14.472 Public Finance II

Asymmetric Information: Empirical Welfare Analysis with Revealed Preference

Amy Finkelstein

Fall 2025

Empirical analysis of selection

- Theory: adverse selection can impair efficient operation of insurance markets and create scope for welfare improving government intervention
- Lays out an empirical agenda:
  - Does selection exist in a particular market?
  - What are the efficiency costs of any detected selection and welfare consequences of potential government interventions?

Lecture(s) road map

- Testing for selection
- **[Today]** Empirical welfare analysis I: Using data on choices and claims
- Empirical welfare analysis II: What happens when you can't use choice data
  - Don't trust revealed preference
  - Markets don't exist

Welfare analysis: emphasized in PF

- One distinguishing feature of PF (vs e.g. applied public policy, labor economics etc.) is the attention to welfare (in private markets, of government policy etc)
- But making welfare statements usually requires additional assumptions
  - Do assumptions drive the result? Is result robust to alternative plausible assumptions?
  - How far can we get w the fewest possible assumptions? If we make more assumptions what is it buying us?

Empirical welfare analysis

- Efficiency cost of adverse selection
  - Once know there is private information, want to know how great efficiency cost is
- Welfare consequences of alternative public policies
  - Can public policy improve on adverse selection equbm?
  - Fundamentally an empirical question
    - E.g. Mandates as canonical solution to adverse selection (underinsurance) problem.
    - However, whether they are welfare improving is an empirical question

Welfare an empirical question

I

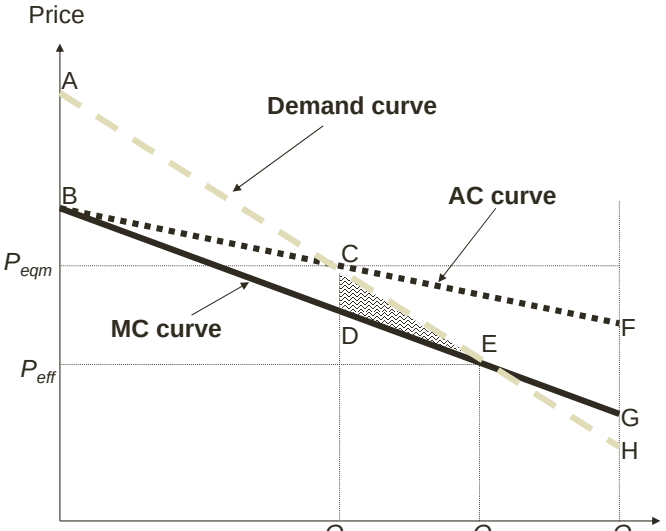

Welfare inferences from extent of pos correlation?

- Some markets with private information about risk type appear more adversely selected than others
  - i.e. larger vs smaller positive correlation
  - Are these markets where efficiency costs likely to be greater?
- Cannot even make qualitative statements about where efficiency cost of adverse selection are likely to be larger vs smaller based on magnitude of reduced form correlation between insurance coverage and risk type
  - Play with the graphs: holding AC of insured vs uninsured same, can rotate demand to get v different welfare costs.

Welfare inferences from extent of pos correlation?

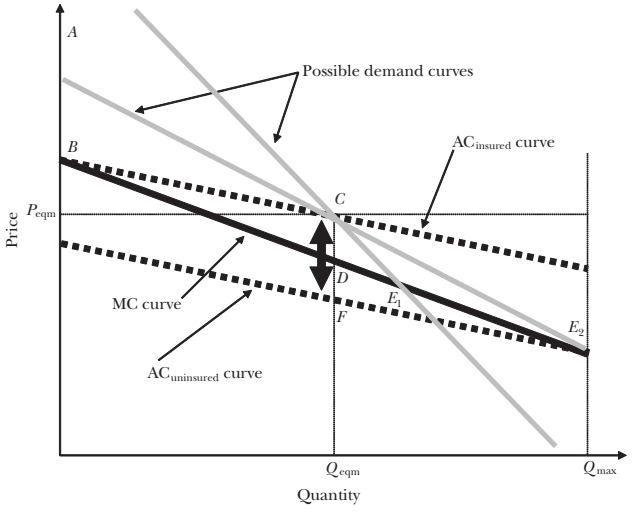

How to estimate welfare cost of selection?

- Need more than the reduced form (positive correlation)
- Two main approaches taken by the literature (see Einav, Finkelstein, Mahoney 2021 Handbook Chapter on Selection Markets for more detail)
  - modeling willingness to pay
  - modeling 'deeper primitives'
- Key theme: trade-offs between two approaches
  - Latter requires more assumptions but allows for richer counterfactuals

Notation

- Consumers indexed by $i \in I$
- Products indexed by $j \in J$ with $j = 0$ denoting outside option
- Consumer characteristics denoted by vector ξ_i
- Product characteristics denoted by vector x_j
- Prices (possibly customized) denoted by $p_j = \{p_{ij}\}_{i \in I}$

- Example: health insurance plans with different x_j (e.g. cost sharing, networks, contract length), and prices that vary with age and smoking status

Notation

- Willingness to pay: $v_{ij} = v(x_j, \xi_i)$
- Consumer surplus: $v_{ij} - p_{ij}$
- Consumer chooses product j if and only if $j = \text{argmax}_{k \in J}\{v_{ik} - p_{ik}\}$

Demand and Revenue

- Integrating over consumers, demand for product j is given by

$$D_j\left(\{p_j, x_j\}_{j\in J}\right) = \int 1\left\{j = \text{argmax}_{k\in J}\left\{\nu(x_k, \xi_i) - p_{ik}\right\}\right\} d\xi_i$$

- Product j's revenue is similarly

$$R_j\left(\{p_j, x_j\}_{j\in J}\right) = \int p_{ij} 1\left\{j = \text{argmax}_{k\in J}\left\{\nu(x_k, \xi_i) - p_{ik}\right\}\right\} d\xi_i$$

Costs

- Unlike conventional markets, in selection markets the firms' costs depend on the identity of the consumer

$$\text{Not} \quad c_j = c(x_j) \quad \text{but} \quad c_{ij} = c(p_j, x_j, \xi_i)$$

- Assuming no scale economies or dis-economiest, total costs are given by

$$C_j\left(\{p_j, x_j\}_{j \in J}\right) = \int c_{ij} 1\left\{j = \text{argmax}_{k \in J}\left\{v(x_k, \xi_i) - p_{ik}\right\}\right\} d\xi_i$$

