Equilibrium in Insurance Markets: An Empiricist's View

Bernard Salanié

Columbia University

Limassol, 19 September 2016
EGRIE Conference

In memoriam

Phoebus J. Dhrymes (1932-2016)
(memorial reception at Columbia: Sep 27)

The Goals

Empirical models of (imperfectly) competitive insurance markets would be useful

- for insurers, to simulate market functioning and to forecast the profitability of a new contract
- for regulators when they consider the antitrust and welfare impacts of restrictions to pricing for instance
- for economists, to understand many important questions.

This Talk

- considers progress made since (for instance!) Chiappori and Salanié (*AER P&P* 2008)
- relies heavily on joint work with Chiappori, Escanciano, Gandhi, Jullien, Pouyet, F. Salanié, Yıldız (who bear no responsibility, etc)
- happily skips over the "existence problem" (competition is "imperfect enough")
- and shamelessly underweighs other relevant work.
- We have also learned much from *non-equilibrium* empirical studies, typically based on data from a single company:
    - welfare impact of asymmetric infomation: e.g. Einav–Finkelstein–Levin *ARE* 2010
    - and an exploding literature on health insurance (e.g. Finkelstein and coauthors, and Handel–Hendel–Whinston *Eca* 2015).

A Brief Summary

- much progress has been made in estimating the distribution of preferences of insurees
- some on perceptions of risk, especially as they overlap with preferences
- steady progress in our ability to quantify the impact of moral hazard and adverse selection
- but most of this work is on data from a single insurer;
- the data required for a market-wide model is still very hard to obtain—**suggestions welcome!**

Outline

Start from a standard empirical IO model, then introduce ingredients specific to insurance markets.

Demand model: discrete choice of consumers among insurance contracts.

Ingredient 1: competition in contracts, not simply in prices of predefined varieties.

A contract can be viewed as a nonlinear quantity/price mapping, at a minimum:
loss $L \to$ reimbursement $R(L)$, for a premium P
e.g. $R(L) = \max(L - D, 0)$ for a straight deductible contract.

In more complicated examples there are several types of risk, quality of service matters etc, and dynamic aspects (experience rating, lifetime aspects...)

The crucial point

In most empirical IO work the range of products offered is considered given; it would be very hard to endogeneize as the space of product characteristics is large.

In insurance the product is typically easily described;

endogeneizing the equilibrium product range is both possible and very important for welfare evaluations for instance.

Demand side

The two main drivers of the demand for insurance: perceived risks, and attitudes towards risk.
Ingredient 2: analyze the preferences of consumers for risky prospects.
Ingredient 3: we need to understand how insurees perceive their risks; and what control they may have over them (moral hazard)

Other ingredients may also come into play: dynamics/learning, ex post moral hazard/fraud, etc;

and switching costs (Handel *AER* 2013): a potential headache with cross-sectional data.

Supply side

Empirical IO is usually wary of modeling supply explicitly:

- information on cost functions and their components (input prices and use) is sparse.
- preferred option: use orthogonality between a residual of demand and an instrument.

Insurance is (sometimes) a more favorable case:

- basically, costs are reimbursement+administrative+financial
- *Ingredient 4:* and since the product space has low dimension, we can hope to endogeneize product characteristics.

A Barebones Demand Model for Exclusive Insurance

(Potential) insuree i has CARA index σ_i, faces loss L_i with probability p_i, chooses among straight deductible contracts with (deductible, premium) (D_{kj}, P_{kj}) where k=insurer and j=contract.

Assume $L > \max D_{kj}$ for simplicity, and no dynamics, and no moral hazard.

The certainty equivalent of contract (k, j) for i is

$$C^i_{kj} = -P_{kj} - \frac{1}{\sigma_i}\log\left(1 - p_i + p_i \exp(\sigma_i D_{kj})\right)$$

$$\left(\simeq -P_{kj} - p_i D_{kj} - \frac{\sigma_i}{2} p_i(1-p_i)D^2_{kj} - \frac{\sigma_i^2}{6} p_i(1-p_i)(1-2p_i)D^3_{kj}\right)$$

Choice between contracts

A reasonable specification for a random utility model could be

$$u^i_{kj} = C^i_{kj} + \xi_{kj} + \varepsilon_{ijk}$$

with ε_{ijk} idiosyncratic (say "multinomial logit": iid standard type I extreme value)

and ξ_{kj} = common preferences for contract (k, j),

- unobserved by the econometrician
- on which the insurers may have some information relevant to their strategies, **hence endogenous.**

Reasonable (?) simplification: $\xi_{kj} = \xi_k$ only depends on the insurer.

Risk and risk-aversion

σ_i and p_i are like the random coefficients of the mixed multinomial logit model of empirical IO

but we usually have more information than IO people do:

- insuree-level data, which helps with the distribution of σ_i
- claims data, which informs us on p_i (actual, not perceived...).

In empirical IO parlance, we can do "micro BLP" (Berry–Levinsohn–Pakes *JPE* 2004)
rather than simply "macro BLP" (Berry–Levinsohn–Pakes *Eca* 1995).

