EMPIRICAL WORK ON ASYMMETRIC INFORMATION IN INSURANCE MARKETS[†]

Modeling Competition and Market Equilibrium in Insurance: Empirical Issues

By Pierre-André Chiappori and Bernard Salanié*

In the last decade or so, numerous papers have been devoted to empirical investigations based on contract theory. Many contributions use insurance data, and specifically files provided by firms. A typical paper would analyze the relationship between individual characteristics, the contracts chosen and the corresponding "outcome," as measured by claims.[1]

The natural next step in this research agenda is to model empirically market equilibrium on insurance markets. Empirical models of competitive insurance markets are important in many respects. First, such models are an indispensable first step for the empirical analysis of existing markets. The discussion of optimal pricing strategies or the definition of new insurance contract would greatly benefit from such models. From a policy perspective, the design of any regulation requires estimating its likely impact on the market allocation. For instance, while a ban on specific pricing options (based, say, on gender or age) is often advocated on ethical grounds, a precise assessment of its impact on insurance markets is needed before any decision is made; and an empirical model is required to provide such an assessment.

From a purely theoretical perspective, any description of insurance markets that aims at a modicum of realism needs to come to terms with a host of complex features (horizontal differentiation of products, unobserved heterogeneity of preferences, frictions of various types), the theoretical analysis of which may be forbiddingly complex. A simple model that can be solved or at least numerically simulated may, in that case, be particularly helpful. Finally, a tractable model of insurance equilibrium can be used to run experiments, which should help us understand individual behavior in such strategic settings as competition under asymmetric information.

On the other hand, modeling insurance markets raises several theoretical and empirical issues, starting, of course, with the well-known pitfalls in modeling equilibrium in contracts. The goal of the present paper is to discuss these problems and summarize the knowledge acquired so far. We successively discuss modeling of the demand side, the supply side, and the equilibrium itself.

I. Modeling Demand

Modeling the demand for insurance requires understanding various aspects: the nature of the risk, the characterization of the contracts traded, an assessment of the various frictions involved (horizontal differentiation for instance), the nature of information asymmetries (if any), and, finally, knowledge of the joint distribution of risk, risk aversion, and income in the population of insurees.

A. *Risk*

The simplest theoretical models involve only two contingencies: a negative outcome occurs

[1] One may mention, among many others, Chiappori and Salanié (1997, 2000), Amy Finkelstein and James Poterba (2004) or Alma Cohen and Liran Einav (2007). See Chiappori and Salanié (2002) or Salanié (2003) for early surveys.

with some given probability and implies a given loss with a money equivalent L. This basic framework can be extended in various directions. The loss may be modeled as a continuous variable. In some cases, several correlated risks have to be considered simultaneously (e.g., liability and damage risks in car insurance.) More complex issues arise when utility is state dependent, since the risk then cannot be considered as purely monetary. For instance, the benefits derived from a life insurance contract depend on the current utility, for a person, of a future transfer to the offspring after the person's death. The underlying intertemporal rate of substitution/altruistic motive may be hard to assess, let alone to distinguish from risk aversion. Several lines of business, including health insurance, raise similar questions.

Moreover, many insurance contracts have dynamic features, which raises specific issues. Annuity contracts differ in their degree of front-loading, a property that can be used to screen insurees in a context of asymmetric information on longevity (an aspect discussed by Finkelstein, Poterba, and Casey Rothschild 2006, and Eytan Sheshinski 2007). In this context, the intertemporal distribution of payments implied by each contract must be modeled with care. Many contracts can be renewed each year, and the dynamics involved can be quite complex. For instance, with asymmetric learning, whereby insurees acquire throughout time more information than insurers about their type, changes in the contract purchased may carry information about the agent's updated belief. In a similar vein, issues related to preexisting conditions are probably the main problem facing private provision of health insurance.

Each of these issues requires care when modeling contracts. In the simplest framework, a contract is simply a (premium, deductible) pair. With a continuum of possible losses, or in an explicitly intertemporal framework, a contract is represented as a continuous (or dynamic) schedule; the characterization of optimal or equilibrium contracts therefore requires optimal control theory or dynamic games. Moreover, in a dynamic framework, commitment issues become crucial. Finally, a specific but important aspect of insurance econometrics is the distinction between accident and claims. With experience rating, for instance, the decision to file a claim after an accident is endogenous and will in general depend on the contract, a feature that must in many cases be taken into account.