- Product is adversely selected if and only if persons who purchase the product have higher costs than average

$$E_{\xi_i}(c_{ij} | j, q_{ij} = 1) > E_{\xi_i}(c_{ij} | j)$$

  - Statement about market equilibrium

Welfare

- For many applications, sufficient to recover the joint distribution of $F(v_{ij}, c_{ij})$
  - Total surplus from allocating product j to person i is then given by $v_{ij} - c_{ij}$
  - Can solve for first best allocation, welfare under alternative allocations etc.
- Two main approaches:
  - Models of willingness to pay ('sufficient statistics')
  - Models of 'deeper primitives' (more 'structural')

Modeling willingness to pay

- Directly model WTP $v_{ij} = v(x_i, \xi_i)$ without specifying deeper primitives
- Einav Finkelstein Cullen (2010)
  - Single contract (in addition to outside option)
  - Need to recover joint distribution of ν_i and c_i, which they do by estimating $F_c(c_i|\nu_i)$
    For many applications, sufficient to recover the joint distribution of $F(\nu_{ij}, c_{ij})$
- Basic idea:
  - Rely on standard consumer and producer theory
  - Key feature of selection markets: firms' costs depend on which consumers purchase their products ("endogenous cost curve")
  - price variation can trace out demand & cost curve
- Develop approach and show application to employer provided health insurance
  - Focus: strengths and limitations of approach

Setup and notation

- Only two contracts: H (full coverage) and L (no coverage)
  - Easy to extend to other or more contracts (harder to draw)
  - $p = p_H - p_L$ is the relative price of contract H
- Key assumption: take non-price characteristics of insurance contracts as given
  - As in Akerlof (1970) compared to Rothschild and Stiglitz (1976)
  - Empirically relevant – often observably different individuals offered same menu of contract, just at different prices
- Individuals defined by a vector of attributes $\zeta_i \sim G(\zeta)$, and have to choose a contract H or L
  - ζ_i includes preferences, information set (i.e. expected claims) etc.
    - these are the 'deeper primitives' once can try to estimate
  - Key here is that we will try to do (some) welfare analysis w/o estimating ζ

Setup and notation (con't)

- $\pi(\zeta_i)$ is willingness to pay for H (i.e., $v_H(\zeta_i, \pi(\zeta_i)) = v_L(\zeta_i)$)
- c (ζ_i) is the expected insurable costs under H
  - Cost to insurance company of insuring the individual (ignoring any administrative costs)
  - Abstract from moral hazard for now for notational simplicity (will come back to)

Theory: Demand, Supply, and Equilibrium

- Demand:

$$D(p) = \Pr\left(\pi(\zeta_i) \geq p\right)$$

- Supply:
  - $N \geq 2$ identical risk neutral insurance providers, who set prices in a Nash Equilibrium (a-la Bertrand)
  - Average cost (AC):

$$AC(p) = E\left(c(\zeta)|\pi(\zeta) \geq p\right)$$

  - Marginal cost (MC):

$$MC(p) = E\left(c(\zeta)|\pi(\zeta) = p\right)$$

- Additional (standard) assumptions $\rightarrow$ Equilibrium exists, unique, and given by the lowest break-even price:

$$p^* = \min\left\{p : p = AC(p)\right\}$$

Welfare definitions

- Total surplus from allocating H to individual i is

$$TS(\zeta_i) = \pi(\zeta_i) - c(\zeta_i)$$

- First best allocation: individual i purchases insurance if and only if

$$\pi(\zeta_i) \geq c(\zeta_i)$$

- Constrained efficient allocation: maximizes social welfare subject to the constraint that price is the only instrument available for screening.
  - Constrained efficient: individual i purchases insurance if and only if

$$\pi(\zeta_i) \geq E(c(\widetilde{\zeta})|\pi(\widetilde{\zeta}) = \pi(\zeta_i))$$

Welfare cost of adverse selection

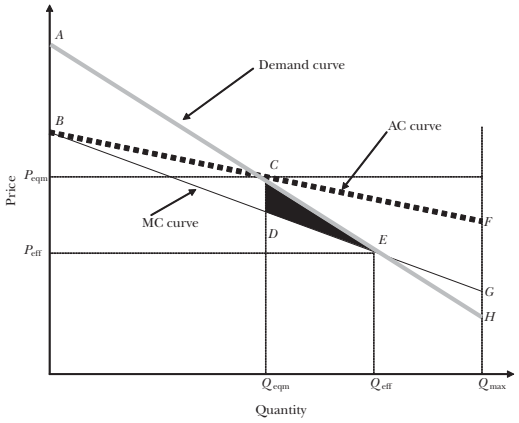

- If have estimated these curves, have welfare cost of selection (CDE).
- Could also evaluate consequences of: subsidies, mandates, pricing on X's...

Sufficient statistics: demand and cost curve

- Graphical analysis illustrates that demand and cost curves are sufficient statistics for welfare analysis of pricing of contracts
- Empirical approach: estimate demand and cost curves but remain agnostic about underlying primitives that give rise to them
- We remain agnostic about underlying primitives (ζ_i) that give rise to demand and cost curve
    - e.g. active vs passive selection generating cost curve?

Estimation

- Sufficient statistics for welfare analysis are:
  - the demand curve $D(p)$
  - the average cost curve $AC(p)$
- Estimation:

$$\begin{aligned} D_i &= \alpha + \beta p_i + \epsilon_i \text{ for everyone} \\ c_i &= \gamma + \delta p_i + u_i \textbf{ for those who endogenously chose } H \end{aligned}$$

- Requires
  - To estimate $D(p)$ variation in p exogenous to demand & quantity
  - To estimate $AC(p)$: *same* variation in p & cost data for sample who endogenously choose H
- Conceptually, variation in p identifies all curves non-parametrically. In practice, likely that need to make functional form assumptions.

Estimation (con't)

- From $D(p)$ and $AC(p)$ we can back out $MC(p)$:

$$MC(p) = \frac{\partial \left(AC(p) \cdot D(p) \right)}{\partial D(p)} = \left(\frac{\partial D(p)}{\partial p} \right)^{-1} \frac{\partial \left(AC(p) \cdot D(p) \right)}{\partial p}$$

- Conceptually, variation in p identifies all curves non-parameterically. In practice, likely that need to make functional form assumptions.
  - Here structure could be useful to guide functional form
  - But graphs highlight which parts of curves are important to "get right"
- Key requirement: Need variation in p that is exogenous with respect to demand and cost

What about moral hazard?