What we have learned on risk preferences (a personal summary)

(from Barseghyan et al *AER* 2013, Chiappori et al 2016, and a lot of survey-based evidence starting with Barsky et al *QJE* 1997)

- if we stick with expected utility,
  - CRRA fits better than CARA—but requires information on "wealth"
  - we need a lot of heterogeneity to make sense of choices
- if we allow for nonlinear probability weightings, heterogeneity becomes less important.

And on actual vs perceived risk

Without data on both, it is hard to distinguish non-expected utility from misperceptions.
Such data is only available in small surveys (but see Finkelstein-McGarry *AER* 2006)
→ what follows is from the literature about preferences.

There is evidence that small risks are overweighted (Barseghyan et al *AER* 2013, *QE* 2016),
with important heterogeneity (Gandhi–Serrano-Padial *REStud* 2014)
probably much less distortion for larger probabilities — hard to see on insurance data.

We certainly need to include some flexibility in modeling weighted p_i.

What of asymmetric information?

Adverse selection: it's in there, via the heterogeneity in σ and p.
Implicit equilibrium concept: Nash in contracts, that "is" Rothschild-Stiglitz.

Moral hazard: a more complicated animal.
In the barebones model, maybe now the certainty equivalent of (P, D) is

$$C(P,D) = \max_e \left(-P - \frac{1}{\sigma_i} \log\left(1 - p_i(e) + p_i(e) \exp(\sigma_i D)\right) - e \right)$$

Say $p_i = q(p_i^0, e)$; the optimal effort is $e^*(D, \sigma_i, p_i^0)$,
the optimal value is $C(P, D) = -P - F(\sigma_i, p_i^0, D)$,
with F increasing in p_i^0, D and decreasing in σ_i.

With luck we have a proxy of the ex post risk $p_i^* = q(p_i^0, e^*(D, \sigma_i, p_i^0))$.
This does not seem very promising,

unless we have auxiliary information about $q(\cdot, \cdot)$ and p_i^0, or perhaps effort.

Do we need asymmetric information?

Pouyet-Salanié-Salanié *BEJTE* 2008: asymmetric information on preferences σ_i *alone* does not matter if competition is perfect

But of course

- it isn't;
- and there is also risk p_i;
- and possibly moral hazard.

Why don't we test for asymmetric information first. . . ?
Chiappori-Salanié *JPE* 2000, Chiappori-Jullien-Salanié-Salanié *Rand* 2006, and many others.

see Chiappori-Salanié *Handbook of Insurance* 2014 for a summary.

A refresher on testing for AI

Take a contract (P_1, R_1), and a contract (P_2, R_2) that offers more coverage: $(R_2(L) - R_1(L))$ is non-decreasing.
Then

- if risk-averse i chose contract 1, then (up to choice errors, and even with non-expected utility)

  $$P_2 - P_1 \geq \tilde{E}_{i1}(R_2(L) - R_1(L))$$

  where $\tilde{E}_{i1}$ is "i's perceived distribution of her losses under contract 1."
- if profits $\pi_2 \geq \pi_1$ (the "normal" case), then we can add the two inequalities to generate a testable inequality
- and if $\pi_j = P_j - \ldots$ the premia cancel out (a good thing since premia are only observed for chosen contracts.)

Predicts positive correlation of risk and coverage *only in simple cases*
but gives a useful test in many more.

Moral hazard vs adverse selection

Hard to disentangle, both predict "positive correlation."

Testing for moral hazard is hard because it affects both choice of contract (the *treatment equation*) and claims (the *outcome equation*).

Escanciano-Salanié-Yıldız 2016: use **exogenous** differences in the menu of contracts offered.
Under the null of no moral hazard, such differences should not affect the distribution of claims.
Also gives an evaluation of the impact of moral hazard (for the difference considered.)

Most differences in the contracts offered are not exogenous; some differences induced by regulation are more likely to be.

Non-exclusive insurance

Coverage can add up across insurers: e.g. life insurance.

On the theory side:
see Attar–Mariotti-F. Salanié *Eca* 2011, *TE* 2014, WP 2016
non-exclusivity makes convex pricing wrt coverage difficult *within one contract*.

only the *Jaynes-Hellwig-Glosten* allocation can be an (unregulated) equilibrium: each agent buys basic coverage, and some complement it to various degrees by buying complementary insurance.

For the empiricist: within-firm pricing can give a misleading view of market equilibrium
e.g. low-coverage contracts may attract high-risk types who buy several contracts.

Having comprehensive data about sources of coverage is crucial (and hard!)

Data

> *"Data!data!data!" he cried impatiently. "I can't make bricks without clay."*
>
> Arthur Conan Doyle, *The Adventure of the Copper Beeches*.

A wishlist: (so far unfulfilled)

- we need data on an insurance market with simple contracts (like car insurance)
- with insuree covariates, and claim data
- *on a representative section of firms*, with at least the main large competitors.