Risk Adjustments in Health Care Markets

Liran Einav
(Stanford and NBER)

ES World Congress, Montreal
August 21, 2015

What is this talk about?

- Some (most?) of you must be wondering what I am doing in a session titled "Dynamic Implications of Mechanism Design" … so do I … ☺

- In a room that is presumably full of theorists, talking about some of my current/recent empirical work would be running the risk of putting everyone to sleep …

- Instead, my plan is to provide general background, and some teasers, to get more theorists interested in a topic, which (IMO) is becoming increasingly important, but with a few exceptions is relatively unexplored form a theory perspective

Overview

(a) increasing share of healthcare provision around the world takes place within "managed competition" markets

+

(b) the "big data revolution" that takes over many businesses (and governments)

→ Risk adjustments will play an increasingly important role in healthcare markets

While often viewed as a *statistical* object (better risk adjustment = greater predictive power such as R^2), I'll try to argue that this is in fact an important *economic* object, and could/should serve as an important market design instrument

As such, optimal risk adjustment system should be a solution to solving a social objective function, taking demand, supply, market equilibrium, and presumably other economics /incentive constraints as given. Of course, more precise prediction would often (but not always!) make things better, but in a "second best" world other considerations should be in play as well

This problem raises many interesting theoretical and empirical issues, which are mostly unexplored; seems related to optimal income tax problems, except that distortions have externalities so the IO of the market becomes important

Outline

- What is risk adjustment, why and how it is being used
- Typical risk adjustment systems
- Simple empirical framework to think about risk adjustment
- Selected issues worth thinking about

Outline

- **What is risk adjustment , why and how it is being used**
- Typical risk adjustment systems
- Simple empirical framework to think about risk adjustment
- Selected issues worth thinking about

Adverse Selection

(Einav, Finkelstein, and Cullen, 2010)

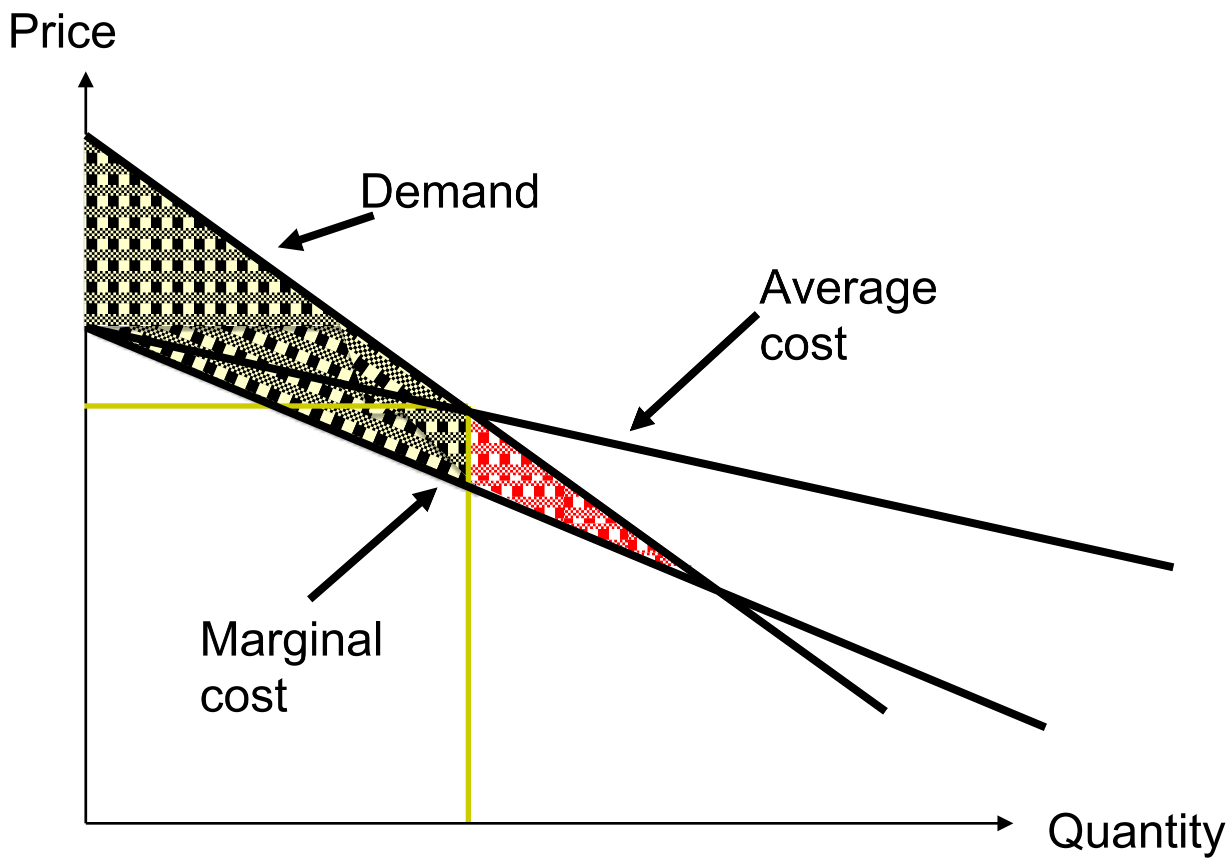

Risk-Based Pricing

Typical way to combat adverse selection in insurance markets is to price risk, and charge high-risk individuals higher premiums

But this is not a great solution in the context of health insurance:

- Regressive … Health and income are negatively correlated, so we'd generally want to tilt pricing the other way
- High expected cost likely associated with high absolute risk, so risk-based premium may price out of the market exactly those we would like to cover the most
- Risk-based premium would introduce "reclassification risk," which may lead to inefficiencies that are even larger than those generated by adverse selection (Handel, Hendel, and Whinston, 2015)

Risk Adjustments

- Fortunately, and at least in part due to the above reasons, healthcare markets rarely operate as completely free markets, and are often "sponsored" (by governments, states, or employers)

- When healthcare is not directly provided by the sponsor, the typical market setting is one of "managed competition"
  - Highly regulated set of rules regarding eligible coverage contracts and price setting
  - Subsidies (often in some form of vouchers) to address redistribution, affordability, and perhaps also adverse selection
  - A risk adjustment system that attempts to reduce selection and cream-skimming concerns

- This is a way to implement uniform (or constrained) premiums and still compensate insurers for covering a (predictably) riskier consumer pool
  - Progressive … a form of cross subsidy from healthy to sick
  - When the sponsor subsidizes the entire market (as is often the case), take up is less of an issue even for the healthy
  - Reclassification risk is not an issue anymore

Benchmark: perfect risk adjustment

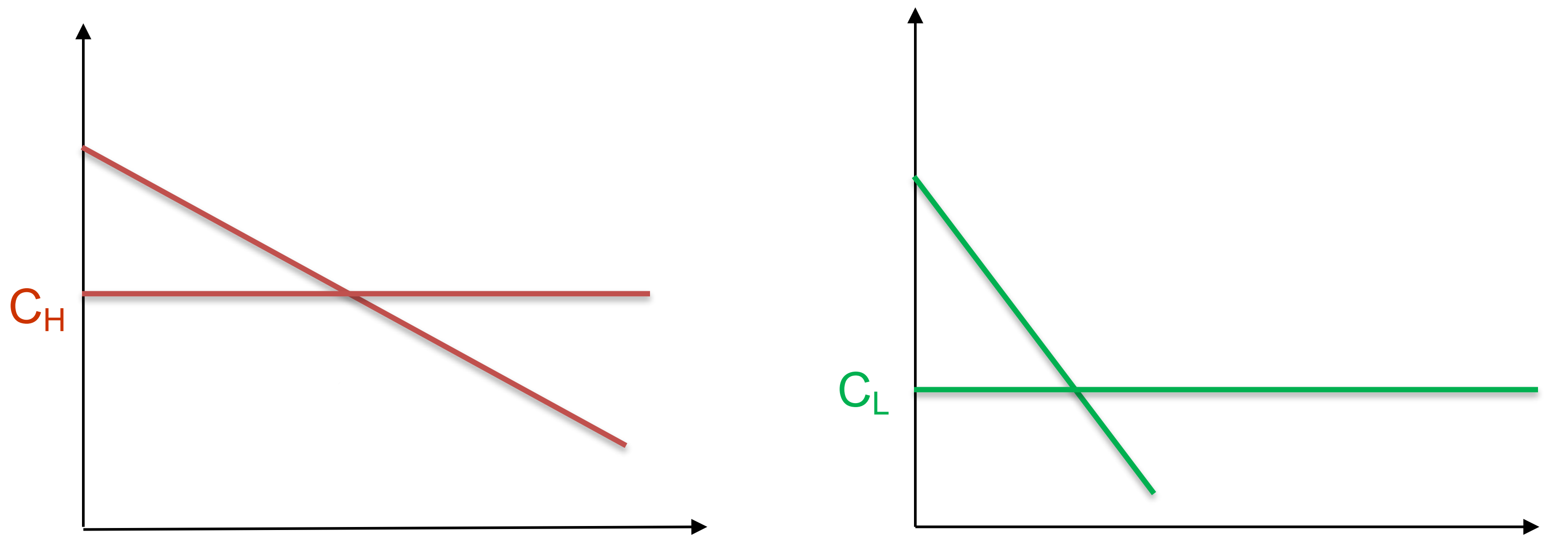

- If we could tell perfectly type H from type L, we can compensate insurers with an extra $\Delta C = C_H - C_L$ and this would completely eliminate the adverse selection problem

- Competing insurers would set a market price $P = C_L$, would not care about their customer composition, and we'd get the efficient outcome