- Welfare analysis takes moral hazard effects as given
- Government generally has no comparative advantage in combating moral hazard effects
  - Part of the "technology" that we take as given
- Analysis of welfare / policy under adverse selection should take moral hazard environment as given
- NB: enormous empirical literature estimating mh effects of social insurance programs
  - Recall: this speaks to optimal level of private or social insurance

Moral hazard

- Since costs are a function of insurance coverage, useful to define $c^H \geq c^L$
  - c^j is expected cost of insurance coverage H when behavior is as under j coverage
  - correspondingly two average cost curves (AC^H and AC^L) and two marginal cost curves (MC^H and MC^L)
- To explicitly recognize moral hazard in preceding analysis, replace c, AC, and MC with superscript "H"
  - Recall that cost curve estimated on sample of individuals who endogenously choose H
- Estimate a second pair of cost curves AC^L and MC^L. Difference between MC^H and MC^L measures moral hazard for each customer time.

Moral hazard: estimation

- What would we do if we wanted to actually estimate moral hazard in an application?
- Regress costs on insurance coverage
  - Instrument for insurance coverage using exogenous variation in prices

Moral hazard: implications for welfare analysis of selection

- Preceding welfare analysis goes through.
  - Note that the c we defined earlier is c^H – i.e. the relevant cost curve is the actual costs of coverage given the moral hazard effect of coverage on expected costs
- Intuition: Why doesn't c^L matter for analysis:
  - Firm: only behavior of insured individuals matters (c^H). How would behave if not insured (c^L) not relevant
  - Individual: gap between c^H and c^L does matter but incorporated into effect on WTP (π)
  - (Caveat: when L is partial coverage, need to account for any "moral hazard externality". see e.g. Cabral and Mahoney (AEJ 2019))

Final comment on moral hazard

- What if what creates the downward slope of the cost curve is heterogeneous moral hazard?
  - i.e. those who have high WTP for insurance have higher behavioral response to the contract
- "Selection on moral hazard" (Einav et al. 2013)
  - implications for current welfare analysis based on mispricing
  - implications for combatting adverse selection through monitoring

Empirical application

- Employer provided health insurance at Alcoa
  - Choice between high and low deductible plan
- Rich data, including all relevant aspects of option set (vs. e.g. "networks")
  - The menu of health insurance options available to each employee
  - The premium associated with each option
  - Employee choices
  - Employee (and dependents') subsequent medical expenditure
  - Rich demographics – everything price setter likely to observe
- Variation in relative price charged for high vs low deductible plans offered to different business units within Alcoa
  - How good is this variation?

Empirical constructs

- $p_i = p_i^H - p_i^L$ where p_i^j is employee i's annual contribution for coverage j
- $D_i = 1$ if i chose H; $D_i = 0$ if i chose L
- m_i is employee i's vector of medical cost during 2004
- $c(m_i; j)$ is the insurer's cost of covering m_i under coverage j
- $c_i = c(m_i; H) - c(m_i; L)$ is the incremental insurer's costs from covering i with H vs. L (holding behavior m_i fixed)
    - Note one will be counterfactual so need to construct (both ideally) using plan rules

Specification

- We estimate (using OLS):

$$\begin{aligned} D_i &= \alpha + \beta p_i + \epsilon_i \text{ for everyone} \\ c_i &= \gamma + \delta p_i + u_i \text{ for those who chose } H \end{aligned}$$

  recall $c_i = c(m_i; H) - c(m_i; L)$

- Marginal cost derived from these without additional estimation

Raw data with basic findings

(Relative) price ($) (1)	Number of employees (2)	Fraction chose contract H (3)	Average incremental cost ($) for those covered under	
			Contract H (4)	Contract L (5)
384	2,939	0.67	451.40	425.48
466	67	0.66	499.32	423.30
489	7	0.43	661.27	517.00
495	526	0.64	458.60	421.42
570	199	0.46	492.59	438.83
659	41	0.49	489.05	448.50

Results: estimates

Dependent variable (sample)	1 if chose High (both High and Low) (1)	Incremental cost (only High) (2)
Panel A: Estimation results		
Relative price of High (US$)	−0.00070	0.15524
	(0.00032)	(0.06388)
	[.034]	[.021]
Constant	0.940	391.690
	(0.123)	(26.789)
	[.000]	[.000]
Mean dependent variable	0.652	455.341
Number of observations	3,779	2,465
R^2	.008	.005

Standard errors (in parentheses) clustered on state

p values in [square brackets]

The welfare cost of adverse selection

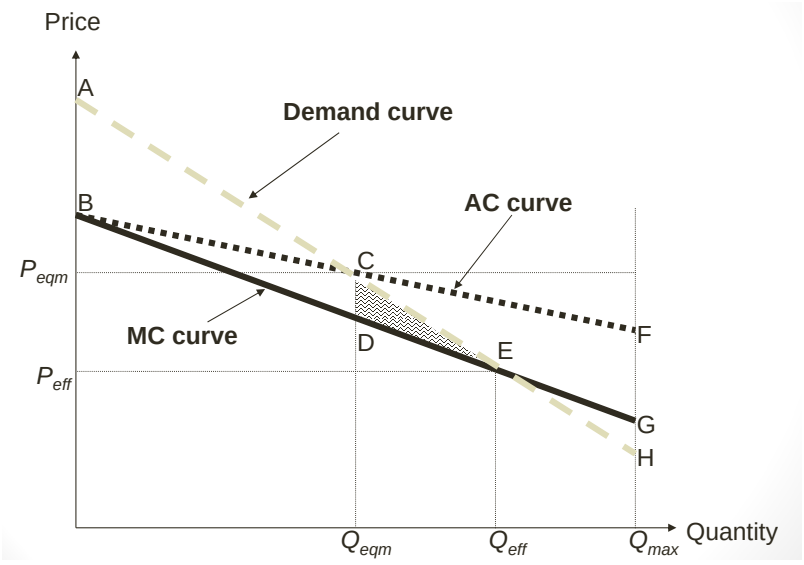

Results: graphical illustration

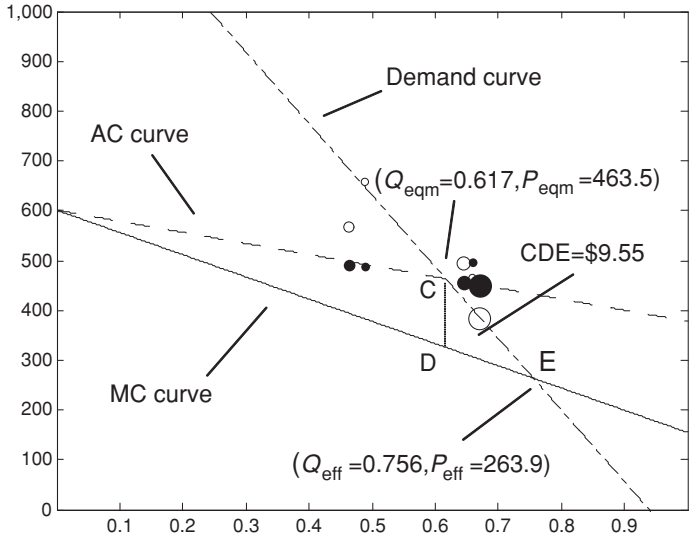

Results: welfare benchmarks

- Estimated demand and cost curves can also provide benchmarks to help provide context
- Preferred benchmark:
  - Cost of price subsidy required to achieve efficient price – i.e. $\lambda(P_{eq} - P_{eff})Q_{eff}$ – is about 5 times welfare gain from moving from adverse selection equilibrium to efficient price.
- Other benchmarks (much more out of sample)
  - Welfare cost of mandatory coverage by H is about 3 times equilibrium welfare cost of adverse selection
  - Welfare cost of adverse selection ~3% of total surplus at stake from efficient pricing

