VU Research Portal

Adverse selection and moral hazard in insurance: Can dynamic data help to distinguish?

Abbring, J.H.; Chiappori, P.A.; Heckman, J.J.; Pinquet, J.

published in

Journal of the European Economic Association
2005

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Abbring, J. H., Chiappori, P. A., Heckman, J. J., & Pinquet, J. (2005). Adverse selection and moral hazard in insurance: Can dynamic data help to distinguish? *Journal of the European Economic Association*, 1, 512-521.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Adverse Selection and Moral Hazard in Insurance: Can Dynamic Data Help to Distinguish?*

Jaap H. Abbring[†] Pierre-André Chiappori[‡]
James J. Heckman[§] Jean Pinquet[¶]

Keywords: insurance, moral hazard, selection, empirical tests of contract theory.

JEL-codes: D82, G22, C41, C14.

Abstract

A standard problem of applied contracts theory is to empirically distinguish between adverse selection and moral hazard. We show that dynamic insurance data allow to distinguish moral hazard from dynamic selection on unobservables. In the presence of moral hazard, experience rating implies negative occurrence dependence: individual claim intensities decrease with the number of past claims. We discuss econometric tests for the various types of data that are typically available. Finally, we argue that dynamic data also allow to test for adverse selection, even if it is based on asymmetric learning.

*Support from the NSF (grant #0096516) is gratefully acknowledged.

[†]Department of Economics, Free University, Amsterdam. This paper was written while Jaap Abbring was visiting the Department of Economics of University College London. Abbring's research is supported by a fellowship of the Royal Netherlands Academy of Arts and Sciences.

[‡]Corresponding author: Department of Economics, University of Chicago, 1126 E. 59th Street, Chicago IL 60637, USA. Tel. (1) 773 702 98 11, email: pchiappo@midway.uchicago.edu.

[§]Department of Economics, University of Chicago.

[¶]Université Paris X - Nanterre.

1. Introduction

For two decades, contract theory has remained a predominantly theoretical field. However, a number of papers have recently been devoted to empirical applications of the theory.¹ It has been argued that insurance offers a particularly promising field for empirical work on contracts. Individual (automobile, housing, health, life, etc.) insurance contracts are largely standardized. Researchers have access to data bases of insurance companies, which typically contain several millions of such contracts. The information in these data bases can generally be summarized in a reasonably small number of quantitative and qualitative indicators. The 'outcome' of the contract - be it the occurrence of an accident, its cost, or some level of expenditure - is very precisely recorded in the firms' files, together with a detailed history of the contractual relationship (changes in coverage, etc.). Not surprisingly, several recent papers are aimed at testing for the existence and estimating the magnitude of asymmetric information effects in competitive insurance markets.²

A popular strategy for studying asymmetric information is to test, conditional on observables, for a correlation between the choice of a contract and the occurrence or severity of an accident. Under adverse selection on risk, 'high risk' agents are, everything else equal, both more likely to choose a contract with more complete coverage and more likely to have an accident. The basic moral hazard story is very close to the adverse selection one, except for an inverted causality. In a moral hazard context, agents first choose different contracts. Then, an agent facing better coverage and, therefore, weaker incentives will be less cautious and have more accidents. In both cases, the same pattern emerges: controlling for observables, more comprehensive coverage should be associated with higher realized risk - a property that can be tested using appropriate parametric or non parametric techniques.

The conditional correlation approach has several advantages. It is simple and very robust, as argued by Chiappori et al. (2002). Furthermore, it can be used on static, cross-sectional data that are relatively easy to obtain. However, these

¹See Chiappori and Salanié (2003) for a recent survey.

²Puelitz and Snow (1994), Dionne and Vanasse (1992), Chiappori and Salanié (1996, 2000), Dionne, Gouriéroux and Vanasse (1997), Richaudeau (1999) and Dionne et al. (2001), to name only a few, analyze automobile insurance contracts, while Holly et al. (1998), Chiappori, Durand and Geoffard (1998), Chiappori, Geoffard and Kyriadizou (1998), Cardon and Hendel (1998) and Hendel and Lizzeri (1999) use health or life insurance data, and Poterba and Finkelstein (2001) consider annuity contracts.

qualities come at a cost. The past history of the relationship influences both the current contract (through experience rating) and the agent's behavior, and this effect is hard to take into account with cross-sectional data. More importantly, the correlation is not informative on the direction of the causality, which makes the two stories (moral hazard and adverse selection) very hard to distinguish. Still, such a distinction is crucial, if only because the optimal form of regulation of insurance markets varies considerably with the context.³

