Empirical Models of Demand for Insurance

Liran Einav (Stanford and NBER)

Cowles Lunch Talk, Yale University
September 18, 2013

Why insurance?

- Insurance markets make a good target for economists
  - Large and important
  - Potential market failures
  - Common government interventions via regulation or provision
  - Large and detailed data sets are increasingly available
- Extremely influential theory work on the potential for market failures (e.g., Arrow 1963; Akerlof 1970; Rothschild and Stiglitz 1976)
- Early influential empirical work on insurance in the 1990s (e.g., Chiappori and Salanie 2000) mostly focused on testing. Underlying issues more subtle:
  - On the demand side, increasing evidence of important and relevant individual heterogeneity
  - On the supply side, increased concerns about "lemon dropping" and "cherry picking" and the more general need to better understand competition, market power, and supply chains in these "contract/selection markets"

Plan for today's talk

As the title suggest, I will mostly focus on demand and say very little about competition or market equilibrium:

- To fix ideas, start with a general (fairly standard) framework of demand in insurance markets
  - Emphasize the key aspect that distinguishes demand in insurance: insurance providers directly care about who bought the product and how he/she behaves after the purchase
- Continue by describing three specific examples of demand models in different insurance contexts, all starting with a model of "realized utility"
- Connect these models to models that are based on indirect utility, which are more common in "traditional" IO
- Relate demand models to firm's problem of pricing and plan design

(see also our review at Einav, Finkelstein, and Levin, 2010)

Plan for today's talk

- **To fix ideas, start with a general (fairly standard) framework of demand in insurance markets**
- Continue by describing three specific examples of demand models in different insurance contexts, all starting with a model of "realized utility"
- Connect these models to models that are based on indirect utility, which are more common in "traditional" IO
- Relate demand models to firm's problem of pricing and plan design

General framework: notation

- ζ: Vector of consumer characteristics (risk, preferences, income, information, etc.)
  - $\zeta = (x, \nu)$: Useful to decompose into observables x and unobs. ν
- (ϕ, p): Contract, i.e. coverage characteristics ϕ and a price (premium) p
- S: Set of possible outcomes (e.g., accident)
- A: Set of actions available to the consumer during coverage period (e.g., driving care)
  - Actions could be contingent, thus covering "ex-post moral hazard"

General framework: expected utility

- Adopt a standard expected utility framework
- Consumer's valuation of a contract (ϕ, p) is

$$v(\phi, p, \zeta) = \max_{a \in A} \sum_{s \in S} \pi(s|a, \zeta)\, u(s, a, \zeta, \phi, p)$$

- Let also $a^*(\zeta, \phi, p)$ denote the consumer's optimal behavior, and $\pi^*(\cdot|\phi, p, \zeta) = \pi(\cdot|a^*(\zeta, \phi, p), \zeta)$
- Convenient assumption: $v(\phi, p, \zeta)$ is additively separable in p
    - Equivalent to CARA in the standard framework

1. Changes in the premium do not affect consumer behavior a^* or the outcome probabilities π^*; helps with identification
2. Also leads to a natural choice of social welfare function

General framework: demand and cost

- Insurer (expected) costs:

$$c(\phi, \zeta) = \sum_{s \in S} \pi^* (s|\phi, \zeta)\, \tau(s, \phi)$$

  where insurance payments $\tau(s, \phi)$ depend on outcome and coverage

- Considering a set of insurance contracts J, consumer with characteristics ζ finds contract $j \in J$ optimal iff

$$v(\phi_j, p_j, \zeta) \geq v(\phi_k, p_k, \zeta) \text{ for all } k \in J$$

- An important point I will get back to: while the primitives include $\pi(\cdot)$, $u(\cdot)$, S, and A, for many questions of interest knowledge of $v(\cdot)$ and $c(\cdot)$ suffices

The distinguishing feature of insurance markets

- The key (I think) aspect which makes insurance (or credit) markets different from more traditional markets is that they are "selection markets"
- That is, unit costs depend on the *composition* of consumers (i.e. enrollee characteristics ζ) rather than just the quantity of consumers
  - Purchase and consumption/utilization are separated in time
  - Consumption/utilization enter the insurance company's profits
- Notes:
  - The first feature is not uncommon in many "traditional" markets (e.g., Dubin and McFadden, 1984)
  - Even the second is sometimes important, but perhaps not as first order as in insurance

Plan for today's talk

- To fix ideas, start with a general (fairly standard) framework of demand in insurance markets
- **Continue by describing three specific examples of demand models in different insurance contexts, all starting with a model of "realized utility"**
- Connect these models to models that are based on indirect utility, which are more common in "traditional" IO
- Relate demand models to firm's problem of pricing and plan design