Discussion: key assumptions of framework

- "Valid" pricing variation
- Revealed preference
  - Or at least a particular behavioral model
- Fixed contracts
  - Estimating inefficiency selection causes via mispricing
    - Not capturing welfare cost of adverse selection from distortion of contract space (Rothschild-Stiglitz 1976)
  - Policy analysis limited to changes in prices of existing contract space
    - Preferable "small" price changes that don't expect to trigger endogenous contract response

Many subsequent (better!) applications (and extensions)

- Insurance markets (health, crop, flood, unemployment, worker's compensation, disability) and credit markets
- For more details see Einav and Finkelstein (2023, *Geneva Risk and Insurance Review*) "Empirical analyses of selection and welfare in insurance markets: a self-indulgent survey"

Application: Health insurance subsidies for low income adults

- Finkelstein, Hendren, Shepard (2019)
- Subsidized health insurance exchange in MA introduced in 2006 "Romneycare" reform
  - Precursor to ACA exchanges
  - Subsidies for low-income, non-elderly uninsured adults between 133-300% of FPL
- Quasi-random pricing across individuals
  - Public subsidies - designed to make insurance "affordable"
    - Change at discrete income bins
  - Regression discontinuity design

Quasi-Random Variation in Price

Panel A: Premiums for Cheapest Plan (2009-2013)

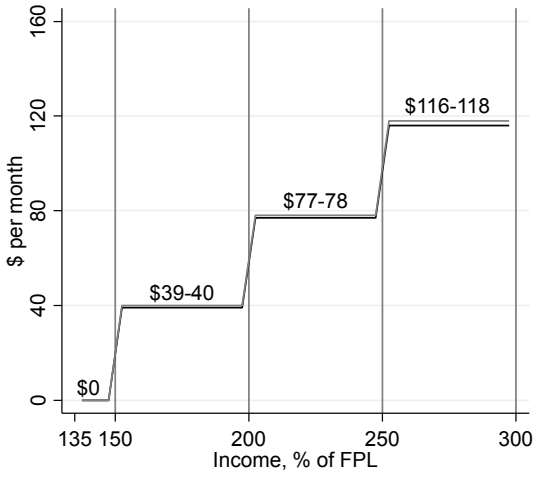

Price Changes Prompt Coverage Changes

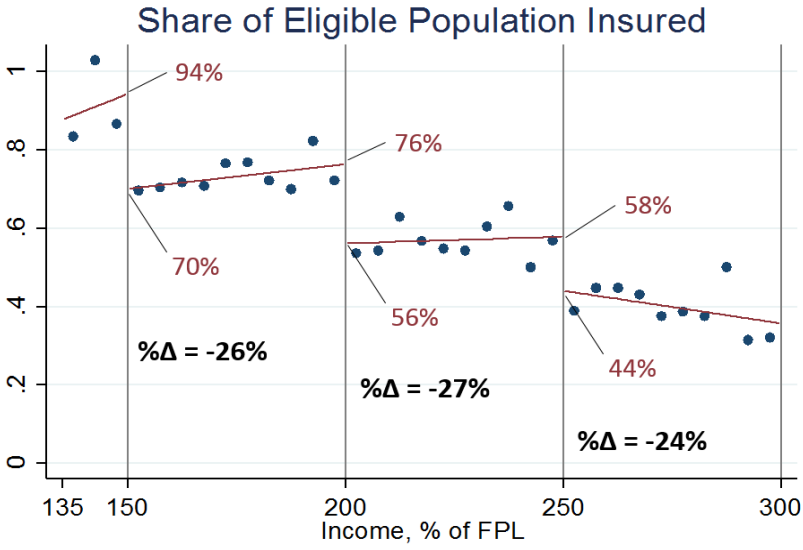

Observed Demand Points

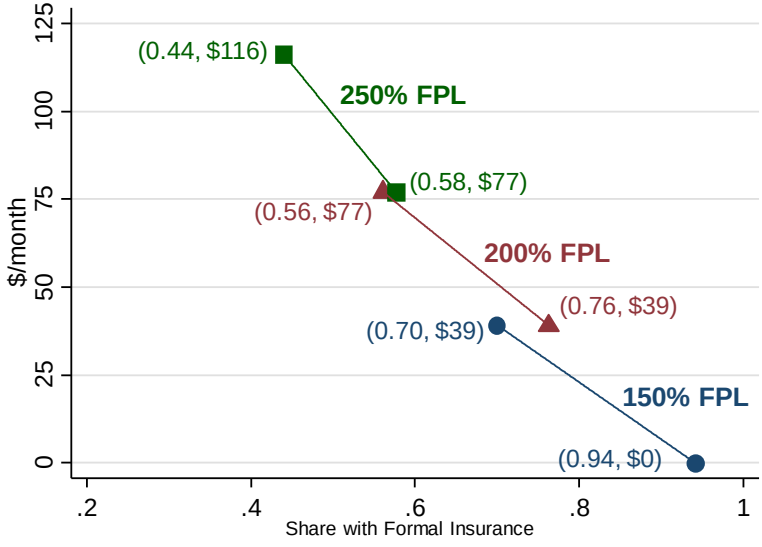

WTP (Adjusted to 150% FPL)

