Theory*, 25, 65-79.

Appendix :

Proof of Proposition 1*:* For any claim L, let $W_i(L) = R_i(L) - L - P_i$ be the resulting wealth under contract C_i. Fix the distribution of claims at F_1^θ and define the demeaned wealths

$$X_i(L) = W_i(L) - \int_0^{+\infty} W_i(L) dF_1^\theta(L)$$

By assumption, $X_i(L)$ is non-increasing. Since C_2 covers more than C_1, $(X_1(L) - X_2(L))$ also is non-increasing. It follows that for any X, the difference

$$\Delta(X) = \Pr(X_1(L) \leq X \mid F_1^\theta) - \Pr(X_2(L) \leq X \mid F_1^\theta)$$

is a function of X that can only change sign once, from positive to negative. Now consider the function

$$D(X_0) = \int_{-\infty}^{X_0} \Delta(X) dX$$

Clearly, D can only be increasing then decreasing. Moreover, $D(-\infty) = 0$, and by integrating by parts it is easily seen that

$$D(+\infty) = \int_0^{+\infty} X_2(L) dF_1^\theta(L) - \int_0^{+\infty} X_1(L) dF_1^\theta(L) = 0$$

Thus D is positive everywhere, which by definition implies that under F_1^θ, $X_1(L)$ is a single mean-preserving spread of $X_2(L)$.

Now agent θ prefers C_1 under F_1^θ to C_2 under any F, and in particular under F_1^θ. By assumption 2, the agent is averse to mean-preserving spreads; the fact that he chooses C_1 thus implies that the expected wealth under (C_1, F_1^θ) is larger than that under (C_2, F_1^θ), i.e.

$$\int_0^{+\infty} W_1(L) dF_1^\theta(L) \geq \int_0^{+\infty} W_2(L) dF_1^\theta(L),$$

which yields the result. ■

Proof of Proposition 5 :

Assume that some type θ buys contract C_i, and chooses a probability F_1^θ under C_1. By a simple revealed preference argument, we must have:

$$\int_0^{+\infty} w_1(L) dF_1^\theta(L) \geq \int_0^{+\infty} w_2(L) dF_1^\theta(L)$$

Aggregating over the types buying C_1, we find that

$$\int_0^{+\infty} w_1(L)dF_1(L) \geq \int_0^{+\infty} w_2(L)dF_1(L)$$

With a similar argument applied on C_2 :

$$\int_0^{+\infty} w_2(L)dF_2(L) \geq \int_0^{+\infty} w_1(L)dF_1(L)$$

Taking the difference between the two inequalities yields the result.