The research program summarized in the present paper relies on the insight that the dynamic aspects of the relationship can help distinguishing between adverse selection and moral hazard. Two approaches can be distinguished. First, the form of optimal dynamic contracts differs considerably between the two cases. Thus, the qualitative properties of observed contracts may provide useful insights into the type of problem they are designed to address. The research program described in this paper concentrates on a second approach, in which the (possibly suboptimal) contracts are taken as given and we concentrate on their implications for observed behavior. In particular, most 'real life' insurance contracts exhibit some form of experience rating. A typical property of experience rating schemes is that the occurrence of an accident shifts the entire incentive scheme the agent is facing. Under moral hazard, this results in a form of autocorrelation in the accident process. Thus, an empirical analysis of this process can be informative on the presence of moral hazard.

In addition, dynamic data allow to address the problem of asymmetric learning. Conventional wisdom suggests that, in many cases, asymmetric information may not be present at the beginning of the relationship (e.g., the relative quality of a young driver is unknown to her and her insurer). Rather, it emerges gradually as a consequence of different learning processes (say, the young driver learns from near misses that are not even observed by the insurer). Then the contractual changes that take place during the relationship may be informative about the agent's riskiness, even if the initial choice of a contract is uncorrelated with residual risk (as found by most studies).

2. Dynamic moral hazard under experience rating: theory

The model is directly borrowed from Chiappori and Heckman (2000) and Abbring et al. (forthcoming). We consider a dynamic version of an insurance model a

³For a review of various attempts to distinguish between moral hazard and adverse selection, see Chiappori (2000).

la Mossin. Time is discrete. In each period t , the agent receives an income normalized to one and may with some probability $(1 - p_t)$ incur a fixed monetary loss L . She is covered by an insurance contract involving a fixed deductible D and a premium Q_t that depends on past experience. Specifically, the evolution of Q_t is governed by the following 'bonus-malus' system:

$$\begin{aligned} Q_{t+1} &= \delta Q_t \text{ if no accident occurred at period } t \\ &= \gamma Q_t \text{ if an accident occurred at period } t \end{aligned}$$

where $\delta < 1 < \gamma$.⁴

The no-accident probability p_t is subject to moral hazard. Specifically, in each period t the agent chooses an effort level $e_t \geq 0$, resulting in a no accident probability $p_t = p(e_t)$ for some increasing, concave function p . The cost of effort is assumed separable, i.e. the agent attaches utility

$$u(x) - e$$

to income x if he exerts effort e , where u is increasing and strictly concave. The horizon is infinite and agents maximize expected discounted utility.

According to the bonus-malus scheme, each accident shifts the incentive scheme faced by the agent upward, thus modifying her incentives. It follows that the 'cost' of an accident, in terms of higher future premia, depends on random events (the sequence of future accidents) and endogenous decisions (the sequence of future efforts). Technically, the agent must solve a stochastic control problem. Here, we simply summarize the main properties of the solution; the reader is referred to Abbring et al. (forthcoming) for a precise analysis. A first result is that past experience matters for the current decision only through the current level of the premium; i.e., Q_t is the only state variable of the control problem. Secondly, the optimal effort is increasing in the premium, at least when both the premium and the deductible are small relative to the agent's income. It follows that the accident probability process of any given agent will exhibit a *negative occurrence-dependence property*. In the absence of an accident, the premium - hence, by our result, the agent's incentives - decreases. Effort is optimally reduced, resulting in a

⁴Proportional bonus-malus scheme of this type are empirically frequent. The French system, which is relevant for our empirical application, corresponds to $\delta = .95$ and $\gamma = 1.25$. In addition, the French system imposes a floor and a ceiling on θ_t , respectively equal to .5 and 3.5. In our discussion of the French system, we ignore the fact that accidents occur continuously but premiums are only updated annually. See Abbring et al. (forthcoming) for a formal discussion.

steady increase of the accident probability. However, the occurrence of an accident generates a discrete jump in the premium, which boosts incentives and ultimately results in a drop in the accident probability. The main testable implication of the model is thus the following:⁵

The accident process exhibits negative occurrence dependence, in the sense that individual claim intensities decrease with the number of past claims.