Example 1: Deductible choice in auto/home insurance

- Based on Cohen and Einav (2007)
  - Main objective: estimates of risk aversion
- Auto insurance (in Israel) contracts characterized by a premium and a (per claim) deductible, (p, d)
  - Consider a binary choice between (p_l, d_l) and (p_h, d_h) with $d_l < d_h$ and $p_l > p_h$
- Key assumptions:
  - No behavioral effect of the contract (that is, risk invariant to coverage)
  - Claims are a realization of a Poisson process (given the contracts, can focus on claim counts and abstract from distribution of claim amount)
  - Individuals know their (objective) individual-specific claim rate λ

Deductible choice (cont.)

- Expected utility from contract $\phi = (p, d)$ over a *short* time period t:

$$v(p, d, w, \lambda, \psi) = (1 - \lambda t)u(w - pt) + (\lambda t)u(w - pt - d)$$

- Assume (heterogeneous) CARA utility $u(x) = -\exp(-\psi x)$
- Model is convenient/transparent as one can obtain a closed-form expression for the deductible choice:

$$Low\ \ Deductible \qquad \Leftrightarrow \qquad \lambda > \frac{\psi\left(p_l - p_h\right)}{\exp(\psi d_h) - \exp(\psi d_l)}$$

- Low deductible more appealing if higher risk or higher risk aversion

Deductible choice (cont.)

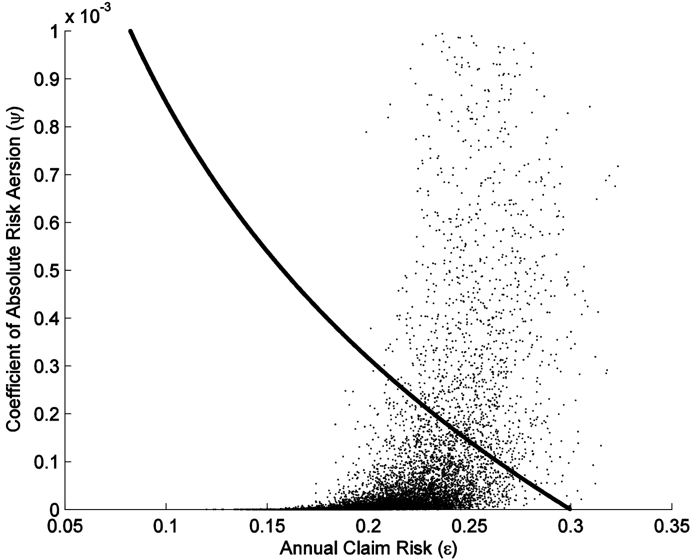

Example 2: Guarantee choice in the context of annuities

- Based on Einav, Finkelstein, and Schrimpf (2010)
  - Main objective: estimate welfare consequences of suboptimal pricing due to private information
- Annuity contracts (in UK) chosen by retirees
  - Focus on a setting with mandatory annuitization, so primary choice is the guarantee period
  - Contracts characterized by pairs of guarantee lengths (always 0, 5, or 10 years in this market) and corresponding annuity rates $z_{10} < z_5 < z_0$
  - Longer guarantees mean getting more if die quickly, but less if not
- Standard annuity framework:
  - Fully rational, forward looking, risk averse retirees
  - Retirees with stock of wealth face stochastic mortality with mortality hazard $\kappa_t(\alpha)$
  - Flow utility given by either utility from consumption (if alive) or bequest (if dead):

$$v(w_t, c_t) = (1 - \kappa_t)\, u(c_t) + \kappa_t b(w_t)$$

Annuity guarantee choice (cont.)

- Without annuity, optimal consumption is the solution to

$$\begin{aligned} V_t^{NA}(w_t) &= \max_{c_t \geq 0} \left[(1-\kappa_t)(u(c_t) + \delta V_{t+1}^{NA}(w_{t+1})) + \kappa_t b(w_t) \right] \\ s.t.\ w_{t+1} &= (1+r)(w_t - c_t) \geq 0 \end{aligned}$$

- With annuity and guarantee g, modified problem is

$$\begin{aligned} V_t^{A(g)}(w_t) &= \max_{c_t \geq 0} \left[\begin{array}{c} (1-\kappa_t)(u(c_t) + \delta V_{t+1}^{A(g)}(w_{t+1})) \\ +\kappa_t b(w_t + Z_t(g)) \end{array} \right] \\ s.t.\ w_{t+1} &= (1+r)(w_t + z_t(g) - c_t) \geq 0 \end{aligned}$$

where $Z_t(g) = \sum_{\tau=t}^{t_0+g} \left(\left(\frac{1}{1+r}\right)^{\tau-t} z_\tau(g) \right)$

- Optimal choice compares $V_0^{A(0)}(w_0)$, $V_0^{A(5)}(w_0)$, and $V_0^{A(10)}(w_0)$

Annuity guarantee choice (cont.)