Benchmark: perfect risk adjustment

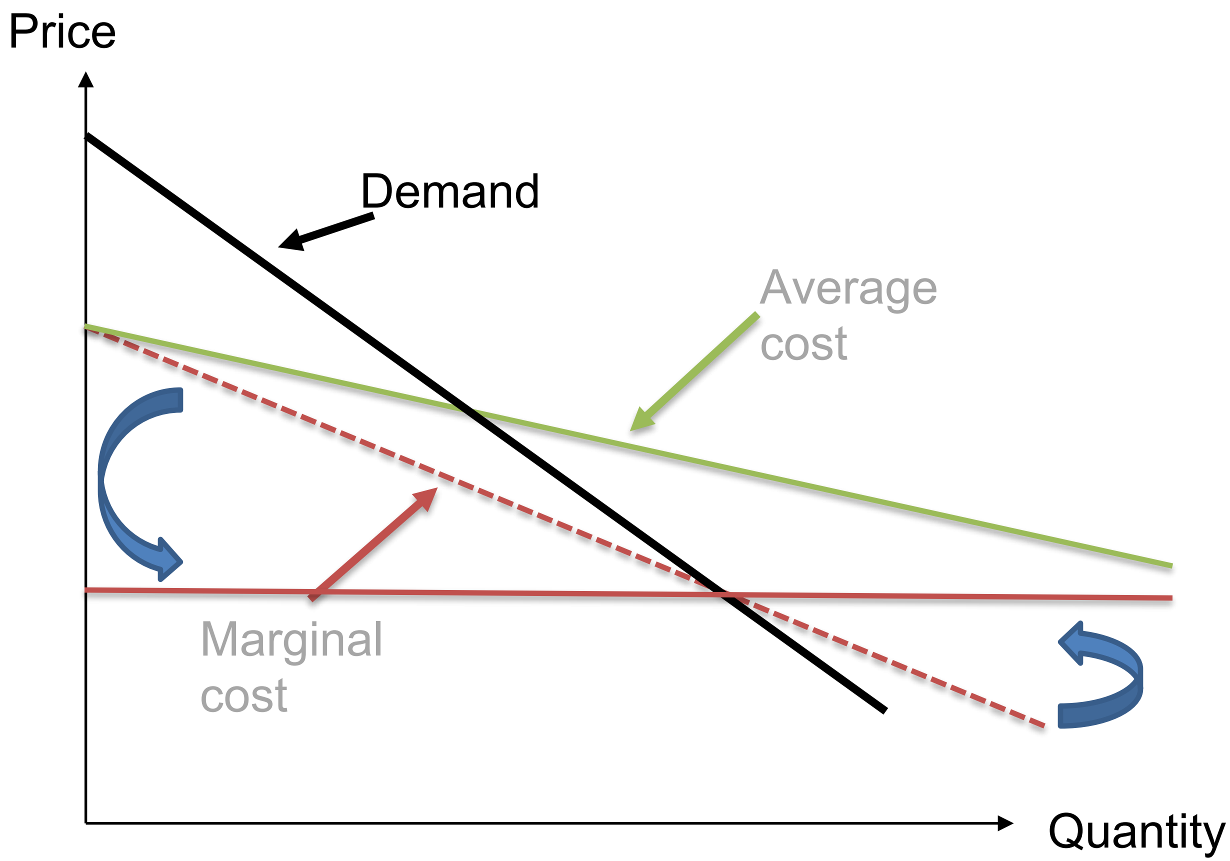

- Unfortunately, reality is more messy …
  - Recent paper by Braverman and Chassang (2015) makes the point that even perfect risk adjustment is in practice imperfect under "Big Data limit"

Outline

- What is risk adjustment , why and how it is being used
- **Typical risk adjustment systems**
- Simple empirical framework to think about risk adjustment
- Selected issues worth thinking about

Actual risk adjustments

- Three ways to implement risk adjustments:
  - Ex-ante: compensate insurers based on forecasted cost of enrollees
  - Ex-post: compensate insurers after the fact, based on realized cost (a bit like "cost plus" regulation)
  - In between: compensate insurers based on "nowcasted" cost of enrollees

- In what follows I will limit attention to ex-ante risk adjustment, which is the most common and the one that provides the least perverse incentives

Common way to implement risk adjustment

- Basic structure begins with some version of a purely predictive model (not always as sophisticated as one would have guessed):
  - Object to be predicted: next year’s healthcare cost
  - Input: some demographics (age, gender, location) and past (often previous year’s) claims

– Basic idea is to map claims to health conditions, and distinguish persistent conditions (cancer, diabetes) from transitory ones (broken arm, infectious disease)

- Then each individual gets a predicted cost, which often gets normalized to some representative agent. All current systems end with a one-dimensional health risk score per individual that feeds into the risk adjustment system
  - That is, systems price the person, not the deal. Will return to this later.

Examples of risk adjustment: Medicare

- Current model uses gender, age, and approximately 70 health conditions in an additive separable way
  - Use past claims to "turn on" relevant indicators
  - Each indicator associated with a health risk coefficient
  - Add up all relevant indicators to obtain a normalized "risk score": average person is 1, a risk score 2 is expected to be twice as expensive, and so on

- Greatest contributions to the score: metastatic cancer (2.28), respirator dependence (1.87), third degree burns (1.42)
- Most prevalent: diabetes w/o complication (13.6%), chronic pulmonary disease (11.4%), vascular disease (11.3%)
- Most important: chronic pulmonary disease (11.4%; 0.4), congestive heart failure (10.3%; 0.41), vascular disease (11.3%; 0.31)