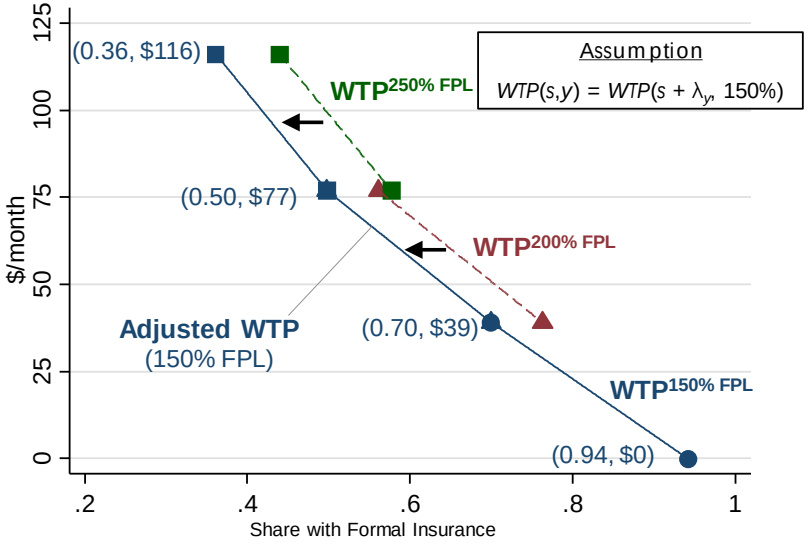

Demand Curve

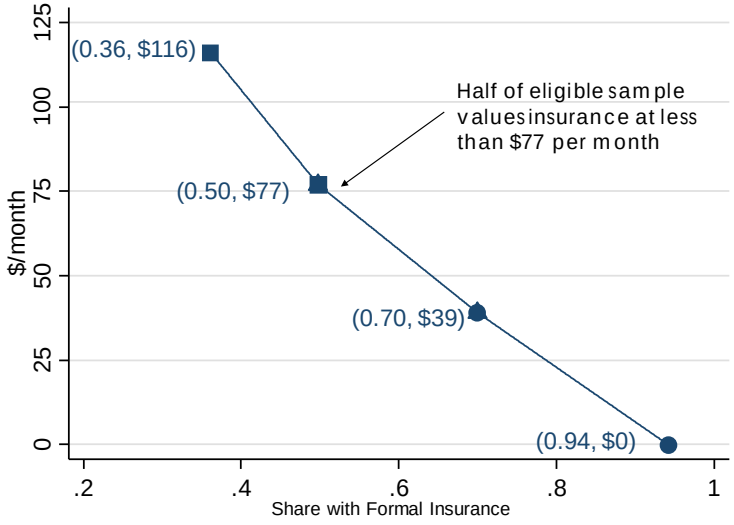

Average Insurer Costs

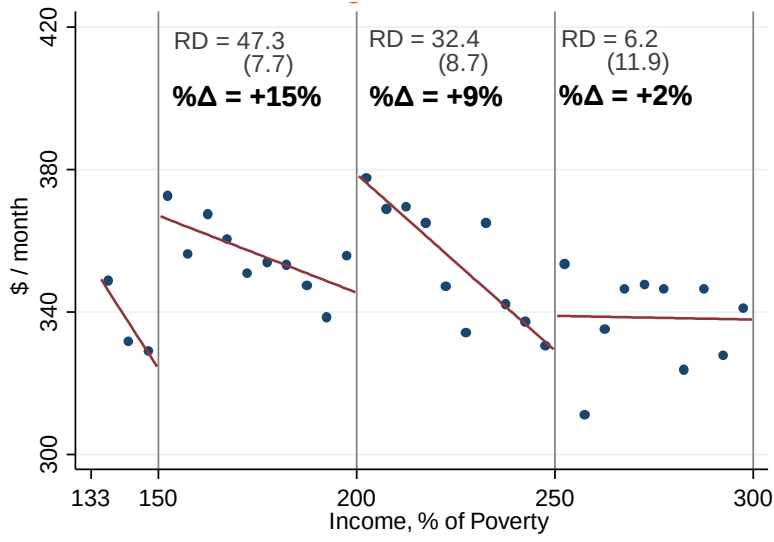

Observed Average Costs

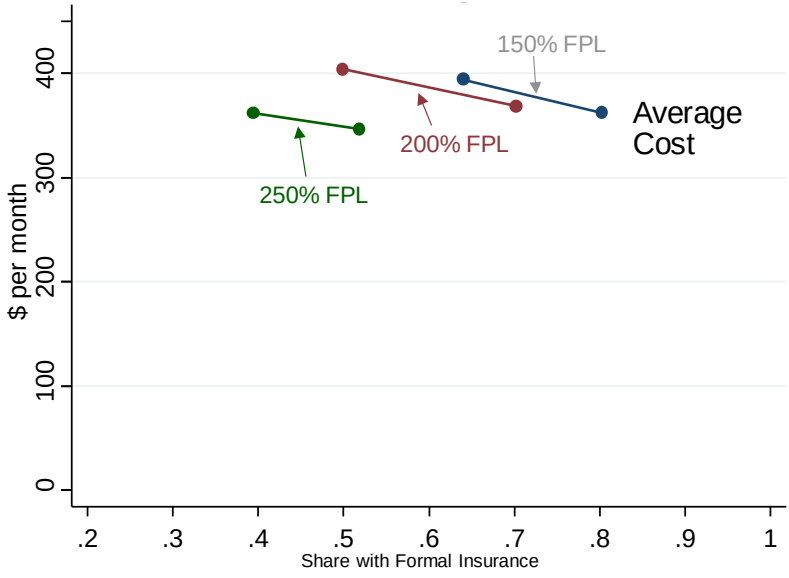

Average Cost (Adjusted to 150% FPL)

