

Asymmetric information in insurance: general testable implications

Pierre-André Chiappori*

Bruno Jullien**

Bernard Salanié*

and

François Salanié***

Several recent articles 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. We show here that the positive correlation property can be extended to general setups: competitive insurance markets and cases where risk aversion is public. We test our results on a French dataset. Our tests confirm that the estimated correlation is positive; they also suggest the presence of market power.

1. Introduction

■ Asymmetric information is present on all markets. Whatever the product or service sold, the seller almost never knows the buyer's preferences, nor the maximum price she would be willing to pay to acquire it. Similarly, the buyer is in general unlikely to have much information about the seller's production technology or marginal costs. Most of the time, however, this asymmetry is irrelevant. In a perfectly competitive setting, the seller would not benefit from a detailed knowledge of the buyer's willingness to pay, because he has to charge the competitive price; the buyer needs no information about the technology, since again all the information she needs is contained in the price. Hence, asymmetric information is in general both paramount and inconsequential.

A common feature of the examples above is that the value of the hidden information is private (in the sense that the payoff of the uninformed party does not depend on it for a given contract). Fagart (1996) proves under weak assumptions that competition in the private-values case always

* Columbia University; pc2167@columbia.edu, bs2237@columbia.edu.

** Toulouse School of Economics (CNRS-GREMAQ and IDEI); bjullien@cict.fr.

*** Toulouse School of Economics (INRA-LERNA); salanie@toulouse.inra.fr.

We thank participants at the European Group of Risk and Insurance Economists (Strasbourg, 2001), Mathias Dewatripont, Victor Ginsburgh, Etienne Wasmer, Lucy White, and especially David de Meza, Pierre Picard, David Webb, the Editor, and two referees for their comments. We are very grateful to Jean Pinquet for giving us access to the dataset we use in Section 4. Parts of this article were written while Bernard Salanié was visiting the University of Chicago, which he thanks for its hospitality. Chiappori gratefully acknowledges financial support by the National Science Foundation (Grant No. 0096516).

leads to an equilibrium, which is moreover efficient. For want of a better term, we will call such a case “irrelevant asymmetric information” in the present article. The main innovation of the literature on asymmetric information, as pioneered by Akerlof (1970), Rothschild and Stiglitz (1976), and many others, was to exhibit cases in which, on the contrary, asymmetric information was indeed relevant—and actually had important consequences for the existence and efficiency of competitive equilibrium. The key property driving this conclusion is the presence of “common values,” in the form of a link between an agent’s hidden information and the other party’s payoff. In the market for lemons for instance, the buyer’s payoff depends on the quality of the car, which is known only by the seller. Similarly, an insurer’s profit depends on the risk of the insurees who buy contracts from him.

When considering empirical applications of such models, the previous remarks have important consequences. One is that evidence of information asymmetries, while relatively easy to produce, are often of little interest unless the asymmetries are of the relevant type. To give only one example, agents are often faced with menus of contracts. Menus of contracts are indeed suggestive of asymmetric information. Most of the time, however, this asymmetry is irrelevant. New cars are offered in different colors, which indeed reflects the seller’s ignorance about the buyer’s taste. Still, market equilibrium will typically exist and be efficient as usual, as the buyer’s taste does not directly affect the seller’s payoff. Different levels of insurance coverage may be proposed to insurees, reflecting asymmetric information about risk aversion. Insofar as differences in risk aversion have no impact on the insurer’s profit, however, the Akerlof-Rothschild-Stiglitz conclusions do not apply, and standard analysis is still valid. This simply reflects the fact that in a competitive setting, the insuree’s true risk matters to the insurer, even conditional on the insuree’s contract choice, while risk aversion does not. The former is a case of common values, and the latter a case of private values.

Clearly, one should primarily be interested in testing for asymmetric information in the “relevant” case. The main purpose of the article is precisely to propose robust empirical tests of relevant information asymmetries. Throughout, we concentrate on the particular case of insurance contracts, both because the main theoretical contributions to competition under adverse selection (starting with Rothschild and Stiglitz’s seminal article) used this framework, and because a large fraction of existing empirical literature deals with insurance contracts. However, our conclusions are general, and the methodology developed here could be useful in other cases.

In the literature on insurance, the general notions just sketched lead to a well-known property on which recent empirical work has largely focused.¹ Under both moral hazard and “relevant” adverse selection, one should observe 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. In the Rothschild and Stiglitz (1976) model of competition under adverse selection, where riskiness is an exogenous and unobservable characteristic of agents, the correlation stems from the fact that “high-risk” agents are ready 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 and 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.

The “positive correlation” prediction is appealing, but its robustness may however be questioned—a standard problem facing any empirical work on the topic. Theoretical models of asymmetric information typically use oversimplified frameworks, which can hardly be directly transposed to real-life situations. Rothschild and Stiglitz’s model assumes that accident probabilities are exogenous (which rules out moral hazard), that only one level of loss is possible, and more

¹ See, for instance, Chiappori and Salanié (2000) and the references in Chiappori (2000). This type of property has actually been tested in different contexts. For example, one of the first articles to test the Stiglitz and Weiss (1981) theory of credit markets 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.

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, moral hazard can hardly be discarded *a priori* (and interacts with adverse selection in a nontrivial way, as precaution depends on risk and preferences²); losses are continuous variables, often ranging from small amounts to hundreds of thousands of dollars; and preference heterogeneity is paramount and largely unobserved.