- Focus on heterogeneity in (mortality) risk and preferences (for bequest). Longer guarantee is more appealing if higher risk or greater bequest preferences

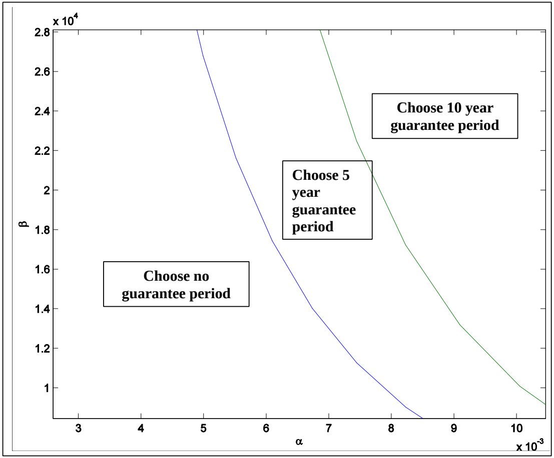

Example 3: Health insurance coverage choice

- Based on Einav, Finkelstein, Ryan, Schrimpf, and Cullen (2013)
  - Main objective: assess the importance of "selection on moral hazard"
- Choice among five employer-provided PPO health insurance plans
  - PPO means that tradeoff is solely financial (as in two previous examples); all choices associated with the same doctors, networks, and other restrictions
  - Contracts characterized by the generosity of the out-of-pocket function $c_j(m)$ and corresponding premium p_j
  - Key aspects of $c_j(m)$ are the annual (cumulative!) deductible level and the oop maximum (although rarely binding)
- Model designed to isolate three distinct determinants of an individual's coverage choice: health risk and risk aversion, which are similar to the two previous examples, as well as "moral hazard type," which is likely more important in the context of health insurance

Health insurance choice (cont.)

- An employee in a given year is characterized by
  - λ (monetized) health realization
  - $F_\lambda(\cdot)$ that govern health risk
  - ψ coefficient of absolute risk aversion
  - ω moral hazard type (price sensitivity)
- Two period model:
  - Period 1: given $(F_\lambda(\cdot), \omega, \psi)$, make optimal plan choice j^* from a plan menu J
  - Period 2: given plan j, health realization λ, and ω, make optimal utilization (spending) choice $m^* \geq 0$
- Note: totally abstract from the dynamics within the coverage year (which is the focus of two other papers of ours: one of which in tomorrow's IO seminar talk)

Health insurance choice: period 2

- Utility in period 2 given by:

$$u(m;\lambda,\omega,j) = \underbrace{\left[(m-\lambda) - \frac{1}{2\omega}(m-\lambda)^2\right]}_{h(m-\lambda;\omega)} + \underbrace{[y - c_j(m) - p_j]}_{y(m)}$$

- Optimal spending given by

$$m^*(\lambda,\omega,j) = \arg\max_{m\geq 0} u(m;\lambda,\omega,j)$$

- With linear contracts ($c_j(m) = c_j m$):

$$m^*(\lambda,\omega,j) = \lambda + \omega(1-c_j)$$

so ω ("moral hazard type") is the utilization difference between full and no insurance

Health insurance choice: period 1

- An individual valuation of plans has a CARA form over period 2's realized utility:

$$v_j(F_\lambda(\cdot), \omega, \psi) = \int -\exp(-\psi u^*(\lambda, \omega, j)) dF_\lambda(\lambda)$$

so optimal plan choice given by

$$j^*(F_\lambda(\cdot), \omega, \psi) = \arg\max_{j \in J} v_j(F_\lambda(\cdot), \omega, \psi).$$

- Optimal choice trades off higher up-front payment for more subsequent coverage
    - More coverage means both higher expected reimbursement and sheds off more risk
- Higher coverage more attractive for "higher" $F_\lambda(\cdot)$ (risk), higher ψ (risk aversion), and higher ω (moral hazard)

