

Abstract

The paper surveys a number of recent empirical studies that test for or evaluate the importance of asymmetric information in insurance relationships. I first discuss the main conclusions reached by insurance theory in a both a static and a dynamic framework. A particular emphasis is put on the testable consequences that can be derived from existing models. I review several studies exploiting these theoretical insights in a static context. Then I briefly consider the dynamic aspects.

Keywords : Insurance, adverse selection, moral hazard, contract theory, tests.

Econometric Models of Insurance under Asymmetric Information*

Pierre-André Chiappori[†]

February 2000

* To appear in the *Handbook of Insurance* (G. Dionne, ed.). The author is indebted to G. Dionne, I. Hendel, B. Salanié and two referees for useful suggestions and comments. Financial support from the Chaire d'Economie de l'Assurance (Paris) is gratefully acknowledged. Errors are mine.

[†] Department of Economics, The University of Chicago, 1126 east 59th Street, Chicago IL 60637. E-mail : pchiappo@midway.uchicago.edu

1 Introduction

Modern insurance economics have been deeply influenced by the developments of contract theory. Our understanding of several crucial aspects, such as the design of optimal insurance contracts, the form of competition on insurance markets or the role of public regulation, just to name a few, systematically refers to the basic concepts of contract theory - moral hazard, adverse selection, commitment, renegotiation and others. Conversely, it is fair to say that insurance has been, and to a large extent still remains, one of the most important and promising fields of empirical application for contract theory.

It can even be argued that, by their very nature, insurance data provide nearly ideal material for testing the predictions of contract theory. Chiappori (1994) and Chiappori and Salanié (1996) remark that most predictions of contract theory are expressed in terms of a relationship between the form of the contract, a 'performance' that characterizes the outcome of the relationship under consideration, and the resulting transfers between the parties. Under moral hazard, for instance, transfers will be positively correlated to but less volatile than outcomes, in order to conjugate incentives and risk sharing; under adverse selection, the informed party will typically be asked to choose a particular relationship between transfer and performance within a menu. The exact translation of the notions of 'performance' and 'transfer' varies with the particular field at stake. Depending on the particular context, the 'performance' may be a production or profit level, the performance of a given task or the occurrence of an accident; whereas the transfer can take the form of a wage, a dividend, an insurance premium and others.

In all cases, empirical estimation of the underlying theoretical model would ideally require a precise recording of (i) the contract, (ii) the information available to both parties, (iii) the performance, and (iv) the transfers. In addition, the contracts should be to a large extent standardized, and large samples should be considered, so that the usual tools of econometric analysis can apply. As it turns out, data of this kind are quite scarce. In some contexts, the contract is essentially implicit, and its detailed features are not observed by the econometrician. More frequently, contracts do not present a standardized form because of the complexity of the information needed either

to characterize the various (and possibly numerous) states of the world that should be considered, or to precisely describe available information¹. In many cases, part of the information at the parties' disposal is simply not observed by the econometrician, so that it is de facto impossible to condition on it as required by the theory. Last but not least, the 'performance' is often not recorded, and even not precisely defined. In the case of labor contracts, for instance, the employee's 'performance' is often the product of a supervisor's subjective estimation, that is typically not recorded in the firm's files (and, in any case, will typically not be available to the econometrician).

In contrast, most insurance contracts fulfill all of the previous requirements. Individual insurance contracts (automobile, housing, health, life, etc.) are largely standardized. The insurer's information is accessible, and can generally be summarized through a reasonably small number of quantitative or qualitative indicators. The 'performance' - whether it represents the occurrence of an accident, its cost, or some level of expenditure - is very precisely recorded in the firms' files. Finally, insurance companies frequently use data bases containing several millions of contracts, which is as close to asymptotic properties as one can probably be. It should thus be no surprise that empirical tests of adverse selection, moral hazard or repeated contract theory on insurance data have recently attracted renewed attention.

In what follows, I shall concentrate on empirical models that explicitly aim at testing for or evaluating the importance of asymmetric information in insurance relationships. This obviously excludes huge parts of the empirical literature on insurance, that are covered by other chapters of this volume. Also, I will leave aside the important literature on fraud - a topic that is explicitly addressed by the contribution of R. Derrig. Similarly, since the undoubtedly major field of health insurance is comprehensively surveyed by M. Pauly, I shall only allude to a few studies relating to information asymmetries in this context.

The structure of this contribution is as follows. Section 2 discusses the main conclusions reached by insurance theory in a both a static and a dynamic framework. A particular emphasis is put on the testable consequences that can be derived from existing models. Section 3 reviews several studies exploiting these theoretical insights in a static context. Section 4 briefly considers the dynamic aspects of the issue. Finally, concluding comments are in

¹This problem, for instance, is frequently encountered with data related to firms' behavior.

the last section.

2 Empirical tests of information asymmetries: the theoretical background

It is by now customary to outline two polar cases of asymmetric information, namely adverse selection and moral hazard. Each case exhibits specific features that must be understood before any attempt at quantifying their empirical importance².

2.1 Adverse selection³

2.1.1 Definition

Adverse selection arises when one party - generally, the subscriber - has a better information than the other party - the insurer - about some parameters that are relevant for the relationship. In most theoretical models, the asymmetry is relative to the level of risk : the client is assumed to know better either her accident probability, or the (conditional) distribution of losses incurred in case of accident, or both. A key feature is that, in such cases, the agent's informational advantage is directly related to the insurer's (expected) cost of providing the contract.

While theoretical models concentrate upon one particular source of adverse selection - the agent's better knowledge of her risk - the empirical relevance of this exclusive emphasis is not always guaranteed. In many real-life applications, risk is not the only possible source of informational asymmetry, and arguably not the most important one. For instance, individuals also have a better knowledge of their own preferences, and particularly their level of risk aversion - an aspect that is often disregarded in theoretical models.

A possible justification for this lack of interest is that, in principle, adverse selection on preferences, per se, has negligible consequences upon the form and the outcome of the relationship, at least in a competitive context. Pure competition typically imposes that companies charge a fair premium, at least whenever the latter can be directly computed (which is precisely the

²For a clear and comprehensive presentation of the various theoretical models, the reader is referred to Salanié (1997).

³See also the contribution by Dionne and Doherty in this volume

problem when the agent's risk is not known). Hence, the equilibrium contract should not depend on the subscriber's preferences, whether the latter are public or private information. To be a little more specific: in a model of perfectly competitive insurance markets with symmetric information, the introduction of hidden information on preferences only will not alter the equilibrium outcome.

This conclusion should however be qualified, for at least two reasons. For one thing, perfect competition is a natural assumption within a simplified theoretical model, but much less so in reality. Fixed costs, product differentiation, price stickiness, switching costs and cross-subsidization are part of the real world; oligopoly is probably the rule rather than the exception. In such a context, firms are able to make positive profit, and the latter is related to the agents' demand elasticity - which directly reflects risk aversion. To take an extreme case, it is well known that in a principal-agent framework - equivalent to some monopoly position of the insurance company - adverse selection on risk aversion does matter for the form of the optimal contract.

A second caveat is that even when adverse selection on preferences *alone* does not matter, it may still, when added to asymmetric information of a more standard form, considerably modify the properties of equilibria. In a standard Rothschild-Stiglitz context, for instance, heterogeneity in risk aversion may result in violations of the classical single-crossing property of indifference curves 'a la Spence-Mirrlees', which in turn generates new types of competitive equilibria⁴. More generally, situations of bi- or multi-dimensional adverse selection are much more complex than the standard ones, and may require more sophisticated policies⁵.

2.1.2 What does the theory predict ?

The previous remarks only illustrate a basic conclusion : when it comes to empirical testing, one should carefully check the robustness of the conclusions under consideration to various natural extensions of the theoretical

⁴See Villeneuve (1996) or Chassagnon (1996). The same remark applies to models with adverse selection and moral hazard, whether adverse selection is relative to risk, as in Chassagnon and Chiappori (1997), or to risk aversion, as in Julien, Salanié and Salanié (1996).

⁵Typically, they may require more instrument than in the standard models. In addition, one may have to introduce randomized contracts, and bunching may take specific forms. See Rochet and Choné (1998)

background. Now, what are the main robust predictions emerging from the theoretical models?

A first distinction, at this stage, must be made between exclusive and non exclusive contracts. The issue, here, is whether individuals are free to buy an arbitrary number of contracts from different insurance companies, or whether the insurer can impose an exclusive relationship. In the field of insurance, both situations coexist; for instance, automobile insurance contracts are almost always exclusive, whereas annuities or life insurance are typically sold without exclusivity.