The first part of our article is devoted to a theoretical analysis of this issue. We show that the positive correlation property derived from Rothschild and Stiglitz extends to much more general models, as already conjectured by Chiappori and Salanié (2000), although its form and robustness vary with the type of competition at stake. Specifically, we extend the property in two directions. First, we consider the case of competitive markets, and we show that relevant asymmetric information (with any combination of adverse selection and moral hazard that generates common values) indeed implies a positive correlation between risk and coverage, for suitably defined such notions. This result is a direct extension of Rothschild and 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 nonexpected utility). Second, 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 must hold. 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, and Salanié (2007). Risk aversion thus is a key parameter whose informational status drives the testable implications of simple models in the presence of market power.

In the second part of the article, we illustrate the theoretical analysis by testing the predictions it generates on a dataset collected by a large French car insurer. We first test the general relevance of our setting and, in particular, of the assumption that agents correctly assess their accident probability. Our test uses a revealed-preference argument that is robust to any assumption on the information structure or the nature of competition. We find that the data strongly corroborate the predicted property, which validates our approach. We then test for the positive correlation property, and we find evidence of a positive (generalized) correlation. A closer examination of the data suggests that the insurer’s profits are probably higher for contracts with a higher coverage, contrary to the predictions of competitive models. This suggests that more work should be devoted to analyzing imperfectly competitive models of insurance markets.

Section 2 builds a general model of insurance under asymmetric information. In Section 3 we apply a revealed-preference argument to obtain a first testable implication that relates the premium differential to expected indemnities. Section 4 analyzes the robust 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 5 tests the properties derived in Sections 3 and 4. Section 6 concludes.

2. The model

■ **The general framework.** Suppose that we observe a population of insurance policy holders, their insurance policies, and their insurance claims. Since we want to derive testable predictions about the insured agents’ choices, we must deal with heterogeneity across agents. Throughout this article we denote by X the characteristics of the insured that are observable by insurers, and we assume that the econometrician observes X (from the insurers’ files). A group of agents sharing the same values for X are *ex ante* indistinguishable from the insurers’ viewpoint and so must have access to the same set of insurance contracts. To derive predictions on their choice of a particular contract, we focus on such a group, which simply means that we work conditionally on the value of X ; for convenience we often omit the variable X .

² See Chassagnon and Chiappori (1997), de Meza and Webb (2001), Jullien, Salanié, and Salanié (2007).

Our theoretical framework thus focuses on unobservable heterogeneity. Agents face the risk of an accident, which we assume to be equivalent to a monetary loss $L \geq 0$. Each agent is characterized by a parameter $\theta \in \Theta$, which is his private information and may affect his preferences and/or his risk. An agent of type θ may secretly choose the distribution of losses G in some set \mathcal{G}^θ . This set may be a singleton, as in pure adverse-selection models, or it may include more than one choice, as when agents choose prevention efforts in moral hazard models.

Each agent can also buy an insurance contract $C = (R(\cdot), P)$, which specifies a premium P and an indemnity $R(L) \geq 0$ for every possible claim $L \geq 0$. By definition, $R(0) = 0$. In some situations the agent may also choose not to report a loss, for example if the loss is smaller than the deductible in the contract. For conciseness in notation, we identify claims and losses here, but we shall see that our predictions extend to more complex settings. Hence the final wealth of the agent is $W(L) = W_0 + R(L) - L - P$, where W_0 denotes initial wealth. Throughout the article we concentrate on reimbursement schedules such that a higher loss is always bad news for the agents, in the sense that the net loss $L - R(L)$ is nondecreasing with L . Not only is this property satisfied empirically, but it also relies on compelling theoretical arguments.³

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 nondecreasing.

In the oft-treated case of two events $L \in \{0, \bar{L}\}$, the condition reduces to $R_2(\bar{L}) \geq R_1(\bar{L})$. In the empirically prevalent case of contracts with straight deductibles $R_i(L) = \max\{L - d_i, 0\}$, the definition amounts to $d_2 \leq d_1$. Similarly, for contracts offering constant copayment $R_i(L) = \alpha_i L$, our criterion is equivalent to $\alpha_2 \geq \alpha_1$. As it turns out, most real-life contracts (including the contracts considered in our empirical application) belong to these classes.

From a more theoretical viewpoint, the definition requires that C_2 offer relatively greater indemnities when L is higher; one implication is that a contract with more coverage reduces the risk on final wealth, in a precise sense that is discussed below. Obviously, the “order” induced by our notion of coverage is only partial; i.e., for two arbitrary contracts, it may well be the case that none covers more than the other. Then an agent’s decision (and the empirical tests based on it) cannot be analyzed without a precise knowledge of the agent’s preferences.

□ **Behavioral assumptions.** Within this very general setup, we state the following assumptions:

Assumption 1. Each agent’s preferences can be represented by a state-independent preference ordering over the distribution of final wealth,⁴ monotonic with respect to first-order stochastic dominance.

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

Assumption 3 (“Realistic Expectations”). When agents choose a contract, they correctly assess their accident probability and loss distribution; i.e., they use the true loss distribution G .

The first two assumptions are very weak. Assumption 3 is stronger but indispensable for empirical applications, since the agents’ subjective assessments of their risk are not observable. Our first task will be to provide a test allowing us to check the empirical relevance of this assumption.

³ Contracts for which $R(L)$ increases faster than L are almost systematically ruled out because of usual risk-sharing arguments, as well as because of their perverse incentive properties. Indeed, under such contracts, the agent gains in exaggerating the loss, a type of fraud that is extremely difficult to detect. This point has been widely emphasized by the literature on insurance fraud (see Picard, 2000).

⁴ This implicitly rules out situations in which the agents’ choices are partly driven by extrinsic, unobserved characteristics such as the reliability of the insurer, its location, etc. It fits well with empirical studies involving choices between different contracts offered by the *same* company, as is the case in the application below; it may, however, be more problematic when different companies are considered.