Plan for today's talk

- To fix ideas, start with a general (fairly standard) framework of demand in insurance markets
- Continue by describing three specific examples of demand models in different insurance contexts, all starting with a model of "realized utility"
- **Connect these models to models that are based on indirect utility, which are more common in "traditional" IO**
- Relate demand models to firm's problem of pricing and plan design

Connecting examples to models of indirect utility

- In all three examples, the model primitive was the realized utility, while typically in IO we model valuation (or indirect utility) directly
- Consider the auto insurance example. Choice given by

$$Low\ \ Deductible \qquad \Leftrightarrow \qquad \lambda > \frac{\psi\left(p_l - p_h\right)}{\exp(\psi d_h) - \exp(\psi d_l)}$$

  or, equivalently, the *incremental* "indirect utility" from upgrading to a low deductible coverage is

$$v_i = \frac{\lambda_i}{\psi_i}\left(\exp(\psi_i d_h) - \exp(\psi_i d_l)\right) - \left(p_l - p_h\right)$$

- Instead, one could directly (and perhaps more flexibly) specify v as a function of price and deductible

$$v_i = f\left(d_{i,LD}, d_{i,HD}; \theta_i\right) - \left(p_l - p_h\right)$$

  where θ_i is a vector of random coefficients, allowed to be correlated with subsequent outcomes to account for possible selection.

Possible benefits from modeling realized utility

What is the benefit, then, from modeling these "deeper" primitives?

1. Analyzing decision-making under uncertainty is one of the most fundamental topics in microeconomic theory
   - Economic theory provides structure, guidance, and sometimes restrictions on contract valuations
   - For the same reason, the utility parameters of the "deeper" primitives (e.g., risk aversion) are often of independent interest, as they may apply more broadly
   - In contrast, economists are less likely to care about the deeper (psychological?) primitives that explain why, say, individuals value convertible cars

Possible benefits from modeling realized utility

What is the benefit, then, from modeling these "deeper" primitives?

1. Analyzing decision-making under uncertainty is one of the most fundamental topics in microeconomic theory
2. A model of deeper primitives allows for more guidance regarding out of sample extrapolation and for richer counterfactual exercises
   - Consider choices among health insurance plans with deductibles that range from zero to $1,000. A model of realized utility can provide guidance regarding extrapolating to, say, $2,000, and *only* a model of realized utility can tell us something about w.t.p. for, say, a plan with a "donut hole"

Possible benefits from modeling realized utility

What is the benefit, then, from modeling these “deeper” primitives?

1. Analyzing decision-making under uncertainty is one of the most fundamental topics in microeconomic theory
2. A model of deeper primitives allows for more guidance regarding out of sample extrapolation and for richer counterfactual exercises
3. Perhaps also more elegant: connection between ex-ante risk and ex-post cost is more clearly and transparently specified

But even in insurance, this is not always necessary

- For many questions of interest (pricing or welfare from similar contracts), however, a model of contract valuation v may suffice
  - Indeed, Einav, Finkelstein, and Cullen (2010) take such approach
- In other insurance settings, products also vary in non-financial aspects. Then, the appeal of a realized utility model diminishes
  - Consider, for example, the setting of Bundorf, Levin, and Mahoney (2012) who model health insurance choice between HMO and PPO
  - Indeed, they revert to a model that looks much more like the "standard" indirect utility model:

$$v_{ij} = \phi_j + z_i' \beta_j + f(r_i + \varepsilon_i; \gamma_j) - \alpha p_j + \epsilon_{ij}$$

with a key distinction being the importance of the correlation between some of the random coefficients (in this case ε_i, as well as r_i which is observed) and insurer's cost

Informal notes on estimation and identification

- Both approaches lead to estimates of demand and claims rates and how they covary with consumer preferences. Moreover, both recover the essential information needed to analyze consumer surplus or explore the implications of many policy interventions
- For either approach, source of identifying variation is key, and its credibility should be evaluated in a given context
  - Much work on general identification results for v, but not much work on formal identification results for general cases of u
  - Initial results in two recent papers by Perrigne and Vuong
- Two aspects that are important to pay attention to in insurance:
  - Because offers and choice sets are highly customized to individual consumers, need disaggregated data at the choice-set level (at least), otherwise may confuse demand and supply
  - Not impossible, but difficult to make empirical progress w/o some data on subsequent outcomes. Simple (but loose) intuition for identification is of a 2-to-2 mapping (or 3-to-3 in the context of third example)

Plan for today's talk

- To fix ideas, start with a general (fairly standard) framework of demand in insurance markets
- Continue by describing three specific examples of demand models in different insurance contexts, all starting with a model of "realized utility"
- Connect these models to models that are based on indirect utility, which are more common in "traditional" IO
- **Relate demand models to firm's problem of pricing and plan design**