Medicare risk adjustment

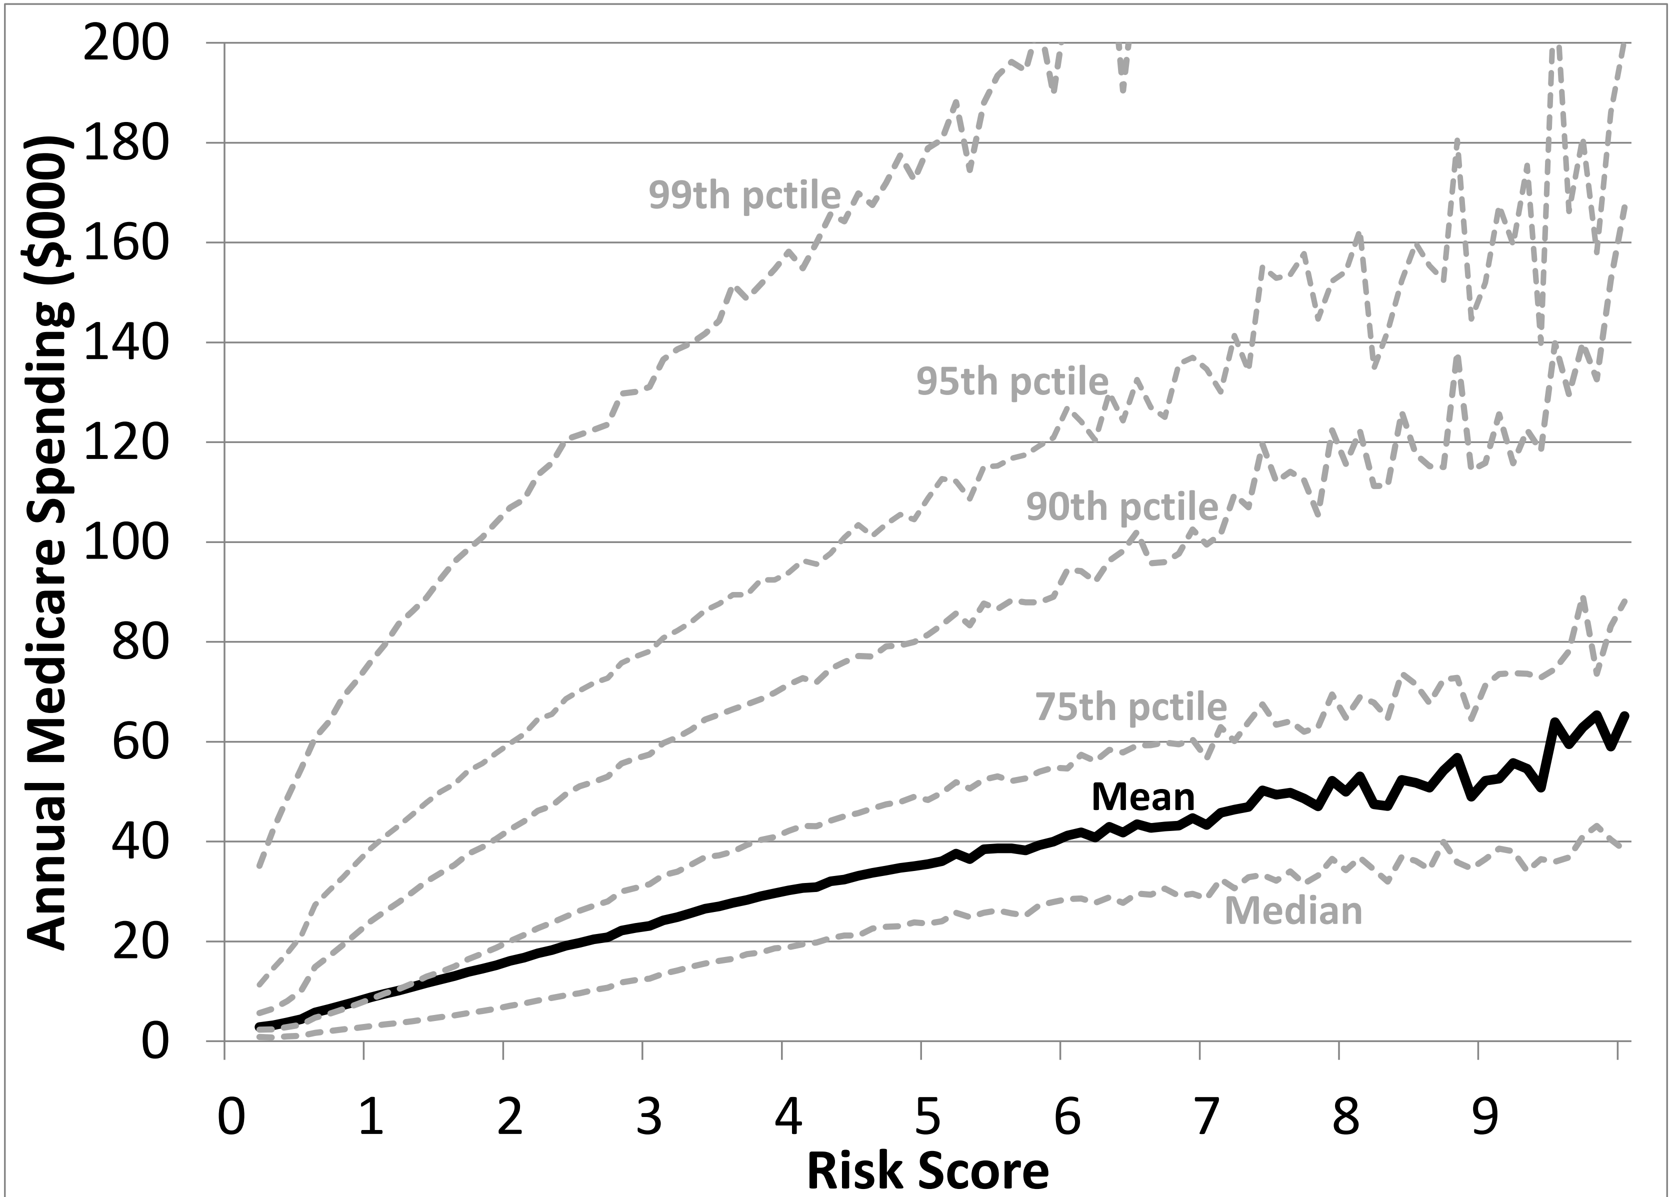

Medicare risk adjustment

- Demographics alone explain a tiny fraction (1-2%), current model about 11%

- Conventional wisdom is that "fully predictive," unconstrained model (see later) could explain ~25%

Outline

- What is risk adjustment , why and how it is being used
- Typical risk adjustment systems
- **Simple empirical framework to think about risk adjustment**
- Selected issues worth thinking about