Before describing such a test, one should first emphasize that the class of models we consider encompasses most existing contributions, including the following works, which all assume a von Neumann-Morgenstern utility function $u^\theta(W, G)$:

- (i) *Pure adverse selection* (Rothschild and Stiglitz, 1976, or Stiglitz, 1977). Here \mathcal{G}^θ is a singleton. The von Neumann-Morgenstern utility function u^θ does not depend on G , but it may depend on θ as in the multidimensional model of Landsberger and Meilijson (1999).
- (ii) *Pure moral hazard* (e.g., Arnott and Stiglitz, 1988). Here Θ is a singleton, since agents are *ex ante* identical, while \mathcal{G}^θ has at least two elements.
- (iii) *Moral hazard plus adverse selection on prevention cost* (Chassagnon and Chiappori, 1997). Here $u^\theta(W, G) = v(W) - c^\theta(G)$, where v is common to all types of agents.
- (iv) *Moral hazard plus adverse selection on risk aversion*. In de Meza and Webb (2001), utility takes the form $u^\theta(W, G) = v^\theta(W) - c(G)$; in Jullien, Salanié, and Salanié (2007), $u^\theta(W, G) = v^\theta(W - c(G))$. In both models, c is common to all types of agents, which differ only through their utility of wealth v^θ .

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 nonexpected utility framework as well, provided agents are risk averse.

3. Testing the realistic expectations assumption

■ A test of Assumption 3 must check that the agent’s subjective assessment of her own risk is not at odds with her true loss distribution. Since subjective representations are not directly accessible, they must be inferred from the agent’s decision, i.e., contract choice. Hence, ultimately, the test must boil down to some consistency property between agent’s choices and true underlying probabilities. Ideally, the test should be valid irrespective of the information structure (i.e., whether information is symmetric or not) and the competitive context. We now present a general test that meets these criteria.

□ **A revealed-preference argument.** The test relies on a simple but robust revealed-preference property. Assume that an agent prefers a contract C_1 to a contract C_2 that covers more. 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_1 would violate first-order stochastic dominance. Risk aversion allows this bound to be strengthened on the premia differential:

Lemma 2. Assume that an agent prefers contract C_1 to C_2 , and C_2 covers more than C_1 . Let G be the distribution of claims as anticipated by the agent under C_1 . Then under Assumptions 1 and 2,

$$P_2 - P_1 \geq \int [R_2(L) - R_1(L)]dG(L). \tag{1}$$

Proof. See Appendix A.

The proof relies on a standard argument. Under Definition 1, for any given distribution G , the final wealth under C_1 is riskier than the final wealth under C_2 . Since the insured dislikes mean-preserving spreads, it must then be that the final wealth under C_1 has a higher expectation to compensate for its higher riskiness; otherwise the agent would gain by buying C_2 while keeping the same distribution of claims.

In words, if an agent chooses one contract over another with better coverage, the decrease in premium must be sufficient for the expected income of the agent to increase at unchanged behavior. As this result uses only revealed preference, it is very general. For instance, it still holds if there is some compulsory insurance, as it involves only the comparison between two available

contracts, conditional on the fact that the agent buys a contract. Also, it is valid with symmetric or asymmetric information, and it does not require perfect competition—the property holds under monopoly or oligopoly as well.

Lemma 1 was obtained under the assumption that the agent reports all losses. However, it turns out that it can be extended to more complex settings. We refer the reader to the Appendix for a proof that Lemma 1 still holds when this assumption is relaxed.

□ **A testable prediction.** Lemma 1 was phrased in terms of a revealed-preference argument, valid at the agent's level. To obtain a testable prediction, let us consider a group of indistinguishable agents, with the same values for the observable variables X . Assume that the insurers' data show that some of them bought contract C_1 while others bought contract C_2 . The premia P_2 and P_1 are known, and the data also allow us to compute an empirical distribution of claims $F_1(L)$ for contract C_1 . We then have the following result.

Proposition 1. Suppose that contract C_2 covers more than contract C_1 , and that both contracts are sold to indistinguishable agents with realistic expectations. Then under Assumptions 1, 2, and 3,

$$P_2 - P_1 \geq \int [R_2(L) - R_1(L)] dF_1(L). \quad (2)$$

In most cases, insurers set the levels of premia according to observables X but leave indemnities unchanged for each contract. Then the prediction must hold for any group of indistinguishable agents, that is, for any value of X , and one may test the following family of inequalities:

$$\forall X, \quad P_2(X) - P_1(X) \geq \int [R_2(L) - R_1(L)] dF_1(L | X).$$

To illustrate the prediction, let us apply it to the case when there are only two events $L \in \{0, \bar{L}\}$. In this case let q_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 (2) obviously gives

$$P_2 - P_1 \geq q_1(R_2(\bar{L}) - R_1(\bar{L})). \quad (3)$$

The result may also be applied to the case of two contracts with straight deductibles $R_i(L) = \max\{L - d_i, 0\}$, with $d_1 \geq d_2$. From the data we can compute the probability $q_1 = \text{Prob}(L > d_1)$ that a positive indemnity is transferred under C_1 . Then $(R_2 - R_1)$ is zero in the absence of claim or when $L < d_2$, is nonnegative for $L \in [d_2, d_1]$, and is equal to $(d_1 - d_2)$ otherwise. We get the following.

Corollary 1. Suppose that C_2 and C_1 are two straight deductible contracts, and C_2 covers more than C_1 ($d_1 \geq d_2$). Let q_1 be the probability that L is above d_1 under C_1 . Then

$$P_2 - P_1 \geq q_1(d_1 - d_2). \quad (4)$$

Proposition 1 provides a test of the realistic-expectations assumption, in the sense that the property should hold under the null that agents perfectly know their distribution of losses. Its power should be qualified, however, for the following reason. Assume that agents are subject to biases in their risk perceptions. Then (1) applies only to the distribution G as perceived by each agent, and not to the agent's true distribution of claims F . Suppose however that agents are pessimistic and overestimate the risk, in the sense that $G \leq F$. Then

$$\int [R_2(L) - R_1(L)][dG(L) - dF(L)] = \int \frac{d}{dL} [R_2(L) - R_1(L)][F(L) - G(L)] dL \geq 0$$

and (1) remains valid when replacing G by F . Hence Proposition 1 and the above applications still hold when agents are pessimistic.