Non exclusive contracts and price competition Non exclusivity strongly restricts the set of possible contracts. For instance, no convex price schedule can be implemented : if unit prices rise with quantities (which is typically what adverse selection requires), agents can always 'linearize' the schedule by buying a large number of small contracts from different insurers⁶. The same holds true for quantity constraints, which can be considered as a particular form of price convexity. To a large extent, the market will in that case entail standard (linear) price competition.

In this context, adverse selection has well-known consequences. Since all agents face the same (unit) price, high risk individuals are de facto subsidized (with respect to fair pricing), whereas low risk agents are taxed. The latter are likely to buy less insurance, or even to leave the market. A first prediction of the theory is precisely that, in the presence of adverse selection, the market typically shrinks, and the high risk agents are over-represented among buyers. In addition, purchased quantities should be positively correlated with risk; i.e., high risk agents should, everything equal, buy more insurance. Both predictions are testable using insurers' data.

Finally, the presence of adverse selection will have an impact on prices. Because of the over-representation (in number and in quantities) of high risk agents in the insurers' portfolios, unit prices will, at equilibrium, exceed the level that would obtain in the absence of adverse selection. Although the latter is not observable, it may in some cases be computed from the average characteristics of the general population. A typical example is provided by annuities, since the distribution of life expectancy conditional on age is well documented. It is in principle possible to compute the fair price of a given

⁶The benefits of linearization can be mitigated by the presence of fixed contracting costs. For large amounts of coverage, however, this limitation is likely to be negligible.

annuity, and to compare it to actual market price. A difference that exceeds the 'normal' loading can be considered as indirect evidence of adverse selection (provided, of course, that the 'normal' level of loading can be precisely defined).

Exclusive contracts In the alternative situation of exclusivity, theoretical predictions depend, among other things, on the particular definition of an equilibrium that is adopted - an issue on which it is fair to say that no general agreement has been reached. Using Rothschild and Stiglitz's concept, equilibrium may fail to exist, and cannot be pooling. However, an equilibrium a la Riley always exists. The same property holds for equilibria a la Wilson; in addition, the latter can be pooling or separating, depending on the parameters. Referring to more complex settings - for instance, game-theoretic frameworks with several stages - does not simplify the problem, because the properties of equilibria are extremely sensitive to the detailed structure of the game (for instance, the exact timing of the moves, the exact strategy spaces, ...), as emphasized by Hellwig (1987).

These remarks again suggest that empirically testing the predictions coming from the theory is a delicate exercise; it is important to select properties that can be expected to hold in very general settings. Following Chiappori and Salanié (2000), one can argue that three conclusions seem fairly robust, namely :

1. under adverse selection, agents are likely to be faced with menus of contracts, among which they are free to choose
2. contracts with more comprehensive coverage are sold at a higher (unit) premium; or, more precisely, the marginal premium rate (i.e., the increase in premium required for each additional dollar of coverage) should increase with coverage (*convex pricing*).
3. contracts with more comprehensive coverage are chosen by agents with higher expected accident costs.

The first prediction is essentially qualitative. It is also quite general, and holds true for different types of adverse selection (i.e., agents may differ by their risk, but also by their wealth, preferences, risk aversion, etc.), at least within an imperfect competition context. The second prediction, in most circumstances, is a direct consequence of individual rationality in general : even

in the absence of adverse selection, an agent will not choose a contract with higher deductible (or more coinsurance) unless its unitary price is lower, at least if pricing is approximately fair⁷. Testing for this property is an interesting perspective; however, it requires an explicit and adequate estimation of the firm's pricing policy. Such a task may, in practice, reveal quite difficult. While abstract models generally assume proportional costs, in real life fixed costs and (dis)economies of scale and scope play an important role, and are responsible for many non linearities in the pricing policy. The latter may be very difficult to disentangle from those due to asymmetric information. If, because of adverse selection, the price schedule is less concave than it would otherwise be, there is little hope that such a subtle impact can be empirically identified at all.

On the contrary, testing the third property does not require an estimation of the firm's pricing policy. If agents, facing an identical menu of contracts (sold at identical fares), self select on the basis of some private information they have about their riskiness, then a positive correlation between coverage and expected costs should be observed, whatever the prices. It should be noted that this prediction seems quite robust. For instance, it does not require single crossing, and it holds when moral hazard or multidimensional adverse selection are introduced.

This claim must however be qualified, or at least clarified. What must be stressed, at this point, is that this prediction is valid within a group of *observationally identical* agents. In practice, insurance companies use observable characteristics to categorize individual risks. As far as pricing *across* classes thus constructed is concerned, the previous conclusions are totally irrelevant. Some agents may be offered contracts entailing both higher unitary premium and larger deductible⁸; the point being that they cannot choose the class they will be categorized into. The self-selection issue applies only *within* such classes. The empirical translation is that one must systematically consider probability distributions that are *conditional on all observables*. Although this requirement is in principle straightforward, how this conditioning is actually performed on 'real' data is one of the key problems of this line of empirical investigation.

Finally, although most models predict a positive correlation between e-

⁷This needs not be true when loading is high, because agents with lower risk will typically prefer less coverage, even at a (slightly) higher unit price. However, insurance companies are unlikely to charge a higher unit price to less risky customers in any case.

⁸This is typically the case of automobile insurance for young drivers, for instance.

equilibrium coverage and risk, the literature has identified a few cases in which this correlation can be reversed. One is when the informational advantage is on the insurer's side, as studied by Villeneuve (1996). The problem, here, is that the insurer's claim that a client's risk is high can only be credible if it does not result in increased profit for the insurer. This revelation constraint, in general, requires partial coverage of risky agents and full coverage of safer ones. Another case is the 'cherry-picking' model of de Meza and Webb (1999). In this context, risk averse agents are both more willing to buy insurance and more likely to adopt a cautious behavior that results in smaller accident probability. The authors show that this combination of adverse selection on risk aversion and moral hazard may, in the presence of loading, generate a negative correlation between risk and coverage at equilibrium.

2.2 Moral hazard

Moral hazard occurs when the probability of a claim is not exogenous, but depends on some decision made by the subscriber (e.g., effort of prevention). When the latter is observable and contractible, then the optimal decision will be an explicit part of the contractual agreement. For instance, an insurance contract covering a fire peril may impose some minimal level of firefighting capability, or at least adjust the rate according to the existing devices. When, on the contrary, the decision is not observable, or not verifiable, then one has to examine the incentives the subscriber is facing. The curse of insurance contracts is that their mere existence tends to decrease incentives to reduce risk. In the extreme case of complete insurance (when the insured's welfare does not depend at all on the occurrence of an accident), incentives are killed, resulting in maximum accident probabilities. In general, different contracts provide different incentives, hence result in different observed accident rates. This is the bottom line of most empirical tests of moral hazard.

An additional distinction that is specific to insurance economics is between an *accident* and a *claim*. Moral hazard is ex ante when the consequence of the agent's effort is a decrease in *accident* probability or severity (this is typical of prevention). But insurance companies are interested in claims, not in accidents. Whether an accident results in a claim is the agent's decision, and is as such influenced by the form of the insurance contract. Of course, the previous argument holds for both case : more comprehensive coverage discourages accident prevention *and* increases incentives to file a claim for small accidents. However, the econometrician will in general be

eager to distinguish between 'true' moral hazard (resulting in changes in the accident rates) and incentives to filing a claim, if only because the welfare implications are typically very different. For instance, a deductible is more likely to be welfare increasing when it reduces accident probability than when its only effect is to discourage victims from filing a claim (unless, of course, the technology of the insurance industry - say, the presence of fixed costs - makes the processing of small claims very inefficient⁹).

Quite interestingly, the basic moral hazard story is very close to the adverse selection one, except for an inverted causality. Under adverse selection, people are characterized by different levels of risk (that will ex post be translated into different accident rates); because of these discrepancies, they choose different contracts. In a context of moral hazard, agents first choose different contracts; then they are faced with different incentive schemes, hence adopt more or less cautious behavior, which ultimately results in heterogeneous accident probabilities. In both cases, however, the conclusion is that, controlling for observables, the choice of a contract will be correlated with the accident probability - again, more comprehensive coverage being associated to higher risk.

This suggests that it may be difficult to distinguish between adverse selection and moral hazard in the static framework (i.e., using cross-sectional data). An econometrician may find out that, conditionally on observables, agents covered by a comprehensive automobile insurance contract are more likely to have an accident. Deciding whether they chose full coverage because they knew their risk was higher, or they became more risky because the comprehensive contract they chose (for some exogenous and independent reason) killed most incentives to drive safely, is a much harder task.