B. *Preferences*

Individual risk aversion is a key determinant of the demand for insurance. Several papers have recently been devoted to the estimation of the corresponding parameters (see, for instance, Robert B. Barsky et al. 1997; Cohen and Einav 2007; Chiappori and Monica Paiella 2006 (CP below)). Some clear conclusions emerge from this line of work. First, risk aversion is extremely heterogeneous; even conditional on observable characteristics, risk attitudes differ much more across individuals than risk itself. Second, constant relative risk aversion provides a reasonably good approximation of individual attitude toward risk, at least in an expected utility setting. In particular, CP, using panel data on portfolio composition, show that changes in income result in compensating variations in portfolio composition that leave the fraction invested in risky assets constant. The same studies suggest that relative risk aversion (RRA) is negatively correlated with income, although the correlation is very small; therefore, assuming that RRA is distributed independently of income is (surprisingly) an acceptable approximation. Finally, the distribution of RRA seems to be fit rather well by a lognormal, although the coefficients governing this distribution vary between studies; for instance, the mean of log RRA is estimated at 0.5 by CP, while it is calibrated at almost 2 by Barsky et al. Interestingly, the standard deviation is close to 1 in both studies.

In many studies, however, a constant absolute risk aversion (CARA) approximation is used. Such an approximation is convenient when no data are available on individual income; it is justified when the risk is small with respect to income. Note, however, that if we accept the CRRA context, a CARA approximation introduces a correlation between ARA and income y, since $\log(ARA) = \log(RRA) - \log y$.

Once individual risk, contracts, and risk aversion have been modeled, one can compute for each potential customer i the certainty equivalent of each available contract k, say D_k^i for $k = 1, \ldots, K$, as well as that of no insurance, denoted D_0^i. Given the joint distribution of these three determinants, the investigator then

obtains the distributions of these certainty equivalents[2].

C. *Demand*

In order to calibrate individual demand, one must allow for various frictions. In particular, the law of one price is clearly violated, even for apparently quite homogeneous products such as the liability guarantee in car insurance. A random utility model is a natural way to take this into account. Insurees are supposed to choose the option that maximizes a sum $D_k^i + \varepsilon_k^i$, where the ε_k^i are i.i.d. random perturbations that follow a type-1 extreme value distribution.

Then, the probability of agent i choosing contract k for $k = 0, \ldots, K$ is

$$P_k^i = \frac{e^{\gamma D_k^i}}{\sum_s e^{\gamma D_s^i}},$$

where $1/\gamma$ is a "viscosity" parameter. This form can be further enriched by assuming that

$$P_k^i = \frac{A_k e^{\gamma D_k^i}}{\sum_s A_s e^{\gamma D_s^i}},$$

where the (unknown) A_k parameters capture differing brand recognitions, density of insurance retail outlets and so on. Note that A_k may be endogeneized—i.e., "purchased" by the insurer through advertising for instance. Similarly, the viscosity coefficient $1/\gamma$ may be product- or company-specific, and may vary with insurees' observed or unobserved characteristics. A particularly interesting situation obtains when demand from long-time clients is less elastic; such "customer loyalty" or switching costs generates rich pricing dynamics, since future profits are correlated with current market shares.

Finally, aggregate demand and the average accident rate for a given contract obtain by summing the probabilities of losses of the insurees who choose this contract, weighted by the probability that they do so, over their observed characteristics and unobserved heterogeneity.

---

[2] As in portfolio choice, the assumption that preferences are CARA and disturbances are normally distributed is attractive here, as it makes computations easier; but even so, numerical integration must be used.

D. *Information*

A final and crucial component of the model relates to the information available to each consumer. It is natural to assume that risk aversion is the client's private information. Regarding risk, two polar cases can be considered—but they hardly exhaust the possibilities. In one case, individuals know exactly their risk class, and the insurer cannot use that information, say, because regulation prohibits the use of specific indicators (race, gender, age, …). Then, the model entails bidimensional adverse selection. Alternatively, we may assume that individuals have no informational advantage over insurers; then, adverse selection bears on risk aversion only, and is relevant only to the extent that competition is imperfect (Chiappori et al. 2006). Intermediate cases (e.g., agents receive a private signal about their accident probability) can be considered as well. Finally, the introduction of moral hazard requires, in addition, modeling a prevention cost within the utility function of each consumer. This raises very thorny issues: is the cost homogenous, publicly known, correlated with risk, or with risk aversion? To our knowledge, little progress has been made in that direction.

II. Modeling Supply

While modeling demand on insurance markets is a challenge, modeling supply is much easier. The space of product characteristics is usually low-dimensional; e.g., in car insurance a product is mostly defined by whether it carries damage insurance, liability, and a deductible and premium (and perhaps some minor variants on fire and theft insurance). Similarly, many elements of the cost structure of insurers are fairly simple: cost per contract, per product, and per client, and variable costs when payments are made. Most of these costs can be estimated, calibrated, or simulated. Costs of entry and exit, as elsewhere, are a more difficult problem of course.

On the other hand, dynamic issues in insurers' strategies are quite interesting, for several reasons. As mentioned above, switching costs can be readily be introduced by assuming that the viscosity parameter $1/\gamma$ varies with a client's seniority. In some cases, the parameter can moreover be directly influenced by the insurer

(say, through advertising). Commitment issues are more complex. When an insurance company sells a contract that commits it for the long term (e.g., a life insurance policy with a guaranteed renewal clause), it is de facto offering an option, the pricing and hedging of which raise delicate problems. For instance, selective attrition may be a serious concern; and computing provisions is known to be a difficult problem in life insurance. Finally, regulation typically constrains the dynamics of capital and profit, since past losses, insofar as they reduce available capital, limit the current size of operations.