An empirical framework: primitives

1. F insurers ($f = 1,2, \ldots, F$), each offers a set of J_f coverage contracts that are taken as given, each associated with a base rate/price p_j

2. Individual i is defined by $\{w_i, c_i, r_i\}$:
   - Willingness to pay for each contract: $w_i = \{w_{i1}, w_{i2}, \ldots, w_{iJ}\}$
   - Expected cost (by insurer) for contracts: $c_i = \{c_{i1}, c_{i2}, \ldots, c_{iJ}\}$
   - Predicted cost (by sponsor) for contracts: $r_i = \{r_{i1}, r_{i2}, \ldots, r_{iJ}\}$
     - Natural to assume that $c_i = r_i + e_i$ with $E(e_i) = 0$
     - Predictions get statistically better as $Var(e_i)$ gets smaller, and are perfect when $e_i = 0$

An empirical framework: primitives (cont.)

3. A risk adjustment (and subsidies) system:
   - A voucher for each individual: $\{v_{i1}(p), v_{i2}(p), \ldots, v_{iJ}(p)\}$
   - A transfer system: $\{t(r_{i1}, p), t(r_{i2}, p), \ldots, t(r_{iJ}, p)\}$
     - Important/practical design question: should these depend on prices (trading off information against some incentive issues)
     - One could imagine even richer settings, where transfers depend on composition of pool

→ Subsidies and risk adjustments are often treated separately, even though they probably shouldn't … one shifts demand, the other shifts costs, but all we care is about the difference

An empirical framework: equilibrium

Demand: Individual i chooses contract j (that is, $d_{ij} = 1$) iff

$$w_{ij} - \min\{0, p_j - v_{ij}(p)\} \geq w_{ik} - \min\{0, p_k - v_{ik}(p)\} \qquad \forall k$$

(No system allows it, but one could imagine letting individuals pocket the voucher)

Profits: *(Expected)* Profits from covering *i* by contract *j* given by

$$\pi_{ij} = p_j + t(r_{ij}, p) - c_{ij}$$

Equilibrium: Nash in "prices" would be natural, but obviously one could introduce other forces or other concepts

– Bayesian Nash might be more "realistic" (in many healthcare markets prices are set simultaneously without complete knowledge of rivals' prices), but a full information version (as in Curto, Einav, Levin, and Bhattacharya, 2015) seems (IMO) a decent (and much simpler to work with) approximation
– Could also imagine dynamic pricing due to consumer inertia (as in Miller, 2014)

An empirical framework: market design

Need to decide about a clear objective *W*, and then set $v_i(p)$ and $t(r,p)$ (<u>and define the right geographic/demographic markets!</u>), which would maximize this objective, subject to practical and incentive constraints

Seems natural to think about

$$W = \lambda_C CS + \lambda_P PS - \lambda_G GS$$

where

- *PS* is naturally $\sum_{ij}(d_{ij}\pi_{ij})$
- GS is given by $\sum_{ij}\{d_{ij}[min\{p_j, v_{ij}(p)\} + t(r_{ij}, p)]\}$
- *CS could be the usual sum of wtp, but could also involve other considerations* (e.g., *greater weights on poorer/sicker individuals, or attempt to account for behavioral issues, or paternalistic value)*

Notes:

- In PE it seems common to set $\lambda_G = 0.3\lambda_C = 0.3\lambda_P$ although $\lambda_C >> \lambda_P$ seems more natural in this context
- The 0.3 number is also less obvious in this context (but I don't have a constructive alternative)

Outline

- What is risk adjustment , why and how it is being used
- Typical risk adjustment systems
- Simple empirical framework to think about risk adjustment
- **Selected issues worth thinking about**

Issues I: How restrictive are one-dimensional scores?

- I defined cost types as $c_i = \{c_{i1}, c_{i2}, \ldots, c_{iJ}\}$ and scores as $r_i = \{r_{i1}, r_{i2}, \ldots, r_{iJ}\}$, reflecting the fact that coverage levels and incentives vary across contracts *and* individuals are likely heterogeneous in the nature of health issues they face and their response to price

- Yet, in practice risk scores are always one-dimensional, with occasional (uniform) adjustments for coverage levels

- How important is it? And are there natural way to enrich it in a simple way?

Example from Medicare part D

(from Einav, Finkelstein, Kluender, and Schrimpf, 2015)

- Convex kink in Medicare Part D contracts should lead to "bunching" due to "moral hazard" (as in Saez's and follow-up work on labor supply)