Technically, we thus test the null of realistic expectations against the assumption that agents are optimistic, in the sense that they systematically underestimate their risk. Note, however, that optimism is frequently assumed in the literature. For instance, in a model analyzed by Koufopoulos (2005), competitive equilibria are such that optimistic agents choose a low coverage but do not make any prevention effort, while pessimistic agents choose a high coverage and exert some effort to reduce their probability of accident. Then both (1) and the positive correlation property below may be violated, if some agents are optimistic enough.

4. The positive correlation property

■ The result in Proposition 1 provides a test that does not rely on the market structure, but which obviously does not translate 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 competitive context or the information structure. Below we study two cases 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.

□ **A prediction in a 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 assume that competition, whatever its particular form, leads to profits that are nonincreasing in coverage. Also, we assume away proportional costs; moreover, while fixed administrative costs are allowed for, we assume they are identical across contracts.⁵ Technically, let F_i be the empirical distribution of claims under C_i , and let $\pi(C_i)$ be the profit the insurer makes on contract C_i . Then

$$\pi(C_i) = P_i - \int R_i(L)dF_i(L) - \Gamma,$$

where Γ denotes fixed costs associated with the contract. We then assume the following.

Nonincreasing profit assumption (NIP). If C_2 covers more than C_1 , then $\pi(C_2) \leq \pi(C_1)$.

This assumption trivially holds in the Rothschild and Stiglitz model and in fact in most theories of competitive equilibrium that have been proposed in the literature, as competition drives profits to zero on every contract. It is, however, more general than a standard zero-profit assumption, since it also holds in the cross-subsidies model of Miyazaki (1977), 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). Of course, it need not hold in noncompetitive models such as Stiglitz (1977) or Jullien, Salanié, and Salanié (2007).

Now the NIP assumption may be rewritten as

$$P_2 - P_1 \leq \int R_2(L)dF_2(L) - \int R_1(L)dF_1(L),$$

and replacing in (2) we get the following proposition.

Proposition 2. Assume that Assumptions 1, 2, 3, and NIP hold. If two contracts C_1 and C_2 are bought in equilibrium, and C_2 covers more than C_1 , then

$$\int R_2(L)dF_2(L) \geq \int R_2(L)dF_1(L). \tag{5}$$

⁵ Note that this assumption cannot be satisfied in the special case when contract C_1 is the no-insurance contract: then administrative costs are not incurred for C_1 , and this changes the expression for the actuarial premium. de Meza and Webb (2001) indeed offer a model in which agents choose between insurance and no insurance, and in which insured agents may be less risky *ex post* than uninsured agents.

The result states 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. 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.

Note that in models of pure adverse selection on preferences (risk aversion for instance), the distribution of claims is identical across contracts. Just as in models with symmetric information, (5) then is obvious, as F_1 and F_2 are identical. This is the “irrelevant” case of private values discussed in the Introduction. Agents do have some private information, but this information is of no relevance for the insurer: it is clear from the expression of profits that only hidden information that affects the distribution of risk matters to him.

Of course, the interesting part of Proposition 2 is that whenever information asymmetry on the distribution of risk is involved and a menu of different contracts is offered and sold at equilibrium, a strict inequality must hold. In that case, the asymmetry is relevant, in the sense that the insuree’s information matters for the insurer’s profit. All the specific features derived in the literature (cream-skimming, nonexistence or inefficiency of equilibrium, etc.) only can occur in this case. This obviously includes adverse selection on risk, but also adverse selection on risk aversion in the presence of moral hazard.

It is easy to derive consequences of (5). Of particular theoretical interest is the case in which contracts specify a fixed level of reimbursement for any accident. Denote again q_i the probability of a claim under contract C_i . Then the empirical riskiness must be positively correlated with the coverage, which is the test performed in Chiappori and Salanié (1997, 2000):

Corollary 2. Assume that Assumptions 1, 2, 3, and NIP hold 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 $q_1 \leq q_2$.

In the case of general loss distributions, let C_1 and C_2 be straight deductible contracts, assume that losses smaller than the value of the deductible are not reported, and define the expected claims under contract C_i as $E_i[L] = \int L dF_i(L)$. Then we obtain

$$E_2[L] - E_1[L] \geq d_2(q_2 - q_1). \quad (6)$$

Thus if contract C_2 leads to a higher probability of a claim, it must also generate larger expected claims. In particular, a contract with full insurance, if available, must generate larger expected claims than any other straight deductible contract (apply (6) with $d_2 = 0$).

□ **Testing the prediction.** Testing Proposition 2 or its corollaries only requires observing the variables X in the insurer’s files, two contracts, one of which has higher coverage, and being able to estimate the conditional distributions of claims. In particular, it does not require knowing the premia under the two contracts. Rejecting the prediction would thus tend to indicate that our first prediction (2) and/or the NIP assumption do not hold.

In practice, the definition of profits may need to be refined. Taxation of premia (or indemnities) may have to be taken into account. Similarly, it is often reported that insurers compute premia by applying a constant “loading factor” λ to the expectation of indemnities. In such cases, one has to rewrite the definition of profits accordingly:

$$\pi(C_i) = \frac{P_i}{1+t} - (1+\lambda) \int R_i(L) dF_i(L).$$

Then the NIP assumption still allows us to compute an upper bound for the premia differential $P_2 - P_1$. Replacing in (1) then yields a testable prediction,

$$\int R_2(L) dF_2(L) - \int R_2(L) dF_1(L) \geq K(E_1[L] - E_2[L]), \quad (7)$$

denoting $K = (1 + t)(1 + \lambda) - 1$. This prediction still does not depend on premia and can be tested given reasonable values of t and λ , as discussed in the next section.