III. Modeling Equilibrium

When modeling market equilibrium in an insurance context, a clear distinction should be drawn between exclusive and nonexclusive contracts. When contracts are not exclusive (as in the case of annuities), pricing schemes that are convex in coverage cannot be used, since the insurees can arbitrage over them. This affects the strategic nature of competition, particularly in the presence of asymmetric information. In that case, screening can take other forms, such as specific designs of contract dynamics (see Rothschild 2007).

The case of exclusive contracts is more familiar, if only because of the large theoretical literature devoted to this situation. Still, the mere choice of an equilibrium concept is not clear. For instance, should equilibrium require nonnegative profit for each contract, or are cross subsidies allowed? Should one refer to equilibria à la Rothschild and Stiglitz, à la Wilson/Spence/Miyazaki, à la Hellwig? And what should one do when the equilibrium fails to exist? From a theoretical perspective, the existence issue is especially difficult in the framework described above, because of its sheer complexity. Adverse selection is multidimensional (on risk and risk aversion), and the distribution of risk aversion at least is typically continuous. The various frictions introduced (finite elasticity, cost per contract and product, etc.) tend in general to make the existence of equilibrium more likely; still, robust examples of nonexistence can be found. An analytic derivation of the existence and uniqueness of the equilibrium (let alone its qualitative properties) is in general out of reach. The model can then be used to derive the results by simulations instead.

IV. A Simulation Game

Finally, a simulation game can be constructed using the ideas presented above; the reader is referred to Chiappori and Salanié (2008) for a precise description. The main features can be summarized as follows. Each player is an insurance company, faced with a demand simulated by the computer. Risk is modeled à la Mossin, i.e., as the occurrence of a single monetary loss. Various risk classes (corresponding to specific accident probability) are defined; in the simplest version, there are two of them, although the extension to continuous distributions is straightforward. Within each class, a random sample of risk aversions is drawn from some distribution, thus generating a set of consumer types; each type can have a specific weight. The behavior of a given type follows the ideas presented above.

In the static version of the game, a first stage is organized assuming symmetric information; in practice, each risk class is offered specific contracts. At each round, insurers compete in "contracts," each of which is defined by a premium and a deductible; the model computes market shares and profits for each contract. Players/insurers are each provided with a simulation tool based on the true characteristics of the virtual market; i.e., they can enter "guesses" about the other players' strategies and test the market shares and profits generated by any contract they may offer. This feature allows them to concentrate on the strategic aspects of the game. The second stage follows the same logic, but differentiation based on risk is prohibited; companies may still offer several contracts, but each contract is freely available to any consumer. Finally, various dynamics can be introduced, along the lines discussed above. The game is currently being tested on a sample of players.

REFERENCES

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro. 1997. "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study." *Quarterly Journal of Economics*, 112(2): 537–80.

Chiappori, Pierre-André, Bruno Jullien, Bernard Salanié, and François Salanié. 2006. "Asymmetric Information in Insurance: General

Testable Implications." *The RAND Journal of Economics*, 37(4): 783–98.

Chiappori, Pierre-André, and Monica Paiella. 2006. "Relative Risk Aversion Is Constant: Evidence from Panel Data." Unpublished.

Chiappori, Pierre-André, and Bernard Salanié. 1997. "Empirical Contract Theory: The Case of Insurance Data." *European Economic Review*, 41(3): 943–50.

Chiappori, Pierre-André, and Bernard Salanié. 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy*, 108(1): 56–78.

Chiappori, Pierre-André, and Bernard Salanié. 2002. "Testing Contract Theory: A Survey of Some Recent Work." In *Advances in Economics and Econometrics Vol 1*, ed. M. Dewatripont, L. Hansen, and S. Turnovsky, 115–49. Cambridge: Cambridge University Press.

Chiappori, Pierre-André, and Bernard Salanié. In progress. "Playing an Insurance Game." Unpublished.

Cohen, Alma, and Liran Einav. 2007. "Estimating Risk Preferences from Deductible Choice." *American Economic Review*, 97(3): 745–88.

Finkelstein, Amy, and James Poterba. 2004. "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market." *Journal of Political Economy*, 112(1): 183–208.

Finkelstein, Amy, James Poterba, and Casey Rothschild. 2006. "Redistribution by Insurance Market Regulation: Analyzing a Ban on Gender-Based Retirement Annuities." National Bureau of Economic Research Working Paper 12205.

Rothschild, Casey. "Adverse Selection, Linear Pricing, and Front Loading in Annuity Markets." Unpublished.

Salanié, Bernard. 2003. "Testing Contract Theory." *CESifo Economic Studies*, 49(3): 461–77.

Sheshinski, Eytan. "Refundable Annuities (Annuity Options)." Unpublished.