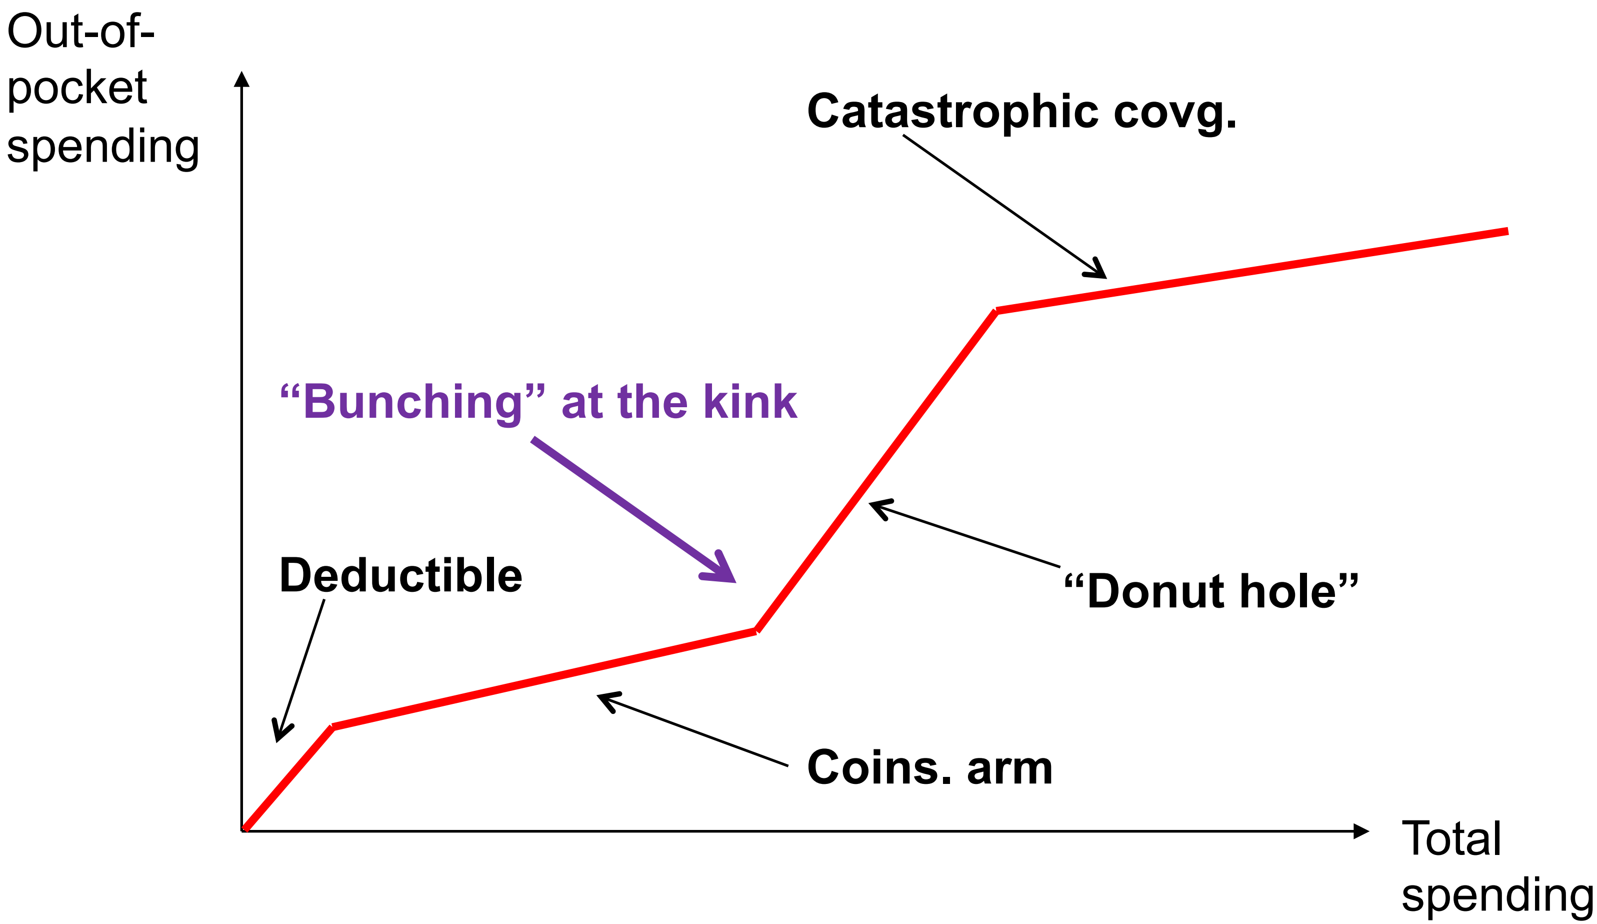

Clear bunching at the convex kink