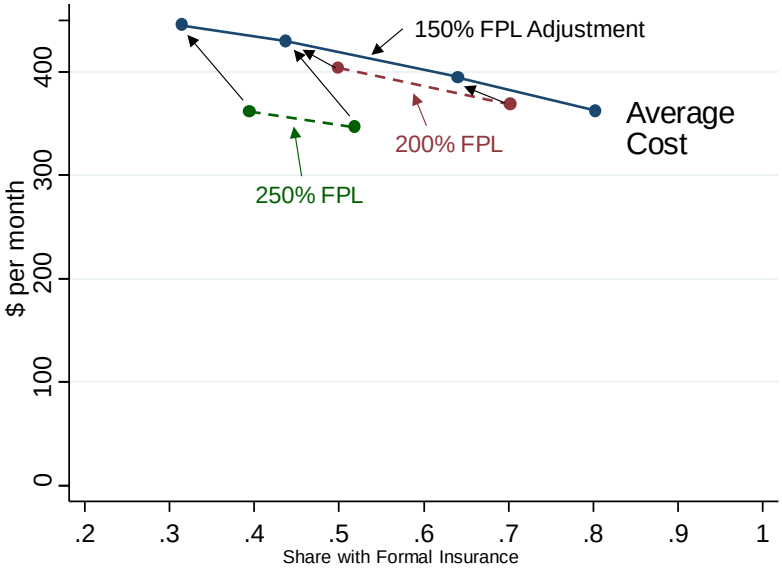

Final WTP and Cost Curves

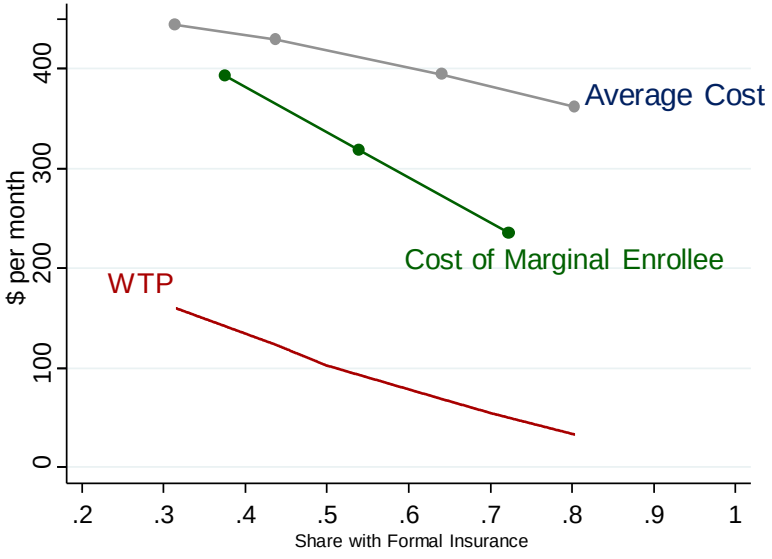

Little Take-up without Large Subsidies

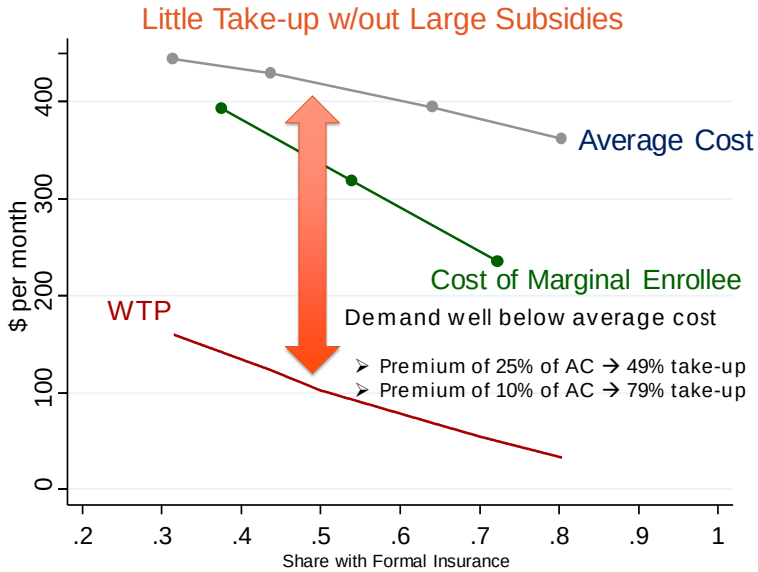

Adverse Selection Alone Cannot Explain Low Coverage

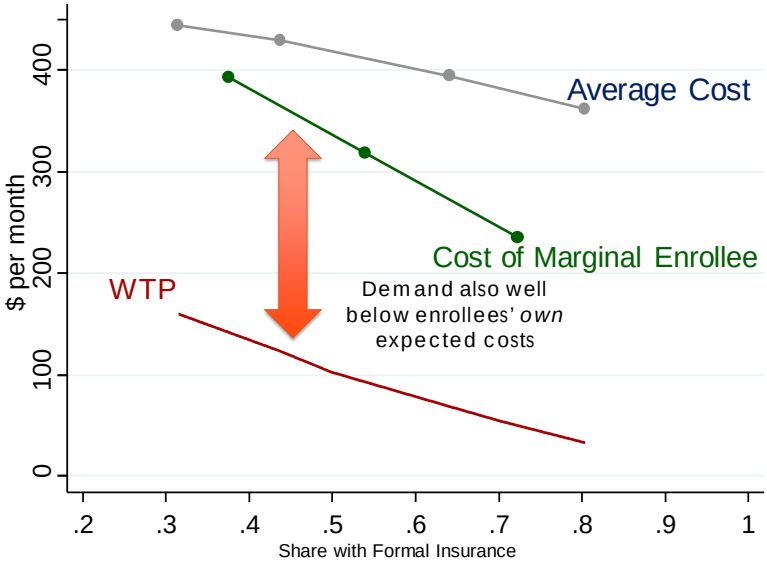

So Why is WTP Below Own Costs?

Application: Credit Markets

- Analagous problem: insurance premium ~ interest rate on loan
- Additional challenges:
  - Assumption that premiums affect contract choice but not costs conditional on product choice allows separation of selection from moral hazard
    - insurance: premiums are sunk, so only threat is income effects
    - credit markets: interest rates can affect downstream probability of default
  - Demand curve estimates maximum quoted rate borrowers are willing to accept
    - Need to adjust down for borrowers' expected default; absent moral hazard adjustment factor based on MC curve (= expected charge off rate)
    - With moral hazard, this is an overadjustment / will overestimate WTP (lower price to get marginal borrower, which lowers default rate on inframarginal borrowers)
- Application: DeFusco et al. (2021) RCT of interest rates for large fintech lender in China

Application: Extensive vs Intensive Margin Selection

- Geruso, Layton, McCormack, and Shepard (2019)
- "Trade-offs between Extensive and Intensive Margin Selection in Competitive Insurance Markets"
- Use EFC (2010) framework to make a simple, important point:
  - once you have more than two choices (e.g. insured vs not and within insured high vs low coverage), then policies that work to reduce selection on one margin can worsen it on the other
  - e.g. insurance mandate penalty
    - can reduce selection on extensive margin
    - but worsen on intensive margin, by bringing in healthy people who lower cost of low coverage plan, can get people dropping out of high coverage plan
- Show calibrated results using demand and cost system from Finkelstein, Hendren and Shepard (2019)
  - Nice example of re-using existing estimates for another purpose

Extension: What if insurance market not perfectly competitive?