Distinguishing adverse selection from moral hazard The adverse selection versus moral hazard puzzle can be solved in different ways. One is to use natural experiments. Assume that the incentive structure that a given population faces is modified for exogenous reasons (typically, a reform of the regulatory framework). Resulting changes in people's behavior (if any) cannot be attributed to adverse selection (since the population remains

⁹A related problem is fraud, defined as any situation where a subscriber files a claim for a false accident or overstates its severity in order to obtain a more generous compensation. The optimal contract, in that case, typically require selective auditing procedures (see the Chapter by R. Derrig in this volume)

unchanged); then moral hazard is a natural interpretation. Ideally, the population would be randomly split into various subgroups, one of which (at least) keeps the old scheme. Such a 'reference' group allows to sort out the changes that are specifically driven by the new regulation. The celebrated Rand study on medical expenditures (see Manning et al. 1987) provides a perfect illustration of such a context.

This basic idea may in some occasions apply even in the absence of an actual 'experiment' of this kind. Any context where similar individuals are facing different incentive schemes can do, provided one can be sure that the selection into the various schemes is not related to risk-relevant characteristics. A few examples of such 'pseudo natural experiments' will be provided later on.

Finally, and perhaps more fruitfully, the distinction between adverse selection and moral hazard may exploit the different dynamic properties of the various cases. Two approaches can be chosen here. One is to assume that existing contracts are optimal. The recent developments of dynamic contract theory can then be used to characterize the main qualitative patterns these optimal contracts should exhibit in each case. Testable predictions obtain, that can be compared to existing contracts. The second approach does not rely on optimality assumptions. Rather, it takes existing contracts as given, and characterizes the agent's optimal response in the various contexts. For instance, one may test for moral hazard by deriving the dynamic properties of the optimal experience rating schemes, and check whether existing schemes correspond to this model. Alternatively, one may take the existing experience rating mechanisms as given, and study the induced dynamics of effort and accidents. Then these can be compared to observed behavior. Note that first approach only requires data on contracts, whereas the second, while more general (since it does not rely on optimality assumptions), is more demanding and necessitates individual data on contracts and accidents.

2.3 The dynamics of optimal contracts

In repeated interactions, the transfers at any period (and in particular the premium paid by the agent) will typically be contingent on the past history of the relationship. While this notion of 'experience rating' is general, the particular form it takes in optimal contracts will in general depend on the type of asymmetry and on the parties' ability to commit.

Repeated adverse selection As a first polar case, consider the case where the basic adverse selection model is repeated. Cooper and Hayes (1987) provide a full characterization of the optimal contract under the assumption that both the insurer and the insured can fully commit on their future behavior. While high risk agents are fully covered by a time-invariant contract, the contract for low-risk individuals entails both partial coverage at each period and experience rating¹⁰. The empirical relevance of this result is however debatable since the corresponding contract is not immune to renegotiation. The agents' type is indeed fully revealed at the first period through her choice of a contract. At any subsequent period, both parties can then increase their (ex post) well-being by agreeing on a different continuation that provides full insurance to low-risk agents as well. If a renegotiation of this type cannot be prevented, however, it will be anticipated by low risk agents at the first period, which aggravates the first-period revelation problem. Hart and Tirole (1988) and Laffont and Tirole (1993) characterize the optimal long term contract under full commitment when renegotiation cannot be ruled out. They show that it entails partial pooling. At any period, two contracts are available, one of which only entails full coverage at a high unit price. Low-risk agents always prefer the partial coverage contract, whereas high-risk agents are indifferent and are randomly distributed across the two contracts. In particular, while the choice of full coverage signals a high risk type with probability one, the partial coverage contract always attracts agents of both types (except for the last period), which avoids the renegotiation problem. This intuition is extended to a competitive insurance framework by Dionne and Doherty (1994). Their model entails one-sided commitment : the insurer can commit to keep its clients, whereas a client is always free to leave the relationship. They show that, in a two-period framework, the optimal renegotiation-proof contract entails semi-pooling in the first period and separation in the second. More importantly, the contract will exhibit 'highballing' features; i.e., the insurance company will typically make positive profits in the first period, compensated by low, below-cost second period prices. An

¹⁰This property apparently contradicts a standard result of the repeated adverse selection literature, stating that the optimal long term contract under full commitment is the repetition of the one-period optimal contract. The difference is due to a particular feature of insurance models : the occurrence of an accident during the relationship provides information about the agent's type *independently of the agent's strategy*. It is then possible to increase efficiency by signing contracts that are contingent on this information, which is the exact definition of experience rating.

important remark is that this property is empirically testable. Moreover, it does not necessarily require individual data, since the highballing prediction can be tested at the firm level¹¹.

Repeated moral hazard Repeated moral hazard constitutes a second natural polar case. A first intuition is that the law of large number should, in this case, be of considerable help. Assuming that the agent's effort strategy is stationary, a 'long enough' observation of accident rates should allow the principal to very precisely infer the action chosen. Then an adequate punishment scheme should lead to outcomes close to the first-best. Rubinstein and Yaari (1983) show indeed that when neither the principal nor the agent has a preference for the present, the first-best can be implemented, at least when the interaction is infinitely repeated.

From an empirical perspective, however, the Rubinstein and Yaari model suffers from several limitations. It requires very patient agents and an infinitely repeated interaction, two assumptions that may seem at least debatable. Another major problem that any repeated contract model should face, at least when empirical applications are at stake, is how to model the agent's access to financial markets. Most theoretical models assume that the payments at each period directly determine the agent's consumption, i.e. that the agent can neither save nor borrow to improve the allocation of her consumption over time¹². Such an assumption has to be considered with suspicion. One can probably accept that some agents may face credit constraints (although the actual importance of these effects seems actually small for a majority of people). In the repeated moral hazard context, however, Rogerson (1985) has shown that the optimal contract with no access to financial markets is such that the agent would actually like to *save more* than what is implied by the contract. Hence ignoring financial markets amounts to assuming that agents cannot freely save, a very strange assumption indeed.

Allowing for the agent's access to financial markets, however, raises other problems. In general, the agent's transactions on financial markets are not observable by the insurer. As a consequence, whenever the risk aversion of the agent depends on her wealth (i.e., her preferences are not CARA), the

¹¹Finally, the no commitment case is particularly complex in the competitive context, in particular because it generates difficult existence problems. See Fombaron (1998) for a recent characterization.

¹²In other word, the insurance contract is used to smooth the agent's consumption not only across states of the world, but also across periods.

second-period relationship entails an additional element of adverse selection on preferences, that is moreover endogenous to the first period strategies. Chiappori et al. (1994) show that this fact dramatically reduces the set of available contracts. They prove, in particular, that only contracts entailing the minimum effort level for all periods but the very first are renegotiation-proof, at least among the set of non randomized contracts. The particular case of CARA preferences and monetary cost of effort, on the other hand, had been previously studied by Fudenberg, Holmstrom and Milgrom (1990). The absence of income effect due to the constant absolute risk aversion assumption turns out to greatly simplify the problem. Fudenberg, Holmstrom and Milgrom show that the (second-best) optimal contract can then be implemented without memory, i.e., through deductible only (and without experience rating). The empirical scope of this result is however limited by the restrictiveness of the two key assumptions.

Symmetric learning A third, and possibly more promising case is when information is symmetric, in the sense that both the customer and the insurance company initially ignore the customer's risk. At each period of the relationship, they learn from the outcome (say, the number of accidents during the period), and symmetrically revise their prior. Assuming that the firm is able to commit whereas the client is not, Harris and Holmstrom (1982) show that in the optimal contract the premium is upwards rigid¹³. Unlucky agents are not penalized for the occurrence of an accident, in the sense that their premium does not increase. On the contrary, lucky agents are rewarded for the absence of accident during the period (this is a direct consequence of the absence of commitment on the agent's side). Note that the existence of this 'bonus' is costly ex ante, because it restricts the agent's insurance coverage. Under bilateral commitment, the optimal contract would not depend on the occurrence of an accident, thus providing insurance against the classification risk as well (i.e., the risk of being believed to be a high risk)¹⁴.

¹³The initial Harris and Holmstrom model consider wage dynamics. The transposition to insurance is straightforward.

¹⁴In principle, the upward rigidity prediction only holds without agent's access to financial markets. Indeed, the first best could otherwise be implemented. The idea is to charge a very high first period premium, then choose at any subsequent period a uniform premium that leaves the agents with the best history just indifferent between staying or leaving (all other agents strictly prefer to stay). Such a contract, however, requires that the agent borrows a large amount at the beginning of the relationship; it cannot be implemented in

These features can be contrasted with the contracts that obtain without commitment from the insurer. Then the occurrence of an accident at period t always worsens the insurer's posterior about the agent's type and results in higher premia in period $t + 1$. Since both types of contract coexist in practice, these qualitative properties can be empirically distinguished, a fact that has been used in the empirical literature¹⁵. Finally, the Harris and Holmstrom model has been recently extended by de Garidel (1997), who analyzes the issue of information transmission between the incumbent insurer and its potential competitors.