Finally, if fixed costs vary between contracts, the condition may still apply, but under a different form. Namely, if

$$\pi(C_i) = P_i - \int R_i(L)dF_i(L) - \Gamma_i,$$

with $\Gamma_1 \neq \Gamma_2$, then NIP must be replaced by the following property: if C_2 covers more than C_1 , then $\pi(C_2) + \Gamma_2 \leq \pi(C_1) + \Gamma_1$. Under zero profit, it follows that $\Gamma_1 \geq \Gamma_2$; this implies, for instance, that a contract involving less coverage has a higher fixed cost. But this property may fail to hold. For example, a contract with higher coverage may also be more comprehensive,⁶ 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 2 becomes

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

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, but one may still estimate the left-hand side from the data and compare its value to a reasonable estimate of the right-hand side.

Thus, the main conclusion of this discussion is that whatever the cost function of the insurers, if we have enough information about it, then we can test some well-defined implication of asymmetric information (which may not look like a positive correlation property any more⁷).

□ **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. To keep the positive 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,⁸ the latter thus being independent of the distribution F . Under this assumption the utility function is determined up to an affine transformation:

There exist functions $v(W)$, a^θ , and c^θ 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 includes the standard models of pure adverse selection à la Stiglitz (1977) and pure moral hazard à la Arnott and Stiglitz (1988), as well as more complex frameworks.

⁶ 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).

⁷ Assume for instance that the payment $R(L)$ actually costs $\kappa R(L)$ to the insurer, where $\kappa > 1$, and that competition drives profits to zero. Then any contract offered at equilibrium must satisfy $P_i = \kappa \int R_i(L)dF_i(L)$. It is well known that in this case, the optimal contract entails only partial coverage (actually, a fixed deductible). With homogenous risk but different risk aversions, more risk-averse individuals will optimally choose a contract entailing more coverage (a smaller deductible). Now, let us introduce some infinitesimal difference in risk that is perfectly correlated with risk aversion; namely, assume that more risk-averse agents are slightly less risky. Then the correlation between risk and coverage is negative—but we can still test (7).

⁸ This is equivalent to assuming that risk aversion is publicly observable and hence is included in the observables X . Notice that when risk aversion depends on wealth, the assumption requires that wealth be observable—or at least that enough information about wealth may be inferred from observables.

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 . For any agent buying C_1 and choosing a distribution of claims G_1 under C_1 , we must have (after simplifying the a and c terms)

$$\int w_1(L)dG_1(L) \geq \int w_2(L)dG_1(L).$$

If agents have realistic expectations, by aggregating over the types buying C_1 , we find that

$$\int w_1(L)dF_1(L) \geq \int w_2(L)dF_1(L),$$

where F_1 is the empirical distribution of losses under C_1 . A similar argument applied on C_2 yields

$$\int w_2(L)dF_2(L) \geq \int w_1(L)dF_2(L).$$

Now subtract these inequalities to get

$$\int (w_2(L) - w_1(L))(dF_2(L) - dF_1(L)) \geq 0. \quad (8)$$

Notice that this result is valid whatever the contracts. If we assume moreover that C_2 covers more than C_1 , then it can be shown that $w_2 - w_1$ is negative, then positive, when L increases. Equation (8) then looks like a positive correlation property: in a weighted sense, F_2 puts more weight than F_1 on high losses and less weight on low losses. Nevertheless, the result involves the utility function v , which is unknown to the econometrician. Its implication is clearest for the case of two events, accident and no accident:

Corollary 3. Assume that risk aversion is public, expectations are realistic, and $L \in \{0, \bar{L}\}$. If two contracts C_1 and C_2 are bought in equilibrium, and C_2 covers more than C_1 , then $q_1 \leq q_2$.

This result was already known in the Rothschild and 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; then the assumption guarantees that agents who are *ex post* riskier indeed prefer contracts with higher coverage.

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 in 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.⁹

⁹ Of course, the situation just described is somewhat specific, because it relies on a strong, exogenous negative correlation between risk and coverage. Assuming, for instance, a positive correlation (risk-averse agents are riskier) would revalidate the positive correlation property. A more interesting approach would endogenize the correlation. For instance, Jullien, Salanié, and Salanié (2001) consider a model where risk aversion is the agent's private information and agents

5. An empirical test

■ Tests of the positive correlation between risk and coverage on insurance contracts have provided mixed results. Most articles on automobile insurance (see, e.g., Chiappori and Salanié, 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 and Snow (1994) was an early exception, but Dionne, Gouriéroux, and Vanasse (2001) attribute 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 and Poterba (2004); Bach (1998) reaches similar conclusions in her study of mortgage-related unemployment insurance contracts.

Since all of these articles rest on a simplified analysis of the insurance market, it is interesting to see whether the more general predictions we obtained in this article fare better when taken to the data. Note that the maintained assumptions of the theory are different for each of our results. Proposition 1 relies only on a revealed-preference argument, while Proposition 2 adds a nonincreasing profit condition and Corollary 3 assumes that risk aversion is public and losses are 0-1.