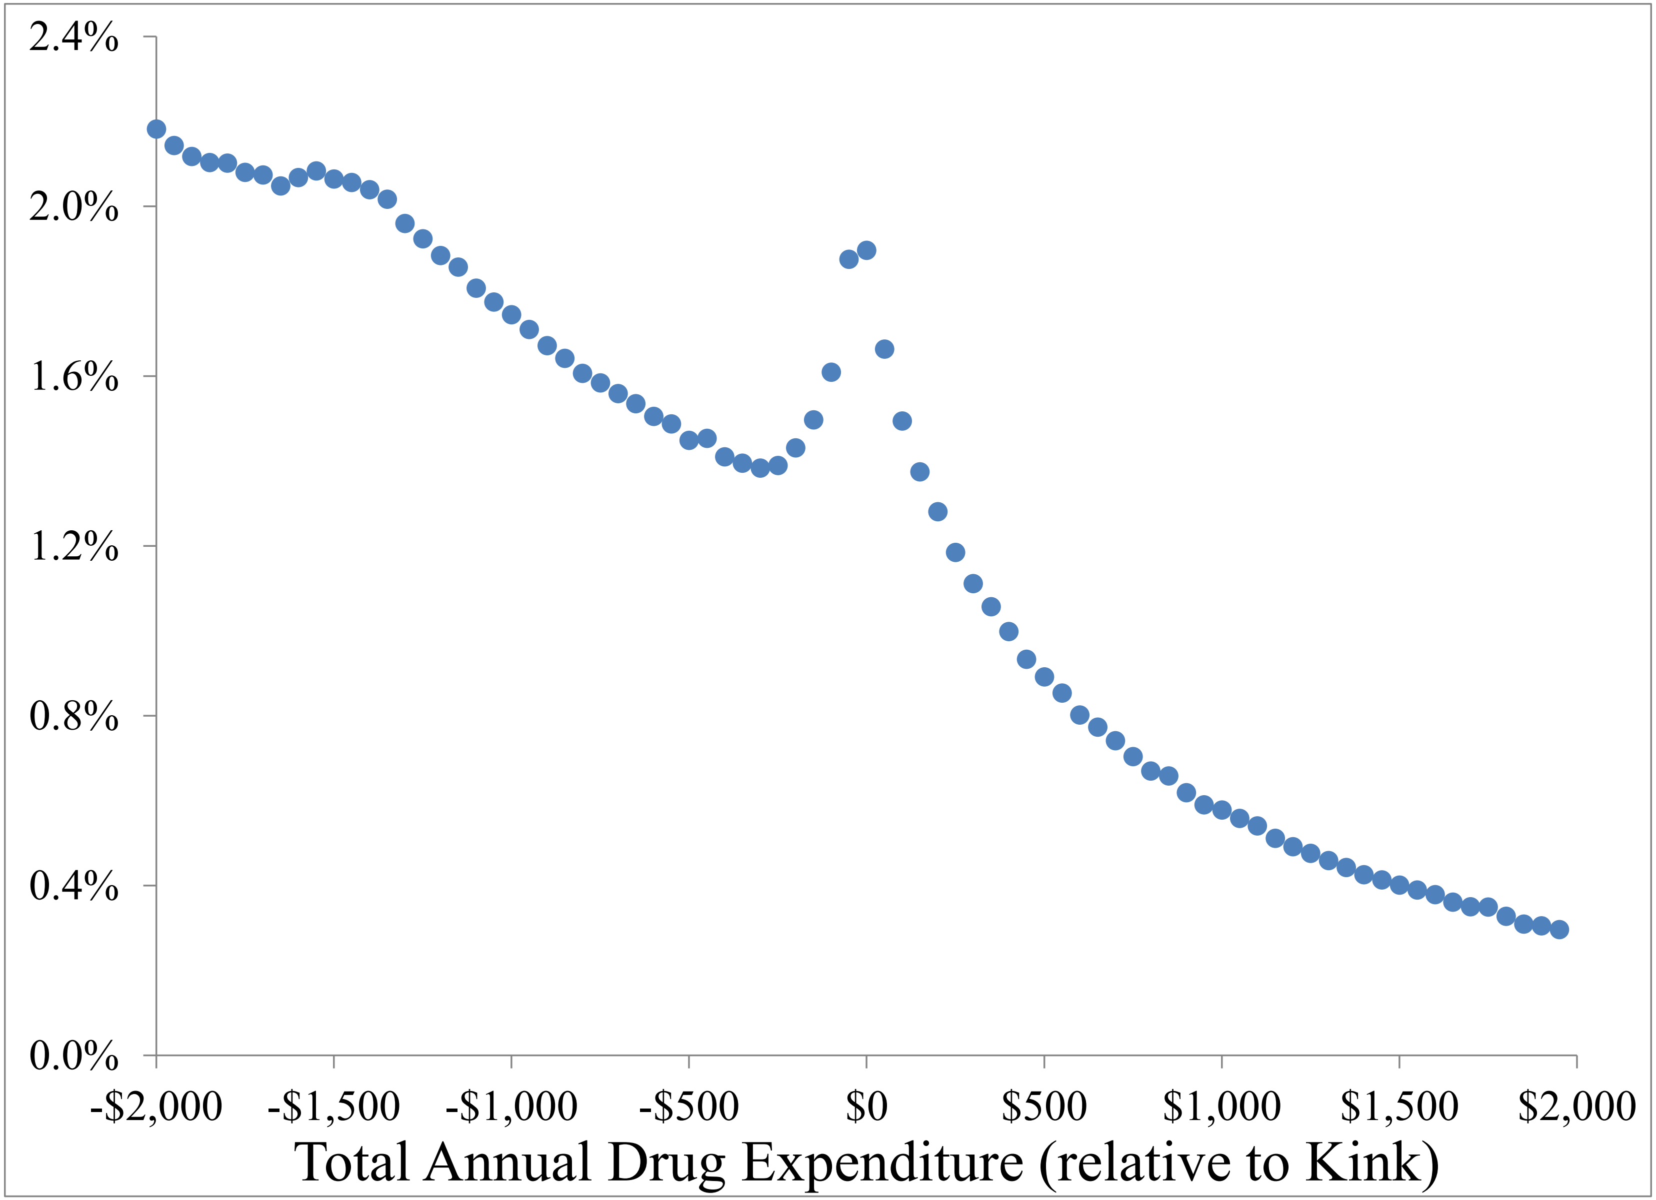

Who is bunching and how they get scored?