Dynamic behavior under (possibly) suboptimal contracts Even when contracts are suboptimal, the dynamics of outcomes (say, accident occurrence) will typically differ under moral hazard and adverse selection. Adverse selection (and more generally any heterogeneity that is residual to the insurance company's classification model) typically results in 'positive contagion' phenomena : more accidents in the past signal a high risk agent, who is more likely to have other accidents in the future. On the contrary, under moral hazard, the experience rating schemes are typically such that the marginal cost of a new accident increases with the number of past accidents. This point is formally established by Chiappori and Heckman (2000), who consider an experience rating scheme for automobile insurance where the premium is multiplied by some constant larger (resp. smaller) than one for each year with (resp. without) accident (the so-called 'bonus/malus' scheme). They show that, under general assumptions, effort at each period is an increasing function of the premium level at the beginning of the period¹⁶. In other words, more accidents in the past result in stronger incentives, more prevention effort and lower accident rates in the future. This 'negative contagion' property can be tested, provided individual, dynamic data are available.

the presence of credit constraints.

¹⁵Interestingly enough, this model also generates strong testable predictions on the dynamics of transfers between the principal and the agent, even when the agent's 'performance' (here, the occurrence and the severity of an accident) is not observable. This property is not utterly appealing in the insurance context, where accidents are typically observed. But it may become quite attractive in different contexts, such as labor contracts. See Chiappori, Salanié and Valentin (1999) for a detailed investigation of the so-called 'late beginner' effect.

¹⁶In practice, the 'bonus/malus' coefficient is often capped. Then the effort monotonicity property obtains only for values of the bonus coefficient that are 'far enough' from the cap.

2.4 Claims versus accidents

As mentioned above, a key feature of insurance data is that the insurer can only observe claims, not accidents. In most cases, the decision to file a claim is made by the subscriber, and must be understood as a response to specific incentives. Should the costs of filing a claim exceed the expected benefits - say, because the expected cost is below the deductible, or experience rating implies that the claim will result in higher future premia - then the insured is always free not to report the accident.

This simple remark has two consequences. One is that the incentives to file a claim should be monitored by the insurance company, particularly when the processing of a small claim involves important fixed costs for the company. A deductible, for instance, is often seen by insurance companies as a simple and efficient way of avoiding small claims. More related to the present topic is the fact that the empirical distribution of claims will in general be a truncation of that of accidents - since 'small' accidents are typically not declared. Moreover, the truncation is endogenous; it depends on the contract (typically, on the size of the deductible or the presence of experience rating), and also on the individual characteristics of the insured (if only because the cost of higher future premia is related to the expected frequency of future accidents). This can potentially generate severe biases. If a high deductible discourages small claims, a (spurious) correlation will appear between the choice of the contract and the the number of filed claims, even in the absence of adverse selection or ex ante moral hazard. The obvious conclusion is that any empirical estimation must very carefully control for potential biases due to the distinction between accidents and claims.

3 Empirical estimations of asymmetric information in the static framework

While the theoretical analysis of contracts under asymmetric information began in the 70s, the empirical estimation of insurance models entailing either adverse selection or moral hazard is more recent. Among early contributions, one may mention Boyer and Dionne (1987) and Dahlby (1983), who does not reject the presence of some asymmetric information. However, Dahlby uses aggregate data only, so that it is not clear whether his results would be robust to the inclusion of more detailed individual data.

3.1 Non exclusivity and price competition

3.1.1 Annuities

In a non exclusivity context, several studies have been devoted to the market for annuities. The latter provide a typical example of non exclusive contracts. Also, the information available to the insurance company is generally rather sparse. Despite the similarities between annuities and life insurances (in both cases, the underlying risk is related to mortality), it is striking to remark that while life insurance contracts (at least above some minimum amount) are contingent to detailed information upon the subscriber's health state, the price of an annuity only depends on the buyer's age. This suggests that adverse selection may play an important role in this context.

A first line of research has concentrated upon prices. In an important contribution, Friedman and Warshawski (1990) compute the difference between the implicit contingent yield on annuities and the available yield on alternative forms of wealth holding (in that case, US government bonds). Even when using longevity data compiled from company files, they find the yield of annuities to be about 3% lower than US bonds of comparable maturity, which they interpret as evidence of adverse selection in the company's portfolio. Similar calculations on UK data by Brugiavini (1990) find a 3% difference, but only when longevity is estimated on the general population.

A related but more direct approach studies the distribution of mortality rates in the subpopulation of subscribers, and compares it to available data on the total population in the country under consideration. Brugiavini (1990) documents the differences in life expectancy between the general population and the subpopulation actually purchasing annuities.. For instance, the probability, at age 55, to survive till age 80 is 25% in the general population but close to 40% among subscribers. In a similar way, the yield difference computed by Friedman and Warshawski (1990) is 2% larger when computed from data relative to the general population.

A final prediction of the theory is that the amount purchased should be positively correlated with (realized) longevity. However, neither paper can test for this property.

3.1.2 Life insurance

Life insurance contracts provide another typical example of non exclusive contracts, although adverse selection might in this case be less prevailing.

In a recent paper, Cawley and Philipson (1997) use direct evidence on the (self-perceived and actual) mortality risk of individuals, as well as the price and quantity of their life insurance. They first find that unit prices fall with quantities, indicative of the presence of bulk discounts. More surprising is the result that quantities purchased appear to be *negatively* correlated with risk, even when controlling for wealth. This strongly suggests that the market for life insurance may not be affected by adverse selection, probably because of the large amount of information available to the insurer in that case¹⁷.

3.2 Exclusive contracts

3.2.1 The hedonic approach (Puelz and Snow 1994)

The alternative approach - i.e., the analysis of competition in exclusive contracts - has attracted renewed attention during the last decade. An important contribution, due to Puelz and Snow (1994), relies upon a hedonic model of insurance pricing. Using individual data from an automobile insurer in Georgia, they build a two-equation model of insurance contracts. The first equation represents the pricing policy adopted by the insurance firm. It takes the form :

$$P_i = g(D_i, X_i, \varepsilon_i)$$

where P_i and D_i are the premium and the deductible in the contract chosen by individual i , the X_i are individual-specific exogenous variables and ε_i is an econometric error term. This allows to directly test our second prediction - namely, that higher premia should be associated to lower deductible. This property is indeed confirmed by the data. However, as argued above, this result, per se, cannot provide a strong support to the existence of adverse selection. Whatever the reason for offering a menu of contracts, one hardly expects rational agents to choose contracts with a higher unitary premium *and* a larger deductible. More interesting is the test they propose for the third prediction - i.e., that the choice of a contract offering a more comprehensive coverage should be correlated with a higher accident probability. For this purpose, they estimate a second equation that describes the agent's choice

¹⁷As argued above, the negative correlation finding can be explained in two different ways. The 'cherry-picking' effect requires moral hazard to play an important role; it may be the case, for instance, that 'timid' agents buy more insurance and invest more in medical prevention. The 'informed principal' argument, on the other hand, is less convincing in this case, since in principle it requires exclusive contracts.

of deductible. The decision depends on the agent's "price of deductible" \hat{g}_D , as estimated from a third regression using instrumental variables, and on his (unobserved) accident probability. Since the latter is unobservable, it is proxied by a dummy variable RT_i that equals one if the individual had an accident and zero otherwise. This leads to an equation of the form :

$$D_i = h(\hat{g}_{Di}, RT_i, X_i, \eta_i)$$

where η_i is another error term. The Rothschild-Stiglitz model predicts that higher risks buy better coverage, i.e. a lower deductible, so that h should decrease in RT . Puelz and Snow specify their pricing equation as a linear model and estimate it by ordinary least squares. Since there are only three levels of deductible in their data set, they estimate the contract choice equation (again linear) by ordered logit; they find a negative coefficient for RT_i (although the choice of deductible does not vary much with the risk type).

3.2.2 Problems with the hedonic approach

There are several problems in the Puelz-Snow approach, that provide an interesting illustration of the difficulties encountered by any attempt at testing the predictions of contract theory. A first (and somewhat technical) one is related to the approximation of the (unknown) accident probability by the dummy variable RT . This procedure introduces a measurement error in the second equation. In linear models, the estimates would be biased towards zero, which would reinforce the conclusion of Puelz-Snow. In an ordered logit, it is not clear which way the bias goes.