We use a dataset on car insurance that was given to us by an association of large French insurers. This dataset covers the full year 1989 and contains information on insuree characteristics, contract variables, and claims for nine insurance companies. The dataset we use in the application comprises 94,251 policies. About half (50,971) of these policies entail comprehensive coverage. These policies are straight-deductible contracts, with varying deductible levels. Table 1 describes the distribution of deductibles.¹⁰ In this dataset they are clustered in a small range (the first and the third quartile only differ by 300 francs, or about 50 dollars.) Thus we pool these contracts as contracts C_2 , setting the premium at the average level for given individual characteristics.¹¹ The other half of contracts are no-deductible, compulsory contracts that do not provide coverage for responsible claims, and we will use these as our contract C_1 .

We start by testing Proposition 1, which tells us that we should have

$$P_2 - P_1 \geq \int [R_2(L) - R_1(L)]dF_1(L).$$

For every claim in France, insurance experts assign responsibilities to the policyholders involved. A policyholder who 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, had only a compulsory coverage, and incurred some damage; there are 1,352 such claims in the data,¹² so that the average q is about .031. (For the sake of comparison, the third column of Table 1 gives the average value of q at various levels of deductibles for comprehensive contracts; the average value is .036).

Remember that our predictions are conditional on all variables X that are observed by the insurer (and hopefully by the econometrician). There is a large number of such variables in the data. As in Chiappori and 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 five 0-1 X variables:

- (i) Whether the policyholder has the best experience rating (a 50% bonus).

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 involve a violation of the positive correlation property.

¹⁰ One franc was about 16 U.S. cents in 1989.

¹¹ We still take into account the value of the deductible in the payments.

¹² Some policyholders have more than one responsible claim; we neglect these multiple occurrences in the following.

TABLE 1 Deductibles for Comprehensive Policies

Range of Deductibles (francs) %	Number of Policies	Average q	Average Claim
0–500	2,743	.034	6,347
501–1,000	34,309	.020	5,075
1,001–1,500	9,381	.042	6,016
1,501–2,000	2,403	.056	4,470
2,001–3,000	1,598	.066	7,761
3,001–10,000	535	.047	5,123

- (ii) Whether his/her car is relatively powerful.
- (iii) Whether his/her car is relatively expensive.
- (iv) Whether the driver is young or old.
- (v) Whether the car is driven in an urban area.

This defines $2^5 = 32$ cells. Each of them holds about 3,000 policyholders on average.

Within each cell we first estimate nonparametrically the distribution F_1 . A difficulty is that for compulsory contracts we observe the liability L' to third parties (which is covered in any case) but not the own damage L (which is not covered); hence our data do not allow us to directly test whether $q_2 \geq q_1$, as was done by Chiappori and Salanié (2000). Moreover, we only derived this theoretical prediction (in Corollaries 2 and 3) when there is only one loss level, while the cost for responsible claims is dispersed in our dataset: some claims are very costly and some very cheap. Thus the corollaries may not apply, and it is of interest to test the generalized positive correlation property, as given in Proposition 2. It is indeed possible that while the comprehensive contract has more claims, they are less costly than under the contract with no deductible. The last column of Table 1 gives some evidence on the size of the average claim under the comprehensive guarantee.

One solution for constructing F_1 is to assume that the distribution of L conditional on L' depends only on the individual's observed characteristics. We can then use the empirical distribution observed on contract C_2 for similar individuals to construct this conditional distribution. Combined with the observed distribution of L' under C_1 , we finally deduce F_1 .¹³

Within each cell X , we compute the quantity


$$T_1(X) = P_2 - P_1 - \int [R_2(L) - R_1(L)] dF_1(L)$$

and a Student statistic by dividing $T_1(X)$ by its standard error. This yields a collection of 32 numbers. Under the null hypothesis that $T_1(X) = 0$ for all X , these numbers should be distributed as a $N(0, 1)$ normal distribution. It turns out that all of these numbers are positive, which gives very strong corroborating evidence for Proposition 1 at any reasonable level of significance.

Comprehensive coverage premia were taxed at a rate of $t = 18\%$ in 1989. Moreover, we need to account for the cost structure of insurers. If we knew the value of the loading factor λ , we could directly test prediction (7). Since we do not know the precise value of λ , we ran the test for values of λ from 0 to 100%, which is a rather wide interval (a value of 20% is a reasonable guesstimate). For each of these values, we compute the Kolmogorov-Smirnov statistic for the null hypothesis that the Studentized statistics are distributed as $N(0, 1)$. We find that the value of the test statistic is

¹³ While this procedure may, under asymmetric information, generate a selection bias, we think the problem is likely to be minor. Note, first, that we condition not only on the agent's characteristics, but also on the size of the damage L' to third parties. If we assume, as seems natural, that severe accidents tend to be costly for all parties involved, then conditioning on L' eliminates most of the bias. Second, if some selection bias remains, it is (if anything) playing against us, since we tend to overestimate the loss probability of uninsured persons, making our test more stringent than it should be.

FIGURE 1
STUDENTIZED ESTIMATES OF $\int R_2(L)dF_2(L) - \int R_2(L)dF_1(L)$



high, with a p-value always below 10%, and reaching 1% at $\lambda = .20$. Figure 1 shows the estimated nonparametric density of these numbers (weighted by cell sizes), along with the $N(0, 1)$ density, for $\lambda = .20$. Our conclusion is that this test gives evidence for the positive correlation property: the null of zero correlation is rejected. By taking into account both the dispersion of claims and the cost structure of the insurer, we are able to corroborate the presence of asymmetric information.

We also investigate the NIP assumption. To test this, we compute profits

$$\pi_i(X) = \frac{P_i(X)}{1+t} - (1+\lambda) \int R_i(L)dF_i(L | X)$$

for $\lambda = .2$ and $t = .18$, and we test $\pi_1(X) \geq \pi_2(X)$. We find strong evidence to the contrary: the data suggest that in fact $\pi_1 < \pi_2$ in most cells and so the NIP assumption is rejected. This is an interesting finding that points toward the presence of imperfect competition in the French car insurance market.