Pricing and plan design

- Easiest to illustrate issues with a *monopoly* provider of a *single* insurance contract, (ϕ, p)
- Normalizing consumer value from no coverage to zero, and maintaining the quasi-linearity assumption, a consumer with characteristics ζ will purchase the contract if $\tilde{v}(\phi, \zeta) \geq p$
- Share of consumers who purchase will be:

$$Q(\phi, p) = \int \mathbf{1}\{\tilde{v}(\phi, \zeta) \geq p\} dF(\zeta)$$

and the insurer's expected costs are:

$$C(\phi, p) = \int \mathbf{1}\{\tilde{v}(\phi, \zeta) \geq p\} c(\phi, \zeta)\, dF(\zeta)$$

- The firm's problem:

$$\max_{\phi, p} \Pi(\phi, p) = p \cdot Q(\phi, p) - C(\phi, p)$$

Pricing and plan design (cont.)

- Fixing the coverage ϕ, the effect of a small increase in price is:

$$\frac{d\Pi(p)}{dp} = Q(p) + \frac{dQ(p)}{dp} \cdot \left(p - \mathbb{E}_\zeta\left[c(\phi,\zeta)\,|\tilde{v}(\phi,\zeta) = p\right]\right)$$

  - The attractiveness of the marginal consumer relative to the average consumer plays a key role

- Choice of ϕ is similar, with the added subtlety that changes in coverage may affect utilization as well as selection:

$$\begin{aligned}\frac{d\Pi(\phi)}{d\phi} &= \frac{dQ(\phi)}{d\phi} \cdot \left(p - \mathbb{E}_\zeta\left[c(\phi,\zeta)\,|\tilde{v}(\phi,\zeta) = p\right]\right) \\ &\quad -\mathbb{E}_\zeta\left[\frac{\partial c(\phi,\zeta)}{\partial\phi}\,|\tilde{v}(\phi,\zeta) \geq p\right]\end{aligned}$$

  - So an additional effect of coverage features is its effect on costs, potentially by inducing behavioral changes as well as mechanically
  - Easy to come up with examples where the marginal consumer w.r.t. p is different from the marginal consumer w.r.t. ϕ

Pricing and plan design (cont.)

- The types of demand models described earlier provide exactly the primitives needed to "fill in" these first order conditions
- One can extend a similar analysis to an oligopolistic setting, although this would raise both conceptual and computational challenges
  - Conceptually, even a game in which firms compete in prices alone may not have the convexity properties typically invoked to assure existence or to justify an analysis based on first order conditions for optimal pricing
  - Computationally, even if an equilibrium does exist and even if it can be characterized in terms of first order conditions, solving numerically for the equilibrium may be challenging

Illustrate using the deductible choice example

- Recall the first (auto insurance) example, where individuals are characterized by (λ_i, ψ_i)
  - λ_i is the individual-specific Poisson claim rate
  - ψ_i is the coefficient of (absolute) risk aversion
- Only λ_i affects insurer's profits. In particular, the *incremental* cost (to the insurer) associated with selling a low deductible contract to individual i is

$$c(d_l, \lambda_i) - c(d_h, \lambda_i) = \lambda_i (d_l - d_h)$$

  thus linking demand to cost
- Consider now how this affects pricing decisions (next slide)

Differential incentives to raise prices due to selection

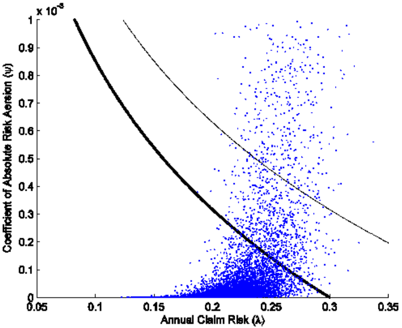

Final remarks

- Described recent empirical models of insurance and their applications
- Did not cover results, but some initial results are somewhat surprising: general finding that certain kinds of welfare losses from asymmetric information, at least in some insurance markets, may be modest
- This is clearly a very early stage of the literature, so lots of room and (I think) interest for more work
- Many many open topics, as well as many other insurance markets that haven't been intensively explored
    - Underwriting and risk selection (Hendren, f'coming)
    - Dynamic insurance provision (Hendel and Lizzeri, 2003; Handel, Hendel, and Whinston, 2013)
    - Consumer search and switching costs (Handel, f'coming)
    - Alternative models of consumer behavior (Barseghyan et al., f'coming)