A second concern is that the data set under consideration comprises individuals of various ages and driving records. This important heterogeneity may be troublesome for two reasons. One is heteroskedasticity. Presumably, the distribution of the random shocks, and especially of η_i , will depend on the driver's seniority. Within a nonlinear model such as the ordered logit, this will bias the estimation. The second and more disturbing problem relates to experience rating. Insurers typically observe past driving records; these are highly informative on probabilities of accident, and, as such, are used for pricing. Omitting these variables will typically generate a bias, that tends precisely to overestimate the level of adverse selection : the corresponding information is treated by the econometrician as being private, whereas it is in fact common to both parties. However, the introduction of past experience is

a quite delicate task, because it is (obviously) endogenous. Not only are panel data required, but endogeneity then raises specific econometric problems that will be discussed in the next section. Thirdly, nothing is done in the paper to distinguish between ex ante and ex post moral hazard. Higher deductibles tend, everything equal, to discourage accident reporting, which has little to do with accident prevention but can generate spurious correlation.

A final (and quite general) problem relates to the use of a highly constrained functional form. In the second equation, in particular, the relationship of the latent variable to the accident probability π and the price \hat{g}_D is taken to be linear. This needs not be the case. To illustrate this point, Chiappori and Salanié (2000) consider the case of constant absolute risk aversion. Then the individual's choice of deductible is of the form :

$$D_i = \frac{1}{\sigma_i} \log \frac{1 - \pi_i - \hat{g}_{D_i}}{\pi_i (1 + \hat{g}_{D_i})}$$

which is highly nonlinear. They argue that, in fact, applying the Puelz-Snow procedure to data generated by a *symmetric* information model, according to this formula, may well result in the kind of negative estimates they get, simply because the accident term captures in fact some of the omitted nonlinearities.

A particularly elegant illustration of this fact is provided by Dionne, Gouriéroux and Vanasse (1998). Their idea is to first run a probit on the “accident” variable, then to introduce the resulting *predictors* $\hat{\pi}_i$ of this probit in the right-hand side of the second equation (for the choice of deductible), together with the dummy RT_i . They find that the $\hat{\pi}$ variable has a large and highly significant negative coefficient, whereas the RT variable is no longer significant. This, obviously, has nothing to do with adverse selection, as $\hat{\pi}_i$ is by construction a function of *observed* variables only; it suggests, a contrario, that the negative influence of RT in the initial model can be spurious and due to misspecification.

3.2.3 Avoiding possible misspecifications

Several studies have attempted to correct for these biases. Chiappori (1994) and Chiappori and Salanié (2000) propose a very general approach, that may potentially apply to most problems entailing adverse selection. The idea is to simultaneously estimate two (non linear) equations. One relates to the choice of the deductible. In the (simplest) case of a binomial decision, it takes the form

$$y_i = \mathbb{I}[f(X_i, \beta) + \varepsilon_i > 0] \tag{1}$$

where, as above, the X_i are individual-specific exogenous variables, the β are parameters to be estimated, and ε_i is an econometric error term. Note that, contrarily to Puelz and Snow, the accident variable RT is *not* included in the right hand side. Nor is the premium; the idea, here, is that the latter is computed as a function of observables only, so that any information it conveys is already included in $f(X_i, \beta)$ - provided, of course, that the corresponding functional form is flexible enough.

The second equation takes the occurrence (and/or severity) of an accident as the dependent variable. In the simplest case, the latter is the dummy for the occurrence of an accident (our previous RT variable), and the equation takes the form :

$$RT_i = \mathbb{I}[g(X_i, \gamma) + \eta_i > 0] \quad (2)$$

Note that this setting can easily be generalized. For instance, a recent contribution by Richaudeau (1999) takes into account the number of accident, modelled as following a negative binomial distribution¹⁸. Equation (2) is estimated using a count data model; the η_i are approximated by their 'generalized residual' counterpart. In the same way, the distribution of accident costs (conditional on occurrence) can be introduced at that stage.

The key idea, then, is to simultaneously estimate the two equations, allowing for general correlation across the error terms. According to standard theory, asymmetric information should result in a positive correlation, under the convention that $y_i = 1$ (resp. $RT_i = 1$) corresponds to more comprehensive coverage (resp. the occurrence of an accident). One obvious advantage of this setting is that it does not require the estimation of the pricing policy followed by the firm, which is probably an extremely difficult task - and a potential source of important bias.

To circumvent the non linearity problems discussed above, as well as the issues raised by experience rating, Chiappori and Salanié consider a subsample of inexperienced drivers, and introduce a large number of exogenous variables, allowing for cross-effects. They use both a parametric and a non parametric approach. The latter relies upon the construction of a large number of 'cells', each cell being defined by a particular profile of exogenous variables. Under the null (in the absence of adverse selection), within each

¹⁸In practice, Richaudeau includes among the regressors of the second equation the generalized residual obtained in the contract choice probit. Under the null of no correlation, the coefficient should be zero. Using a very complete French data base, he cannot reject the null.

cell the choice of contract and the occurrence of an accident should be independent, which can easily be checked using a χ^2 test).

This method can be given a general form. Following the presentation proposed by Dionne, Gouriéroux and Vanasse (1997) and Gouriéroux (1999), a general strategy can be summarized as follows. Let Y , X and Z respectively denote the endogenous variable under consideration (say, the occurrence of an accident), the initial exogenous variables and the decision variables at the agent's disposal (say, the choice of a particular contract within a given menu). Let $l(Y | X, Z)$ denote the probability distribution of Y conditional on X and Z . In the absence of adverse selection, the agent's choice conveys no information upon the endogenous variable. The translation is that :

$$l(Y | X, Z) = l(Y | X)$$

Obviously, this relationship can be given different but equivalent forms, such as:

$$l(Z | X, Y) = l(Z | X)$$

or

$$l(Y, Z | X) = l(Y | X)l(Z | X)$$

(the latter version expressing the fact that, conditionally on X , Y and Z should be independent).

Interestingly enough, in all the empirical applications to automobile insurance just listed (with the exception of the initial paper by Puelz and Snow), independence is not rejected; in other words, these studies find no evidence of adverse selection. The conclusion is that, at least in the context of automobile insurance, adverse selection may not be a crucial issue. One remark must be stressed at this point. According to the previous arguments, the existence of a positive correlation across the residuals cannot be interpreted as establishing the presence of asymmetric information without some precautions : as argued above, any misspecification can indeed lead to a spurious correlation. Parametric approaches, in particular, are highly vulnerable to this type of flaws, especially when they rely upon some simple, linear form. But the argument is not symmetric. Suppose, indeed, that some empirical study does *not* reject the null (i.e., the absence of correlation). Although, in principle, this result might as well be due to a misspecification bias, this explanation is much less credible in that case; for it must be the case that, while (fully conditional) residuals are actually positively correlated, the bias

goes in the opposite direction with the *same* (absolute) magnitude - so that it exactly offsets the correlation. In other words, misspecifications are much more likely to bias the results *in favor* of the presence of adverse selection.

Accidents versus claims A related issue is the distinction between accidents and claims discussed in the theory section. A regression using claims as the dependent variable is likely to generate misleading results, because a larger deductible automatically discourages small accidents reporting, hence reduces the number of claims even when the accident rate remains constant. This is a serious problem, that can be addressed in two different ways. The solution proposed by Chiappori and Salanié (2000) is to discard all accidents where one vehicle only is involved. Whenever two automobiles are involved, a claim is much more likely to be filed in any case¹⁹. A more restrictive version is to exclusively consider accidents involving bodily injuries, since reporting is mandatory in that case; the cost being a drastic reduction in the number of accidents.

Alternatively, one can explicitly model the filing decision as part of the accident process. For any accident, the agent computes the net benefit of filing a claim, and reports the accident only when this benefit is positive (or above some threshold). Although accidents involving no claims are generally not observed²⁰, adequate econometric techniques can be used. Note, however, that these require the estimation of a complete structural model, as in Dionne, Gouriéroux and Vanasse (1998).

3.3 Adverse selection versus moral hazard

Natural experiments As argued above, the previous tests are not specific of adverse selection. Moral hazard would typically lead to the same kind of correlation, although with a different causality. In order to distinguish

¹⁹In principle, the two drivers may agree on some bilateral transfer and thus avoid the penalties arising from experience rating. Such a 'street-settled' deal is however quite difficult to implement between agents who meet randomly, will probably never meet again, and cannot commit in any legally enforceable way (since declaration is in general compulsory according to insurance contracts). We follow the general opinion in the profession that such bilateral agreements can be neglected.