One possibility is that in testing the NIP assumption, we erred in assuming that the cost structure of the insurer is well captured by a loading factor. Higher fixed costs per insuree or higher administrative costs per claim for the contract with higher coverage might account for this finding; still, an easy calculation shows that it would take a very large difference of 2,000 francs in fixed costs per insuree. As this is unlikely, a violation of NIP due to imperfect competition seems to be the most credible culprit.

Finally, let us consider experience rating. It is a legal feature of French car insurance, but so far we have neglected it in these empirical tests. It is easy to see that (provided both insuree and insurer have the same discount rates) introducing experience rating amounts to shifting the reimbursement function $R_i(L)$ by the discounted sum of increase in premia $c_i(L)$ in the objective functions of both insuree and insurer. Since insurees are free to switch contracts after a responsible claim, it seems reasonable to assume that $c_1 = c_2 = c$. Then Proposition 1 is unchanged, while the right-hand side of the inequality in Proposition 2 becomes

$$\int c(dF_2(L) - dF_1(L)),$$

which is positive if $q_2 \geq q_1$. It follows that taking experience rating into account would leave our evidence for Proposition 1 unchanged and weaken our evidence for Proposition 2. We can go a bit further. To the extent that insurees do not change contracts easily, there may be a case for assuming that $c_2 > c_1$ (as the experience rating system in France is a multiplicative bonus-malus system, so that a responsible claim brings a higher increase in premia for the more expensive

contract). It is easily seen that this introduces a term,

$$\int (c_1(L) - c_2(L))dF_1(L),$$

in the right-hand side of the inequality in Proposition 1. But since this new term is negative, this would strengthen our corroborating evidence for Proposition 1.

6. Conclusion

■ A first lesson stemming from this article 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 article provides (somewhat *a posteriori*) a theoretical foundation for many existing empirical articles, 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 2 is characteristic of a competitive setting. However, Proposition 2 must be reformulated with proportional loading or taxation, experience rating, or administrative costs of processing a claim.

Under imperfect competition, the nonincreasing profit assumption may 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 and Stiglitz (1976) or Chassagnon and Chiappori (1997)), then at least with a single claim size, 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.¹⁴

Empirically, most existing articles on automobile insurance do not reject the null of zero correlation. Our analysis shows that a positive correlation hypothesis is accepted when accounting for the dispersion of claims, although the magnitude of the adverse-selection effect remains small for most categories of insurees. More important, our findings—and in particular the strong rejection of the nonincreasing profit (NIP) property—also suggest that more attention should be devoted to the interaction between imperfect competition and adverse selection on risk aversion.¹⁵ In fact, our theoretical conclusions, together with recent results by Jullien, Salanié, and Salanié (2007), strongly indicate that there is a crying need for such models.

Appendix

■ **Proof of Lemma 2.** The proof is based on the following well-known result:

Lemma A1. Suppose that $W_2(L)$ and $W_1(L)$ are two nonincreasing functions of a random variable L , and that $W_2(L) - W_1(L)$ is nondecreasing with L . Then $X_1(L) = W_1(L) - EW_1(L)$ is a mean-preserving spread of $X_2(L) = W_2(L) - EW_2(L)$.

Indeed, denote by $H_i(x) = \text{Prob}(X_i(L) < x)$ the cumulative distribution function for $X_i(L)$. The assumptions in Lemma A1 imply that the difference $(H_2(x) - H_1(x))$ is positive then negative when x increases. Consequently, the function

$$D(X) = \int_{-\infty}^x [H_2(x) - H_1(x)]dx$$

¹⁴ 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.

¹⁵ Using subjective assessments by insurees, Finkelstein and McGarry (2006) find evidence for such a model in long-term care insurance.

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

$$D(+\infty) = E[X_2(L) - X_1(L)] = 0.$$

Thus D is nonnegative everywhere and in fact is positive as soon as $W_1(L)$ differs from $W_2(L)$ with positive probability. The result follows, by definition of a mean-preserving spread. *Q.E.D.*

We can now prove Lemma 1. We know that the agent prefers contract C_1 to contract C_2 , and thus he must prefer C_1 under G to C_2 under G . Let $W_i(L) = R_i(L) - L - P_i$ be the resulting wealth under contract C_i . Notice that under our assumptions on R_1 and R_2 , Lemma A1 applies. Since the agent dislikes mean-preserving spreads, it must be that the expectation of W_1 is above the expectation of W_2 . This yields the result. *Q.E.D.*

□ **On claims and losses.** First consider the case when the agent incurs some observable costs $c(L)$ when the loss L is realized. These costs may represent a monetary value for health, disease or uncovered damages, or some transaction costs incurred when reporting a loss; they may also represent the anticipated increase in premia in subsequent periods, due to experience rating. This amounts to replacing the indemnity $R_i(L)$ under contract C_i with $R'_i(L) = R_i(L) - c(L)$. Definition 1 trivially extends, so that (1) still holds, provided that all assumptions hold for R'_i instead of R_i .