This suggests that we can test for moral hazard by simply testing for negative occurrence dependence in the raw data. One should however be careful at this point. While moral hazard implies occurrence dependence effects at the individual level, individual claim intensities also vary with observed characteristics (such as age, driving experience, region, etc.) and, more importantly, with unobserved individual heterogeneity factors. In automobile insurance, for example, existing work strongly suggests that unobserved heterogeneity is paramount. It is well known that unobserved heterogeneity results in (spurious) *positive* occurrence dependence in the data. The intuition is that those individuals whose risk is persistently high for unobserved external reasons will be more likely to have had accidents in the past *and* to have accidents in the future (in other words, to the extent that 'bad' drivers remain bad for at least a while, we should expect to find a positive correlation between past and future accident rates). Of course, this effect, which is overwhelmingly confirmed by the data, does not contradict the theoretical analysis sketched above: whatever the distribution of unobserved heterogeneity, it is still true that under moral hazard, the accident probability of *each individual* decreases with *the person's* number of past claims. But any empirical investigation of this property must address the problem of disentangling the 'true', negative dependence induced by the dynamics of incentives from the 'spurious', positive contagion generated by unobserved heterogeneity.

This problem is a manifestation of a general question, namely distinguishing heterogeneity and state dependence. This issue has been abundantly discussed in the labor literature since the seminal contribution of Heckman and Borjas (1980). An interesting side aspect of our research, thus, is that it establishes a link between an existing literature in labor economics and questions that arised recently in applications of contract theory.

⁵An additional (and standard) difficulty comes from the fact that we observe claims, not accidents, and that the decision to file a claim is endogenous. See Chiappori and Salanié (1998) for a precise discussion of this problem.

3. Testing for moral hazard

In most empirical studies in insurance, data are drawn from the files of one (or more) insurance companies. Many relevant characteristics of the driver (age, gender, place of residence, seniority, type of job, ...) and the car (brand, model, vintage, power, ...) are used by companies for pricing purposes. All these are available to the econometrician as well. The same is true for the characteristics of the contract (type of coverage, premium, deductible,...). Finally, each accident - or more precisely each claim - is recorded with all the relevant background information.

The main differences between data sets can be traced back to the way past history is recorded. Existing situations can be gathered in three broad cases:

- In the most favorable situation, the exact date of each accident is recorded. Then the occurrence of an accident can be modelled in continuous time, using event-history models.
- Many experience-rating schemes can be implemented with information on the number of accidents in each contract year only. In such cases, insurance companies will often only provide researchers with individual counts of claims over the years. In some cases, information on whether at least one accident has occurred or not in any year (rather than the exact number of accidents in each year) is sufficient. Then, for each agent we only observe a sequence of 0s (for years without accidents) and 1s (years with accidents).
- Finally, the minimum information that is needed to implement a bonus-malus scheme may be even poorer. If all past accidents are treated symmetrically whatever their exact timing (as in our theoretical model), the computation of a bonus-malus coefficient only requires information on the total number of past accidents. In our model, an agent who has been driving for t periods and has had n accidents will be charged a premium of $\gamma^n \delta^{t-n}$ times her initial premium, whatever the exact timing of each of the accidents. In this case, a single draw from an insurance company's files may only give a cross-section of total counts of accidents for a group of clients that has been driving for periods of varying length. Dynamics can only be studied by comparing across individuals of different (driving) seniority.

In each of these three cases, the dynamics of accidents can be used to test for the presence of moral hazard, against the null that the accident probability

does not depend on the agent's incentives and only evolves according to some predetermined law (possibly depending on observables, such as age of the driver, age of the car, and others).

3.1. Heterogeneity versus moral hazard in continuous-time event-history data

In the first case, the essence of the test is clear. Under moral hazard, the hazard rate of an accident, conditional on observable and unobservable heterogeneity, should be steadily increasing throughout any period without an accident and drop discontinuously whenever an accident occurs. Under the null, however, the hazard rate should not change after an accident. This can either be tested parametrically or non-parametrically. Denote the number of claims up to and including time t by $N(t)$ and let $X(t)$ be some vector of observable covariates (age, gender, etc.) at time t . Abbring et al (forthcoming) assume that the intensity θ of claims, conditional of the claim history $\{N(u); 0 \leq u < t\}$ and the covariate history $\{X(u); 0 \leq u \leq t\}$ up to time t takes the form

$$\theta(t|\lambda, \{N(u); 0 \leq u < t\}, \{X(u); 0 \leq u \leq t\}) = \lambda \beta^{N(t-)} \psi(t) e^{X(t)'\gamma},$$