²⁰Some data sets do, however, record accidents that did not result in claims. Usually, such data sets have been collected independently of insurance companies. See Richaudeau (1999).

between adverse selection and moral hazard, one needs some additional structure. Of particular interest are the situations where a 'natural experiment' takes place. Assume that, for some exogenous reason (say, a change in regulation), a given, exogenously selected set of agents experiences a sudden and exogenous modification in the incentive structure they are facing. Then the resulting changes in behavior can be directly studied; and adverse selection is no longer a problem, since it is possible to concentrate upon agents that remained insured throughout the process²¹. A typical example is provided by the changes in automobile insurance regulation in Québec, where a "no fault" system was introduced in 1978, then deeply modified in 1992. Dionne and Vanasse (1996) recently provided a careful investigation of the effects of these changes. They show that the new system provided strong incentives to increase prevention, and that, as a result, the average accident frequency dropped significantly during the years that followed its introduction. They conclude that changes in agents' behavior, as triggered by new incentives, did have a significant effect on accident probabilities²².

A limitation of any work of this kind is that, strictly speaking, it establishes a simultaneity rather than a causality. What the Dionne and Vanasse study shows is that, on a given period, accident probabilities have changed significantly, and that this evolution immediately followed a structural change in regulation. But, of course, the two phenomena might stem from simultaneous and independent causes. Such a 'coincidence' may be more or less plausible. In the particular case of the Quebec reform, for instance, the incentive explanation remains by far the more convincing one, given both the magnitude of the drop in sinistrality and the absence of other major changes that could account for it during the period under consideration.

Still, an ideal experiment would involve a 'reference' sample that is not affected by the change, so that the effects can be estimated in differences (or more precisely differences of differences), allowing for very convincing tests. A paper by Dionne and St-Michel (1991) provides a good illustration of this idea. They study the impact of a regulatory variation of coinsurance level in the Quebec public insurance plan on the demand for days of compensation. The main methodological contribution of the paper is to introduce a distinction between injuries, based on the type of diagnosis (easy or difficult). This

²¹In addition, analyzing the resulting attrition (if any) may in some cases convey interesting information on selection issues.

²²See Browne and Puelz (1998) for a similar study on US data.

distinction is based on information obtained from the medical literature; it reflects the fact that it is much easier for a physician to detect a fracture than, say, lower back pain. If moral hazard is more prevalent when the information asymmetry is larger, theory predicts that the regulatory change will have more significant effects on the number of days of compensation for those cases where the diagnosis is more problematic. This prediction is clearly confirmed by empirical evidence. A more generous insurance coverage, resulting from an exogenous regulatory change, is found to increase the number of days of compensations, but only for the cases of difficult diagnoses. Note that the effect thus identified is *ex post* moral hazard. The reform is unlikely to have triggered significant changes in prevention; and, in any case, such changes would have affected all types of accidents.

Additional evidence is provided by Fortin et al. (1994), who examine how the Canadian Worker's Compensation (WC) and the Unemployment Insurance (UI) programs interact to influence the duration of workplace accidents. Here, the duration is estimated from a mixed proportional hazard model, where the baseline hazard is estimated non parametrically, and unobserved heterogeneity is taken into account using a gamma distribution. They show that an increase in the generosity of Worker's Compensation in Quebec leads to an increase in the duration of accidents. In addition, a reduction in the generosity of Unemployment Insurance is, as in Dionne and St-Michel, associated with an increase in the duration of accidents that are difficult to diagnose. The underlying intuition is that worker's compensation can be used as a substitute to unemployment insurance. When a worker goes back to the labor market, he may be unemployed and entitled to UI payments for a certain period. Whenever worker's compensation is more generous than unemployment insurance, there will be strong incentives to delay the return to the market. In particular, the authors show that the hazard of leaving WC is 27% lower when an accident occurs at the end of the construction season, when unemployment is seasonally maximum²³.

In two recent papers, Chiappori, Durand and Geoffard (1998) and Chiappori, Geoffard and Kyriadizou (1998) use data on health insurance that display similar features. Following a change in regulation in 1993, French health insurance companies modified the coverage offered by their contracts in a non uniform way. Some of them increased the level of deductible, while

²³See also Fortin and Lanoie (1992), Bolduc et al. (1997), and the survey by Fortin and Lanoie (1998)

other did not. The tests use a panel of clients belonging to different companies, who were faced with different changes in coverage, and whose demand for health services are observed before and after the change in regulation. In order to concentrate upon those decisions that are essentially made by consumers themselves (as opposed to those partially induced by the physician), the authors study the occurrence of a physician visit, distinguishing between general practitioner (GP) office visits, GP home visits and specialist visits. They find that the number of home visits significantly decreased for the 'test group' (i.e., agents who experience a change of coverage), but not for the reference group (for which the coverage remained constant). They argue that this difference is unlikely to result from selection, since the two populations are employed by similar firms, display similar characteristics, and that participation to the health insurance scheme was mandatory.

'Quasi natural experiments' Natural experiments are valuable but scarce. In some cases, however, a static context exhibits specific features that keep the flavor of a natural experiment, although no exogenous *change* of the incentive structure can be observed. The key remark is that any situation where identical agents are, for *exogenous* reasons, faced with different incentive schemes can be used for testing for moral hazard. The problem, of course, is to check that the differences in schemes are purely exogenous, and do not reflect some hidden characteristics of the agents. For instance, Chiappori and Salanié (1997) consider the case of French automobile insurance, where young drivers whose parents have low past accident rates can benefit from a reduction in premium. Given the particular properties of the French experience rating system, it turns out that the marginal cost of accident is reduced for these drivers. In a moral hazard context, this should result in less cautious behavior and higher accident probabilities. If, on the contrary, the parents' and children's driving abilities are (positively) correlated, a lower premium should signal a better driver, hence translate into less accidents. The specific features of the French situation thus allow to distinguish between the two types of effects. Chiappori and Salanié find evidence in favor of the second explanation : the accident rates of the 'favored' young drivers are, other things equal, smaller than average by a small but significant percentage.

A contribution by Cardon and Hendel (1998) extends these ideas in a very stimulating way. They consider a set of individuals who face different menus of employer-based health insurance policies, under the assumption that there

is no selection bias in the allocation of individuals across employers. Two types of behavior can then be observed. First, agents choose one particular policy within the menu at their disposal; second, they decide on the level of health expenditures. The authors identify a fully structural model, which allows them to simultaneously estimate a selection equation that describe the policy choice, and estimate the price elasticity of demand controlling for selection bias. The key ingredient for identifying the specific effects of moral hazard is that while people are free to choose any contract in the menu they face, they cannot choose the menu itself; and different menus involve different coinsurance level. The "quasi-experimental" features stem precisely from this random assignment of people to different choice sets. Even if less risky people always choose the contract with minimum coinsurance, the corresponding coinsurance rates will differ across firms. In other words, it is still the case that identical people in different firms face different contracts (i.e., different coinsurance rates) for exogenous reasons (i.e., because of the choice made by their employer)²⁴. Interestingly enough, the authors find no evidence of adverse selection, while price elasticities are negative and very close to those obtained in the Rand survey alluded to above. This suggest that moral hazard, rather than adverse selection, may be the main source of asymmetric information in that case.

4 Dynamic models of information asymmetries

As indicated above, tests based on the dynamics of the contractual relationship can either study the qualitative features of existing contracts assuming they are optimal in the relevant context, or take existing contracts as given and investigate the testable properties of the induced individual behavior.

Tests assuming optimal contracts Only a few empirical studies consider the dynamics of insurance relationships. An important contribution is due to Dionne and Doherty (1994), who use a model of repeated adverse selection with one-sided commitment. Their main purpose is to test the 'highballing'

²⁴As Cardon and Hendel put it : '... the coinsurance is an endogenous variable... but since different individuals face different choice sets, the premium at which insurance was offered becomes a useful instrument for the coinsurance' (1998, p.21).

prediction, according to which the insurance company should make positive profits in the first period, compensated by low, below-cost second period prices. They test this property on Californian automobile insurance data. According to the theory, when various types of contracts are available, low risk agents are more likely to choose the experience rated policies. Since these are characterized by highballing, the loss to premium ratio should rise with the cohort age. If insurance companies are classified according to their average loss per vehicle (which reflects the 'quality' of their portfolio), one expects the premium growth to be negative for the best quality portfolios; in addition, the corresponding slope should be larger for firms with higher average loss ratios. This prediction is confirmed by the data. Insurance companies are classified into three subgroups. The slope coefficient is negative and significant for the first group (with lowest average loss), positive and significant for the third group, non significant for the intermediate group. They conclude that the 'highballing' prediction is not rejected.