Let us now show that Lemma 2 extends to the case when the agent may decide not to report a loss if the indemnity is below $c(L)$. In such a case¹⁶ the distribution of claims F may differ from the distribution of losses G : the latter is determined by the agent's precautionary behavior and characteristics, while the former may depend on each contract chosen. Suppose that the agent has chosen the distribution G and that he reports a loss only if he finds it advantageous, whatever the contract chosen. Everything is as if the insured was facing an indemnity

$$\hat{R}_i(L) = \max(R_i(L) - c(L), 0)$$

and would declare a loss only if $\hat{R}_i(L)$ is positive. Under a weak assumption,¹⁷ the property in Definition 1 is preserved by this transformation, and higher losses are still bad news. Then the proof of Lemma 2 still applies: we get

$$P_2 - P_1 \geq \int [\hat{R}_2(L) - \hat{R}_1(L)]dG(L).$$

Now recall that $R_2(L)$ is above $R_1(L)$. Denote A the set of losses L that are reported as claims under C_1 , and thus such that $R_1(L) > c(L)$. Observe that since $R_1(L) > c(L)$ implies $R_2(L) > c(L)$, the bracketed term in the above inequality is equal to $(R_2(L) - R_1(L))$ if $L \in A$, and is nonnegative otherwise. Therefore its integral is larger than

$$\int_A [R_2(L) - R_1(L)]dG(L).$$

Now G does not differ from the empirical distribution of losses F on A ; and for $L \notin A$, either $L > 0$ and then $dF(L) = 0$, or $L = 0$ and then $R_2(0) = R_1(0) = 0$. Overall we have shown that

$$P_2 - P_1 \geq \int_A [R_2(L) - R_1(L)]dG(L) = \int [R_2(L) - R_1(L)]dF(L)$$

and Lemma 2 extends to the case where the agent may choose not to report a loss.

References

- AKERLOF, G. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, Vol. 84 (1970), pp. 488–500.
- ARNOTT, R. AND STIGLITZ, J. "The Basic Analytics of Moral Hazard." *Scandinavian Journal of Economics*, Vol. 90 (1988), pp. 383–413.
- AUSUBEL, L. "Adverse Selection in the Credit Card Market." Mimeo, University of Maryland, 1999.
- BACH, K. *Negativauslese und Tarifdifferenzierung im Versicherungs-sektor*. Deutscher Universitäts-Verlag, Schesslitz, 1998.
- BRUGIAVINI, A. "Uncertainty Resolution and the Timing of Annuity Purchases." *Journal of Public Economics*, Vol. 50 (1993), pp. 31–62.

¹⁶ An obvious illustration is that of a contract with a deductible d_i , in which case the agent will not report losses smaller than d_i .

¹⁷ The slope of the function $R_2(L) - c(L)$ must lie between zero and one when its value is positive.

- CAWLEY, J. AND PHILIPSON, T. "An Empirical Examination of Information Barriers to Trade in Insurance." *American Economic Review*, Vol. 89 (1999), pp. 827–846.
- CHASSAGNON, A. AND CHIAPPORI, P.A. "Insurance Under Moral Hazard and Adverse Selection: The Case of Perfect Competition." Mimeo, DELTA, 1997.
- CHIAPPORI, P.A. "Econometric Models of Insurance Under Asymmetric Information." In G. Dionne, ed., *Handbook of Insurance*. New York: North-Holland, 2000.
- AND SALANIÉ, B. "Empirical Contract Theory: The Case of Insurance Data." *European Economic Review*, Vol. 41 (1997), pp. 943–950.
- AND ———. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy*, Vol. 108 (2000), pp. 56–78.
- DE MEZA, D. AND WEBB, D. "Advantageous Selection in Insurance Markets." *RAND Journal of Economics* Vol. 32 (2001), pp. 249–262.
- DIONNE, G., GOURIEROUX, C., AND VANASSE, C. "Testing for Evidence of Adverse Selection in the Automobile Insurance Market: A Comment." *Journal of Political Economy*, Vol. 109 (2001), pp. 444–453.
- FAGART, M.C. "Concurrence en contrats, anti-sélection et structures d'information." *Annales d'Economie et de Statistique*, Vol. 43 (1996), pp. 1–27.
- FINKELSTEIN, A. AND MCGARRY, K. "Private Information and Its Effect on Market Equilibrium: New Evidence from Long-term Care Insurance." *American Economic Review*, Vol. 96 (2006), pp. 938–958.
- AND POTERBA, J. "Adverse Selection in Insurance Markets: Policyholder Evidence from the UK Annuity Market." *Journal of Political Economy*, Vol. 112 (2004), pp. 183–208.
- JULLIEN, B., SALANIÉ, B., AND SALANIÉ, F. "Screening Risk-averse Agents Under Moral Hazard." *Economic Theory*, Vol. 30 (2007), pp. 151–169.
- KOUFOPOULOS, K. "Asymmetric Information, Heterogeneity in Risk Perceptions and Insurance: An Explanation to a Puzzle." Mimeo, University of Warwick, 2005.
- LANDSBERGER, M. AND MEILLISON, I. "A General Model of Insurance Under Adverse Selection." *Economic Theory*, Vol. 14 (1999), pp. 331–352.
- MIYAZAKI, H. "The Rat Race and Internal Labor Markets." *Bell Journal of Economics*, Vol. 8 (1977), pp. 394–418.
- PICARD, P. "Economic Analysis of Insurance Fraud." In G. Dionne, ed., *Handbook of Insurance*. New York: North-Holland, 2000.
- PUELZ, R. AND SNOW, A. "Evidence on Adverse Selection: Equilibrium Signalling and Cross-Subsidization in the Insurance Market." *Journal of Political Economy*, Vol. 102 (1994), pp. 236–257.
- ROTHSCHILD, M. AND STIGLITZ, J. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics*, Vol. 90 (1976), pp. 629–649.
- STIGLITZ, J. "Monopoly, Nonlinear Pricing, and Imperfect Information: The Insurance Market." *Review of Economic Studies*, Vol. 44 (1977), pp. 407–430.
- AND WEISS, A. "Credit Rationing in Markets with Imperfect Information." *American Economic Review*, Vol. 71 (1981), pp. 393–410.
- VILLENEUVE, B. "The Consequences for a Monopolistic Insurer of Evaluating Risk Better Than Customers: The Adverse Selection Hypothesis Reversed." *Geneva Papers on Risk and Insurance Theory*, Vol. 25 (2000), pp. 65–79.