where ψ is a fully non parametric baseline hazard function, $\beta > 0$ a scalar parameter, λ a nonnegative unobservable covariate reflecting unobserved heterogeneity, and γ a vector of parameters. Note that $N(t-)$ is the number of claims up to, but not including, time t . Thus, the parameter $\beta > 0$ captures true occurrence dependence effects. In the bonus-malus system described above, moral hazard leads to a decline in the intensity of claims with the number of previous claims ($\beta < 1$). Without moral hazard, we expect $\beta = 1$. Distinguishing these cases (testing), and estimating β , is the focus of the empirical analysis. Statistical tests are developed and applied to a French sample of 79,684 contracts, of which 4,831 have one claim in the contract year and 287 have two claims or more. The null ($\beta = 1$) cannot be rejected at any conventional level, suggesting that moral hazard is not a major problem in the data under consideration.

3.2. Testing for moral hazard from sequences of accident counts

When only the total numbers of accidents by year are known, we can develop and apply similar methods for testing occurrence dependence in panel count data. Here, we focus on the more challenging case in which we only observe an annual sequence of 0s and 1s, corresponding to respectively years without and years

with at least one accident, for each agent. Econometric procedures for testing for occurrence dependence on such data have been developed by Heckman (1978, 1981 a,b), Honore (1993), Kyriazidou (1997) and Honore and Kyriazidou (2000). They rely on the assumption that each agent's accident probability remains constant throughout the observation period (stationarity).⁶ 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, $A = (1, 0, 0, 0)$ and $B = (0, 0, 0, 1)$, where a 1 (resp. 0) denotes the occurrence of an accident (resp. no accident) during the corresponding year. In the absence of moral hazard, and assuming away learning phenomena, the probabilities of observing either of the two sequences should be exactly the same; 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 accidents happens after an accident, hence when the premium, and consequently the marginal cost of future accidents and the incentives to take care are maximum. In other words, for a given average frequency of accidents, the timing of the occurrences can provide valuable information on the importance of incentives.

The test described here assumes stationarity. The analogy with the methods for continuous-time data of Abbring et al. (forthcoming) discussed earlier suggests that tests can be developed that are informative on moral hazard even if individual accident probabilities may change over time for external reasons. Richer panel-count data, that do not only record whether an accident has occurred at all but also how many accidents have occurred in any year, may be helpful here. This is on our research agenda.

3.3. Testing for moral hazard from total number of accidents only

Even in the case in which information is minimal - i.e., in which only the total number of past accidents is known for each agent in the insurer's data base - it is still possible to test for moral hazard. In this case, we essentially have a cross-section of total accident counts over the periods that agents have been driving (seniority). Under additional stationarity assumptions, one can exploit the variation in seniority in the data set to test for moral hazard. Specifically, assume that (i) individual accident probabilities are constant over time, and (ii) the distribution of unobserved heterogeneity in the population is identical across

⁶In the duration model, this would correspond to a constant ψ .

cohorts (seniority levels). That is, conditionally on observables, the no accident probability p is distributed among the drivers of any given seniority according to some distribution μ that is identical across seniorities. The idea of the test, as developed by Chiappori and Heckman (2000), is the following. Under the null of no moral hazard, for any driver with seniority t and no accident probability p , the probability of having no accident throughout the observation period is p^t . Hence the proportion of drivers with no accident throughout the period is, under the null, equal to

$$m_t = \int p^t d\mu(p),$$

i.e., to the t -th moment of the distribution. It follows that the numbers $m_1, m_2, \dots$ must, under the null, be the successive moments of the same distribution, which generates a first set of restrictions (see Heckman 1978 and in a different context Chiappori 1985). In addition, one can see that again under the null, the proportion, within the subpopulation of seniority t , of agents with exactly one accident is

$$\begin{aligned} m_t^1 &= \int t p^{t-1} (1-p) d\mu(p) \\ &= t (m_{t-1} - m_t). \end{aligned}$$

This provides a set of simple, linear restrictions involving three statistics, namely m_{t-1}, m_t and m_t^1 . Additional restrictions can be derived involving higher numbers of accidents. An analogous analysis for the moral hazard case is required to judge the power of tests based on these restrictions. These are topics for future research. Note that, in any case, these tests involve a comparison of *disjoint* subpopulations and heavily exploit stationarity assumptions.