Recently, Hendel and Lizzeri (1999) have provided very convincing tests of the symmetric learning model a la Harris and Holmstrom (1982) on life insurance data. They use a rich contract data base that includes information on the entire profile of future premiums. Some contracts involve commitment from the insurer, in the sense that the dynamics of future premium is fixed in advance and cannot depend on the evolution of the insuree's health. For other contracts, however, future premia are contingent on health. Specifically, the premium increases sharply unless the insured is still in good health (as certified, for instance, by a medical examination). In this context, the symmetric learning model generates very precise predictions on the comparison between contracts with and without commitment. Contracts with non contingent future premia should entail front loading, representing the cost of the insurance against the classification risk. They should also lock-in a larger fraction of the consumers, hence exhibit a lower lapsation rate; in addition, only better risk types are likely to lapse, so that the average quality of the insurer's client portfolio should be worse, which implies a higher present value of premiums for a fixed period of coverage. Hendel and Lizzeri show that all of these predictions are satisfied by existing contracts²⁵. Finally, the authors

²⁵The main puzzle raised by these findings is that a significant fraction of the population does not choose commitment contracts, i.e., does not insure against the classification risk. The natural explanation suggested by theory (credit rationing) is not very convincing in that case, since differences in premiums between commitment and no commitment contracts are small (less than \$300 per year), especially for a client pool that includes

study accidental death contracts, i.e. life insurance contracts that only pay if death is accidental. Strikingly enough, these contracts, where learning is probably much less prevalent, exhibit none of the above features.

Another characteristic feature of the symmetric learning model is that any friction reducing the clients' mobility, although ex post inefficient, is often ex ante beneficial, because it increases the agents' ability to (implicitly) commit and allows for a larger coverage of the classification risk. Using this result, Crocker and Moran (1998) study employment-based health insurance contracts. They derive and test two main predictions. One is that when employers offer the same contract to all of their workers, the coverage limitation should be inversely proportional to the degree of worker commitment, as measured by his level of firm-specific human capital. Secondly, some contracts offer 'cafeteria plans', whereby the employee can choose among a menu of options. This self-selection device allows the contract to change in response to interim heterogeneity of insurees. In this case, the authors show that the optimal (separating) contract should exhibit more complete coverage, but that the premiums should partially reflect the health status. Both predictions turn out to be confirmed by the data. Together with the results obtained by Hendel and Lizzeri, this fact that strongly suggests the symmetric learning model is particularly adequate in this context.

Behavioral dynamics under existing contracts Natural as it seems, the assumption that contracts are always optimal may for some applications be problematic. For one thing, theory is often inconclusive. Little is known, for instance, on the form of optimal contracts in a repeated moral hazard framework, at least in the (realistic) case where the agent can freely save. And the few results we have either require utterly restrictive assumptions (CARA utilities, monetary cost of effort) or exhibit features (randomized contracts, for instance) that sharply contrast with real life observations. Even skeptics of bounded rationality theories may accept that such very sophisticated constructs, that can hardly be understood by the average insurance salesman (let alone the average consumer), are unlikely to be implemented on a large scale²⁶.

executives, doctors, businessmen and other high income individuals. Heterogeneous risk perception across individuals is a better story, but formal tests still have to be developed. Obviously, more research is needed on this issue.

²⁶A more technical problem with the optimality assumption is that it tends to generate complex endogeneity problems. Typically, one would like to compare the features of the

Another potential deviation from optimality comes from the existence of regulations, if only because regulations often impose very simple rules that fail to reproduce the complexity of optimal contracts. An interesting example is provided by the regulation on experience rating by automobile insurance companies, as implemented in some countries. A very popular rule is the 'bonus/malus' scheme, whereby the premium is multiplied by some constant larger (resp. smaller) than one for each year with (resp. without) accident. Theory strongly suggests that this scheme is too simple in a number of ways. In principle, the malus coefficient should not be uniform, but should vary with the current premium and the driver's characteristics; the deductible should vary as well; etc.²⁷.

Still, one can take this (probably suboptimal) scheme as given, and use theory to derive the main testable features of individual behavior for the various models at stake. In a recent contribution, Chiappori and Heckman (2000) test in this context the 'negative contagion' prediction that can be derived in a moral hazard framework. The idea is that, for experience rating schemes of this kind, the occurrence of an accident typically increases the marginal cost of a future accidents. This should strengthen incentives to drive safely, hence reduce the probability of future accidents. The problem, however, is that any unobserved heterogeneity among individuals will generate an opposite, 'positive contagion' phenomenon, since bad drivers had more accidents in the past and will probably have more accidents in the future. Technically, the 'negative contagion' effect obtains only *conditionally* on agents' characteristics, including unobserved ones. The challenging econometric puzzle, at this point, is to disentangle the two aspects. In principle, this is feasible whenever panel data are available. Chiappori and Heckman use French data, for which regulation imposes that insurers increase the premium by 25% in case an accident occurs; conversely, in the absence of any accident during one year, the premium drops by 5%²⁸. The technique they suggest relies upon

various existing contracts. The optimality approach requires that each contract is understood as the optimal response to a specific context, so that differences in contracts simply reflect differences between the 'environments' of the various firms. In econometric terms, contracts are thus, by assumption, endogenous to some (probably unobserved) heterogeneity across firms, a fact that may, if not corrected, generate biases in the estimations.

²⁷Of course, the precise form of the optimal scheme depends on the type of model. It is however basically impossible to find a model for which the existing scheme is optimal.

²⁸In addition, several non linearities have been introduced; e.g., there exist both a floor and a ceiling on the resulting 'bonus/malus' coefficient.

existing result on the distinction between pure heterogeneity and state dependence (see Heckman (1978)). To get the intuition in a simple way, assume the system is malus only (i.e., the premium increases after each accident, but does not decrease subsequently), and consider two sequences of 4 years records, $A = (1, 0, 0, 0)$ and $B = (0, 0, 0, 1)$, where 1 (resp. 0) corresponds to the occurrence of an accident (resp. no accident) during the year. In the absence of moral hazard, and assuming away learning phenomena, the probability of the two sequences should be exactly identical; in both cases, the observed accident frequency is 25%. Under moral hazard, however, the first sequence is more probable than the second : in A , the sequence of three years without accident happens after an accident, hence when the premium, and consequently the marginal cost of future accidents and the incentives to take care are maximum²⁹. In other words, for a given average frequency of accidents, the precise timing of the occurrences can provide valuable information upon the importance of incentives or disincentives effects. More surprisingly, Chiappori and Heckman show that a precise record of the sequence is not even needed when the distribution of new contracts is stationary (conditionally on observable). The knowledge of each individual's number of years of driving and number of accidents is then sufficient to test for moral hazard.

5 Conclusion

As argued in the introduction, empirical applications of contract theory is likely to become a burgeoning field; and it is a safe bet that insurance data will play an important role in these developments. Several studies have already contributed to a better knowledge of the impact of adverse selection and moral hazard in various markets. In several cases, the importance of information asymmetries has been found to be limited. This by no means implies, however, that such phenomena are of no importance in insurance. For one thing, the existence and the consequences of informational asymmetries vary considerably across markets. For instance, various (alas unpublished) studies have found strong adverse selection effects in private unemployment insurance markets³⁰.

²⁹Interestingly enough, if, as argued by many insurers, there is learning, in the sense that experienced drivers are better drivers, then A is more likely than B .

³⁰This is particularly true for unemployment insurance contracts linked to mortgages. See Chiappori and Pinquet (1998) for an overview in the French case.

In addition, there exist a number of crucial normative issues where our theoretical and empirical knowledge of asymmetric information are likely to play a crucial role. The economics of discrimination in insurance provide an important example. The availability of an always larger range of medical tests allows insurance companies to classify people according to their health risk in a more precise way. A classic remark by Hirshleifer (1971) is that this progress comes with a cost - namely, the classification risk becomes uninsurable. The induced changes can have a major impact on individuals' lives; think, for instance, of the consequences of the introduction of the HIV test, or the forthcoming developments of genetic testing. A possible solution consists in regulating the use of such data by insurance companies. For instance, many countries restrict (and sometimes prohibit) the use of HIV tests for health insurance pricing. For an economist, however, the potential perverse effect of this regulation is to replace explicit discrimination by adverse selection, which may sometimes result either in similar discrimination plus signalling inefficiencies, or even in market collapse. Policy recommendations on this issue must rely on both a theoretical understanding of the issue and an empirical evaluation of the magnitude of the effects. Obviously, these problems are crying out for more work.