- Assumed equilibrium was $P = AC$
  - But since empirical work requires out-of-equilibrium pricing variation, don't actually observe equilibirum
- Could "easily" extend welfare analysis under a different specific assumption about competition
  - Mahoney and Weyl (2017) develop this formally

Mahoney and Weyl (2017)

- Interaction of market power (imperfect competition) with selection
- Example: risk adjustment subsidies to plan (based on difference between average cost of enrollees and average cost in population)
- This flattens AC curve (at population average)
  - Under perfect competition, lowers average costs and and creates higher Q, lower P equilibrium
  - Under imperfect competition, recall firms set price too high relative to social optimum. Adverse selection reduces incentives to mark up prices (because get worse risk pool / higher costs). Risk adjustment, by offsetting adverse selection, undermines this incentive and may lead to higher P, lower Q
- Example of the **theory of the second best**

The Theory of the Second Best

- Lipsey and Lancaster (ReStud 1956)
  "It is well known that the attainment of a Paretian optimum requires the simultaneous fulfillment of all the optimum conditions. The general theorm of the second best optimum states that if there is introduced into a general equilibrium system a constraint which prevents the attainment of one of the Paretian conditions, the other Paretian conditions, although still attainable are, in general, no longer desirable.... From this theorm there follows the important negative corollary that there is no a priori way to judge as between various situations in which some of the Paretian optimum conditions are fulfilled, while others are not."

- Autor (2025) "imperfect automation is not a first step toward perfect automation, anymore than jumping halfway across a canyon is a first step toward jumping the full distance"

The Theory of the Second Best

- Lipsey and Lancaster (ReStud 1956)

  "It is well known that the attainment of a Paretian optimum requires the simultaneous fulfillment of all the optimum conditions. The general theorm of the second best optimum states that if there is introduced into a general equilibrium system a constraint which prevents the attainment of one of the Paretian conditions, the other Paretian conditions, although still attainable are, in general, no longer desirable.... From this theorm there follows the important negative corollary that there is no a priori way to judge as between various situations in which some of the Paretian optimum conditions are fulfilled, while others are not."

- Translation: Full employment program for empirical economists

Extension / Look Back: Ordeals in selection markets

- Ordeals in non-selection markets (recall takeup and targeting lectures)
  - Ordeals as nonfinancial "price" of enrolling
  - Price in standard markets screens out people with low value (demand)
  - Whether ordeal screen out low or high value individuals is an empirical question

- Shepard and Wagner (2025, AER): Do Ordeals Work for Selection Markets? Evidence from Health Insurance Auto-Enrollment
  - Selection markets: heterogeneity across individuals not only in value (demand) but also costs
    - Insurance markets, credit markets, education ...
  - Now want to screen on net social value: social value minus costs

Gains from targeting (ordeals) w/o and w selection

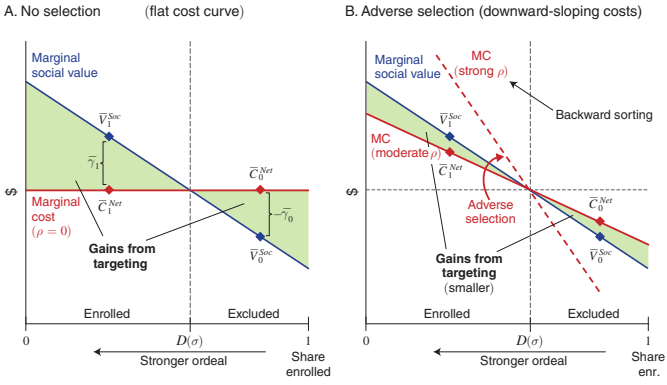

FIGURE 1. GAINS FROM ORDEALS TARGETING WITH NO SELECTION VERSUS ADVERSE SELECTION

Notes: The figure shows the gains from targeting from ordeals in two cases: (i) the "standard" ordeals case without selection (a flat marginal cost curve, panel A) and (ii) with adverse selection (downward-sloping cost curve, panel B). Both panels depict enrollee value and cost curves for marginal enrollees as the ordeal strengthens and enrollment drops (moving right to left), using a setup similar to Einav, Finkelstein, and Cullen (2010). The green shaded areas are the "gains from targeting," which shrink or become negative under adverse selection.

Example: Universality (vs targeting) can be optimal bc of selection

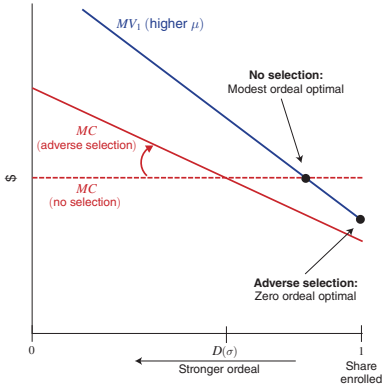

FIGURE 2. OPTIMAL UNIVERSALITY WITH ADVERSE SELECTION

Notes: The figure shows how adverse selection increases the likelihood of "optimal universality" when the social marginal value (MV) curve is shifted upward (relative to Figure 1) due to a higher social welfare weight, μ. With no selection, the new marginal value curve (MV_1) still intersects marginal cost (MC), implying that a (more modest) ordeal is still optimal. With adverse selection, MV_1 lies entirely above MC, implying full enrollment (zero ordeal) is now optimal.

Ordeals in selection markets

- Theory: Adverse selection weakens the classic self-targeting case for ordeals
  - Classic self-targeting: ordeals screen out low demand individuals
  - Adverse selection: low demand indivdiuals tend to be low cost (will discuss at length in Part II of course)
  - Thus adverse selection can undermine ordeals' standard rationale of excluding low-value individuals since they are also low-cost so may not be inefficient to enroll

- Empirics: Impact of minor ordeal (extra step in enrollment process for health insurance)
  - Major impact on enrollment: 33 percent decline
  - Targeting: disproprotionately exclude young, healthy, and economically disadvantages

- Nice cross-pollination of ideas from selection literature into takeup literature
  - Will return to once we've covered selection in more detail
  - Also has **really** nice algebraic and graphical description of self-targeting with and without selection

Discussion of EFC: Attractions

- Model demand and costs but not their primitives (ζ_i) Don't have to take a stand on structure / nature of private information or preferences etc
- Extremely simple to implement
  - Relatedly: transparent. direct mapping from model to data. Makes it easier to see the key empirical assumptions.
- Broadly applicable (as have seen)
  - Data requirements are
    - Demand and cost (as required for pos correlation test)
    - Pricing variation. = key hurdle. But many potential sources
  - Results likely relatively comparable across markets (vs more structural models where model tailored to market)
  - Caveat: settings where fixed contract assmpt seems reasonable
- Bonus: direct test of selection (shape of cost curve)
  - In one package: detect selection and examine welfare cost

Discussion of EFC: Limitations

- Requires good price variation – not always easy to find (but see many possibilities...!)
- Fixed contracts assumption
  - Cannot evaluate welfare from introducing contracts not observed
    - Requires underlying structural primitives (one of the motivations for estimating 'deeper primitives')
    - Welfare analysis limited to policies that change price of existing contracts (mandates; subsidies; restrictions on pricing)
  - Limited to "local" welfare analysis for relatively small price changes if concerned about endogenous contract respond
  - Familiar tradeoff
    - Product-space (e.g. Almost Ideal Demand System) vs Characteristic space (e.g. BLP) approaches to differentiated demand estimation. Latter can be used to evaluate welfare from new goods before introduced.