- But those who bunch are younger healthier males, who would likely spend significantly more under alternative coverages (e.g., in the absence of the kink)

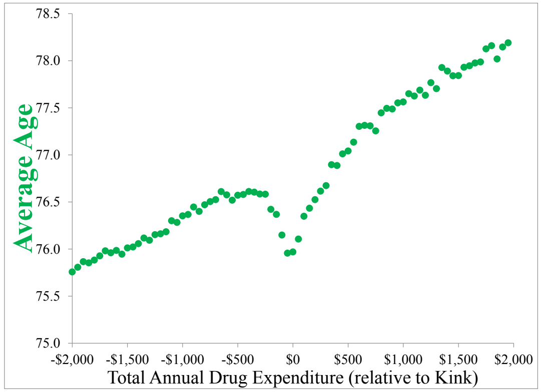

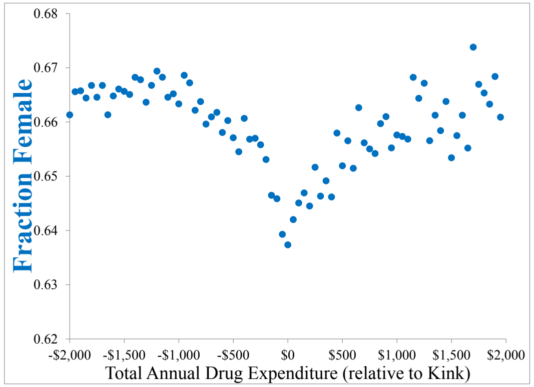

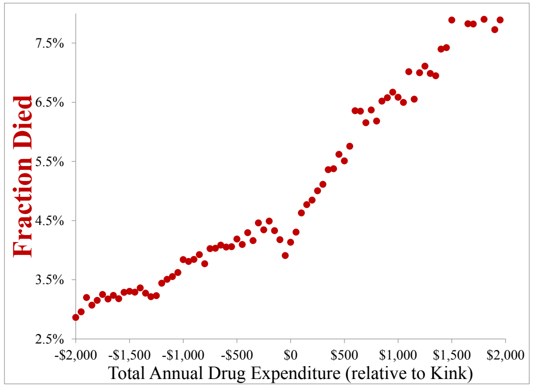

- While risk scores, *by design*, capture none of this:

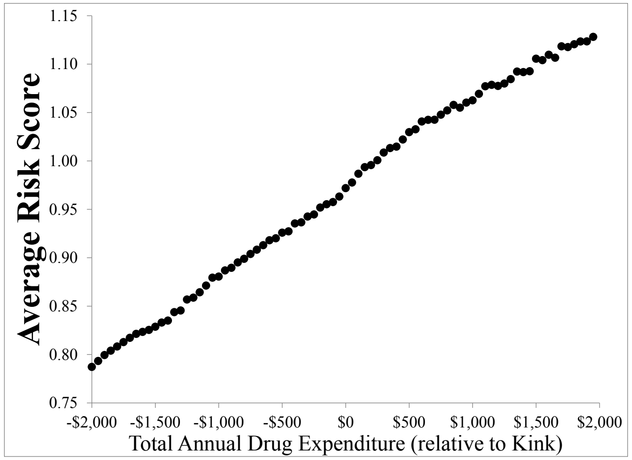

Issues II: Mapping risk scores to transfers

- I defined the transfer system as a function $t(r,p)$. Yet, in practice this is always linear. Is this important?

- When scores are perfect, one-to-one mapping of scores to payments are optimal, but once they are imperfect, it is far less obvious:
  - Glazer and McGuire (2000) show that under certain conditions "over compensating" for risk scores is optimal
  - Layton (2014) and Brown et al. (2014) emphasize the importance of a new objects: instead of focusing on the relationship between *w* and *c*, the focus of market design depends on the relationship between *w* and *c-r*, which could be quite different
  - The quality of the prediction may vary with the level of *r* (e.g., in Medicare it's slightly S-shaped)

Issues III: Endogenous disease coding

- Risk scores are based on claims, and insurers are both those who get paid for covering more risk and those who "produce" claims … so until we get to the point where we have direct, objective, representative measures of health, there are incentive issues

- A fair amount of evidence that insurers, as may be expected, indeed respond to these incentives …
  - When Medicare relied on inpatient claims only, concerns about excess hospitalizations by private Medicare providers emerged (McGuire, Newhouse, and Sinaiko, 2011)
  - Geruso and Layton (2015) provide evidence that private Medicare enrollees have 6-16% higher scores due to more aggressive coding

- What can be done?
  - This imposes some practical constraints on effective risk adjustment, although the optimum is probably to allow some extent of unavoidable "gaming"
  - Raises interesting question about the cost of transparency; a predictive "black box" may be more effective
  - A "black box" may also allow (less equitable?) longer claim histories for scoring, but at least in the US may require a third-party, non-government contractor …

Issues IV: Dynamics …

- So far we viewed healthcare markets as a static object
  - Scores are generated every year based on previous year
  - Then market runs, payment made, and annual coverage begins

- But health depends on behavior, and most sponsored markets have individuals attached to it for many years, so seems natural to use risk adjustments as a way to reward healthy behaviors and penalize less heathy ones …

- Some of this is already in place: ACA exchanges now allow higher premiums for smokers, and some employers (e.g., Safeway) offer premium discounts for health improvements. But it is natural to combine all of this into a risk adjustment system
  - Perhaps certain, more avoidable diseases should get scored less than their predicted impact on cost?

Final slide

- Risk adjustments are likely to play an increasingly important role in healthcare markets

- Way too often, in both policy and academic discussions, they seem to be treated as statistical predictive objects, while in fact they should become an important market design instrument, generating plenty of interesting issues for theoretical and empirical economists to explore

- Although completely unrelated, many of the conceptual issues are quite similar to online platforms, such as Google's AdWords score, or eBay's "Best Match" ranking, which mostly try to predict interest/clickability, but at the same time tailor this to internalize the platform's objective
    - And this all relates more generally to the fairly general question: are more granular/sophisticated pricing lead to better or worse outcomes, and under which conditions?

- Beyond the conceptual interest and the practical relevance and importance, this is also a place for economists to constructively dig into predictive, machine learning algorithms, and potentially incorporate economics into them, so should lead to exciting research opportunities