Finally, a better understanding of actual behavior is likely to require new theoretical tools. The perception of accident probabilities by the insureds, for instance, is a very difficult problem on which little is known presently. Existing results, however, strongly suggest that standard theoretical models relying on expected utility maximization using the 'true' probability distribution may fail to capture some key aspects of many real-life situations. Here again, the application of existing theory to insurance data is likely to reveal an extremely promising research direction.

References

- Bolduc, D., Fortin, B., F. Labrecque and P. Lanoie (1997)**, "Incentive Effects of Public Insurance Programs on the Occurrence and the Composition of Workplace Injuries", CIRANO Scientific Series, Montreal, 97s-24.
- Boyer, M., et G. Dionne (1989)** : "An Empirical Analysis of Moral Hazard and Experience Rating", *Review of Economics and Statistics*, 71, 128-34
- Browne, M., and R. Puelz (1998)**: "The Effect of Legal Rules on the Value of Economic and Non-Economic Damages and the Decision to File", mimeo, University of Wisconsin-Madison
- Brugiavini, A. (1990)** : "Longevity risk and the life cycle", PhD Dissertation, LSE, London
- Cardon, J., and I. Hendel (1998)**: 'Asymmetric Information in Health Insurance : Evidence From the National Health Expenditure Survey', Mimeo, Princeton University
- Cawley. J. and T. Philipson (1997)** : "An Empirical Examination of Information Barriers to Trade in Insurance", Mimeo, University of Chicago.
- Chassagnon, A. (1996)**: "Anti-selection: modèle générique et applications", thèse de doctorat, DELTA-EHESS.
- Chassagnon, A., and P.A. Chiappori (1997)**: "Insurance under moral hazard and adverse selection: the competitive case", mimeo, DELTA.
- Chiappori, P. A. (1994)**: "Assurance et économétrie des contrats : quelques directions de recherche", mimeo, DELTA.
- Chiappori, P. A., F. Durand and P.Y. Geoffard (1998)**, "Moral Hazard and the Demand for Physician Services : First Lessons from a French Natural Experiment", *European Economic Review*, 42, 499-511
- Chiappori, P. A., P.Y. Geoffard and E. Kyriadizou (1998)**, "Cost of Time, Moral Hazard, and the Demand for Physician Services", Mimeo, University of Chicago.
- Chiappori, P. A. and J. Heckman (2000)**, 'Testing for Adverse Selection versus Moral Hazard from dynamic data', Mimeo, University of Chicago.
- Chiappori, P. A., I. Macho, P. Rey and B. Salanié (1994)**, "Repeated moral hazard : memory, commitment, and the access to credit markets", *European Economic Review*, 1527-53.

Chiappori, P. A. and J. Pinquet (1998) : 'Assurance chômage des emprunteurs', *Rapport pour la Direction de la prévision du Ministère de l'économie et des finances*; forthcoming in *Revue Française d'Économie*.

Chiappori, P. A. and B. Salanié (1996), "Empirical Contract Theory: The Case of Insurance Data", *European Economic Review*, 41, 943-51

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

Chiappori, P. A., B. Salanié and J. Valentin (1999), "Early Starters vs Late Beginners", *Journal of Political Economy*, 107, 731-60

Cooper, R., and B. Hayes (1987) : "Multi-period Insurance Policies", *International Journal of Industrial Organization*, 5, 211-31.

Crocker, K. and J. Moran (1998) : "Contracting with Limited Commitment : Evidence from Employment-Based Life Insurance Contracts", Working Paper, University of Michigan, Ann Arbor.

Dahlby, B. (1983), "Adverse Selection and Statistical Discrimination: An Analysis of Canadian Automobile Insurance", *Journal of Public Economics*, 20, 121-130.

Dahlby, B. (1992), "Testing for Asymmetric Information in Canadian Automobile Insurance", in *Contributions to Insurance Economics*, G. Dionne ed, Kluwer.

de Garidel, T. (1997): "Pareto Improving Asymmetric Information in a Dynamic Insurance Market", D.P. 266, LSE, London.

de Meza, D, and D. Webb (1999):"The Timid and the Reckless : Risk Preference, Precaution and Overinsurance", WP, LSE, London.

Dionne, G, Ed. (1992): *Contributions to Insurance Economics*, Kluwer, Boston.

Dionne, G., and N. Doherty (1994), "Adverse Selection, Commitment and Renegotiation: Extension to and Evidence from Insurance Markets", *Journal of Political Economy*, 102-2, 210-35

Dionne, G., C. Gouriéroux and C. Vanasse (1997), "The Informational Content of Household Decisions, With an Application to Insurance under Adverse Selection", W.P., HEC, Montreal.

Dionne, G., C. Gouriéroux and C. Vanasse (1998), "Evidence of Adverse Selection in Automobile Insurance Markets", mimeo.

Dionne, G, and P. St-Michel (1991): 'Worker's Compensation and Moral hazard', *Review of Economics and Statistics*, LXXIII, 236-44

Dionne, G. and C. Vanasse (1992), "Automobile Insurance Ratemaking in the Presence of Asymmetrical Information", *Journal of Applied Econo-*

metrics, 7, 149-65

Dionne, G. and C. Vanasse (1996), "Une evaluation empirique de la nouvelle tarification de l'assurance automobile au Quebec", Mimeo, Universite de Montreal.

Fombaron, N. (19) :

Fortin, B. and P. Lanoie (1992), "Substitution between Unemployment Insurance and Workers' Compensation", *Journal of Public Economics*, 49, 287-312.

Fortin, B., P. Lanoie and C. Laporte (1995), "Is Workers' Compensation Disguised Unemployment Insurance", CIRANO Scientific Series, Montreal, 95s-48.

Fortin, B and P. Lanoie, "Effects of Workers' Compensation: a Survey" (1998), CIRANO Scientific Series, Montreal, 98s-04.

Friedman, B.M., and M.J. Warshawski (1990) : 'The cost of annuities : implications for savings behavior and bequests", *Quarterly Journal of Economics*, 420, 135-54

Fudenberg, D., B. Holmstrom and P. Milgrom (1990), "Short-Term Contracts and Long-Term Agency Relationship", *Journal of Economic Theory*, 51, 1-31

Gouriéroux, C. (1999): "The Econometrics of Risk Classification in Insurance", *Geneva Papers on Risk and Insurance Theory*, 24, 119-139.

Harris M. and B. Holmstrom (1982), "A Theory of Wage Dynamics", *Review of Economic Studies*, 49, 315-333.

Hart, O. and J. Tirole (1988) : "Contract Renegotiation and Coasian Dynamics", *Review of Economics Studies*, 55, 509-40.

Heckman, J.J. (1978): "Simple statistical models for discrete panel data developed and applied to test the hypothesis of true state dependence against the hypothesis of spurious state dependence, *Annales de l'INSEE*, 30-31, p. 227-269.

Hellwig, M. (1987) : "Some recent developments in the theory of competition in markets with adverse selection", *European Economic Review*, 31:154-63.

Hendel, I. and A. Lizzeri (1999): "The Role of Commitment in Dynamic Contracts : Evidence from Life Insurance", Working Paper, Princeton University.

Hirshleifer, J. (1971): "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 61, 561-574.

Jullien, B., B. Salanié and F. Salanié (1999): "Should More Risk-Averse Agents Exert More Effort?", *Geneva Papers on Risk and Insurance Theory*, 24, 19-28.

Laffont, J.-J. and J. Tirole (1993) : *A Theory of Incentives in Procurement and Regulation*, Cambridge, MIT Press.

Manning, W. et al. (1987), "Health insurance and the demand for medical care : evidence from a randomized experiment," *American Economic Review*, 77(3) : 251-277.

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.

Richaudeau, D. (1999): "Automobile insurance contracts and risk of accident: An Empirical Test Using French Individual Data", *The Geneva Papers on Risk and Insurance Theory*, 24, 97-114

Rochet, J.C., and P. Choné (1998): "Ironing, Sweeping, and Multi-dimensional Screening", *Econometrica*, 66, 783-827

Rogerson, W. (1985): "Repeated Moral Hazard", *Econometrica*, 53, 69-76.

Rothschild, M. and J. Stiglitz (1976), "Equilibrium in Competitive Insurance Markets", *Quarterly Journal of Economics*, 90, 629-649.

Rubinstein, A., and M. Yaari (1983): "Insurance and Moral Hazard", *Journal of Economic Theory*, 14, 441-52.

Salanié, B. (1997), *The Economics of Contracts: A Primer*, MIT Press.

Villeneuve, B. (1996), "Essais en économie de l'assurance", thèse de doctorat, DELTA-EHESS