Alternative: Estimating 'deeper primitives'

- Rather than specify WTP funtion directly, derive it from deeper primitives based on product utilization
- In many selection markets, we observe after-purchase behavior
  - We don't observe whether someone drinks milk they bought
  - We do observe whether someone uses their health insurance
- And in selection markets we have theory that guides why consumers like certain products vs others
  - Willingness to pay for insurance can be parameterized by curvature of utility (CRRA or CARA)
  - Why do consumers prefer electric cars (gas savings? less noisy ride? energy conservation?)

Example: Cohen and Einav (AER, 2007)

- Choice of high vs. low deductible in Israeli automobile insurance company
- Each individual i chooses between a high-deductible contract with price and (per claim) deductible of $p_{(i,HD)}$ and $x_{(i,HD)}$, respectively, and a low deductible contract, $p_{(i,LD)}, x_{(i,LD)}$.
- Key assumptions:
  - Claims arrive according to a Poisson process that is not affected by the choice of deductible (i.e. no moral hazard)
  - CARA utility over wealth

Cohen and Einav (2007), cont'd

- Given these assumptions, the expected utility from a contract p_{ij}, x_{ij} is:

$$Eu_{ij} = (1 - \varepsilon_i)u_i(a_i - p_{ij}) + \varepsilon_i u_i(a_i - p_{ij} - x_{ij})$$

  - where ε_i is the individual's Poisson risk rate, a_i is their wealth, and $u_i(w) = -exp(-\psi_i w)$, with ψ_i denoting the coefficient of absolute risk aversion.
  - With CARA preferences, the consumer's wealth does not affect their insurance choices, so the relevant consumer characteristics are given by $\zeta_i = \psi_i, \varepsilon_i$.
    - i.e. risk preference (ψ_i) and risk type (ε_i)

- Key insight for identification:: joint distribution of risk type and preferences identified from relationship between risk experience and insurance choice in the data
  - Ex post risk experience helps identify risk type (assuming no moral hazard)
  - Conditional on risk type, insurance choice (e.g. high vs low deductible) pins down preferences (e.g. risk aversion)

Identifying joint distribution of risk preferences and risk type)

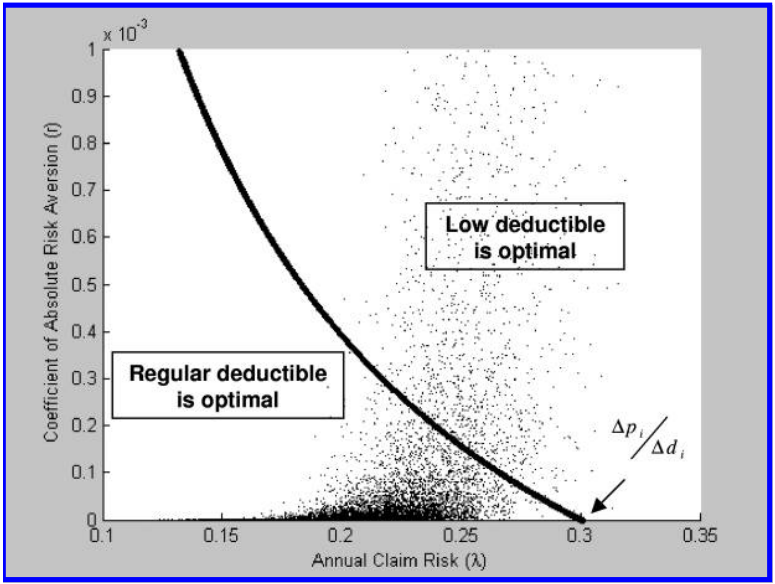

Once you have the deeper primitives...

- Welfare analysis is easy / the sky's the limit
  - Can compute welfare at observed equilibrium
  - Can compare to welfare in counterfactual equilibriums
    - First best (symmetric information). Gives welfare cost of adverse selection.
    - Mandatory social insurance. Gives welfare gain / loss from a particular government intervention.
    - Contracts not observed in equilibrium
  - Key area of discretion: choice of counterfactual (using but not abusing the model)
- Example: Einav Finkelstein and Schrimpf (EMA 2010)

Discussion

- Why would you want to do this?
  - Interested in primitives
    - e.g. Cohen and Einav 2007: interested in risk aversion (average, dispersion, correlates of dispersion...)
  - Want to say something about welfare from contracts not observed in the data
    - Although hopefully not too far out of sample
    - The art of the counterfactual

- Key limitation: without strong assumptions, can rationalize data with very different underlying primitives
  - Fundamentally, risk preferences and private information about risk type are separately identified only by functional form
  - If you see two people with same claims history choosing difference deductibles is it because they have:
    - different risk aversion?
    - different information about underlying risk type?
    - made a mistake?

Empirical welfare analysis: road map

- Thus far: two approaches to empirical welfare analysis
  - Willingness to pay vs deeper primitives
  - More vs. less structure
- Both approaches rely on observing demand and taking a revealed preference approach
- Up next:
  - What if we want to abandon revealed preference / "go behavioral"?
  - What if market doesn't exist / has completely unraveled so we can't observe demand? How do we recover preferences?

Motivation: Small estimated welfare costs of adverse selection

- A large number of studies find small estimated welfare costs of pricing distortion induced by adverse selection in health insurance:
  - Cutler and Reber (QJE 1998), Carlin and Town (2010), Einav, Finkelstein and Cullen (2010) Bundorf, Levin and Mahoney (AER 2012), Einav, Finkelstein, Ryan, Schrimpf and Cullen (AER 2013), Handel (AER 2013).
- Also small welfare costs from pricing distortion induced by adverse selection in annuity markets (Einav, Finkelstein and Schrimpf 2010)

Interpretation?

- Adverse selection not a big deal
  - At least given current policy environment.
  - Perhaps where it WAS a big deal for welfare, that's where policy solutions emerged
    - Tax subsidies for employer-provided health insurance
    - Mandates
    - Social safety net / publicly provided insurance
  - That doesn't mean couldn't design policies that on the margin would create huge adverse selection
- And/or something is missing from the approach (= Next two topics)
  - Can we use observed demand to infer value of insurance?
  - Lampost problem: studying relatively small margins of contract choice in markets that exist.
    - What about welfare costs from complete unraveling of market (ultimate distortion of contract space)?

Lamppost Problem