3.4. Testing for adverse selection

In the three cases considered above, the null (no moral hazard) is consistent with the presence of unobserved heterogeneity, whatever its type. Such heterogeneity reflects the impact of any variable that is *not* observed by the insurance company (and therefore the researcher), whether it is known by the insuree or not. In other words, one does not, under the null, distinguish between *adverse selection* and symmetrically imperfect information. Testing for adverse selection (and particularly asymmetric learning) requires analyzing the joint process followed by accidents and contractual choices.

In a dynamic setting, adverse selection can be modelled in various ways. One way is to assume that each agent is characterized by some constant parameter reflecting her 'quality' as a driver, which is known by the agent but not by the insurer at the beginning of the relationship. In this setting, adverse selection can be tested using the simple, cross-sectional approach described in introduction. In the case of automobile insurance, most existing analyses fail to find positive conditional correlation, at least on populations of young drivers. This suggests that adverse selection, if any, is not adequately described by the 'fixed quality parameter' story.

A more complex but also more convincing version relies on the asymmetric learning argument sketched in the introduction. There, adverse selection gradually emerges during the relationship. A natural empirical strategy is to study the causal relationship between the sequences of accidents and contract choices (or amendments). In particular, agents who learn their risk is above average are more likely to switch to a contract entailing a more comprehensive coverage. The previous (heterogeneity versus occurrence-dependence) perspective must then be extended to a two-dimensional process. Again, this will be the topic of future work.

References

Abbring, J.H., P.A. Chiappori and J. Pinquet (forthcoming), 'Moral Hazard and Dynamic Insurance Data', *Journal of the European Economic Association*.

Cardon, J., and I. Hendel (1998): 'Asymmetric Information in Health Insurance : Evidence From the National Health Expenditure Survey', Mimeo, Princeton University

Chiappori, P.A. (1985), "Distribution of Income and the Law of Demand", *Econometrica*, 53, 109-127.

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

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 moral hazard on dynamic insurance data: theory and econometric tests', Mimeo, University of Chicago.

Chiappori, P. A., B. Jullien, B. Salanié and F. Salanié (2002), "Asymmetric Information in Insurance: Some Testable Implications", mimeo, University of Chicago.

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. and B. Salanié (2003), 'Testing Contract Theory: a Survey of Some Recent Work', in *Advances in Economics and Econometrics - Theory and Applications, Eighth World Congress*, M. Dewatripont, L. Hansen and P. Turnovsky, ed., Econometric Society Monographs, Cambridge University Press, Cambridge, 115-149.

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., M. Maurice, J. Pinquet and C. Vanasse (2001), "The Role of Memory in Long-Term Contracting with Moral Hazard: Empirical Evidence in Automobile Insurance", mimeo, University of Montreal.

Dionne, G. and C. Vanasse (1992), "Automobile Insurance Ratemaking in the Presence of Asymmetrical Information", *Journal of Applied Econometrics*, 7, 149-65

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.

Heckman, J.J. (1981a): "Statistical models for discrete panel data", *Structural Analysis of Discrete Panel Data with Econometric Applications*, C. Manski and D. McFadden, ed., Cambridge, MIT Press.

Heckman, J.J. (1981b): "Heterogeneity and State Dependence", *Studies of Labor Markets*, S. Rosen, ed., Chicago, The University of Chicago Press.

Heckman, J.J. and G.J. Borjas (1980), "Does unemployment cause future unemployment? Definitions, questions and answers from a continuous time model

of heterogeneity and state dependence”, *Economica*, **47**, 247–283.

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

Holly, A., L. Gardiol, G. Domenighetti and B. Bisig (1998), “An Econometric Model of Health Care Utilization and Health Insurance in Switzerland”, *European Economic Review*, 42, 513-522.

Honore, B (1993), “Orthogonality Conditions for Tobit Models with Fixed Effects and Lagged Dependant Variables”, *Journal of Econometrics*, 59, 35-61.

Honore, B, and E. Kyriazidou (2000), “Panel Data Discrete Choice Models with Lagged Dependant Variables”, *Econometrica*, **68**, 839-74.

Kyriazidou, E. (1997), ”Estimation of Panel Data Sample Selection Models”, *Econometrica*, **65**, 1335-64.

Poterba, J., and A. Finkelstein (2003), ”Adverse Selection in Insurance Markets: Policyholder Evidence From the U.K. Annuity Market”, Mimeo, MIT.

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