DS 574
Fall 2026

Health Insurance Markets
Prof. Kira Goldner

Mechanism Design for Social Good

The Process

For each problem:

1. Learn about the area (from domain experts!)—basic rights that people don't have, inefficiencies in government policies, global harms.
2. Hone in on a high-level problem suitable for our tools.
3. Determine the model (with domain experts!)—what to keep, what to abstract.
4. Solve your theory problem.
5. Circle back **with domain experts** about results.

Regulation in Health Insurance Markets

An employer can negotiate with insurance providers and choose which plans to offer. More plans give consumers more options, but making providers compete to be included may lower prices. Could restricting entry improve consumer utility?

Here the tradeoff is between the value of the chosen plan and its price.

The model

There are $n \geq 2$ sellers with zero costs. Seller i posts a nonnegative price p_i. A consumer has values $(v_1, \ldots, v_n)$ drawn independently from the same continuous distribution F on nonnegative values. Each seller can serve any number of consumers.

The consumer must purchase a remaining plan, choosing one that maximizes $v_i - p_i$, even when every available plan gives negative utility. This mandatory-purchase assumption abstracts from an outside option that is prohibitively undesirable.

The market designer chooses between two rules:

1. **Free market:** all sellers enter.
2. **Limited entry:** exclude a highest-priced seller. Break price ties uniformly at random, independently of consumers' values.

Sellers know the rule before choosing prices. We compare expected consumer utility under equilibrium prices, not at a fixed common price vector.

The designer controls entry, sellers control prices, and consumers choose among the surviving plans. Mandatory purchase and zero costs isolate the effect of the entry rule. Correlation across plan values, different costs, and an attractive outside option could change the comparison.

The Limited Entry Setting

If all sellers post zero, no seller can profit by deviating: a positive price makes that seller uniquely most expensive and therefore excluded, while a negative price is infeasible. Thus all-zero prices form an equilibrium.

One seller is then removed uniformly at random. A consumer receives their favorite remaining plan at price zero. Compared with having all n plans available at zero prices, the loss is purely a loss of choice. Compared with the free-market equilibrium, prices can also change; we must analyze sellers' incentives there before comparing utility.

The Free Market Setting: a best response

Fix all other sellers' prices at p, and let seller i try price $q \geq 0$. A consumer buys from i when

$$v_i - q \geq \max_{j \neq i}(v_j - p).$$

Value ties have probability zero. Rearranging isolates the price seller i controls:

$$q \leq v_i - \max_{j \neq i}(v_j - p).$$

Define the consumer's *effective willingness to pay* for seller i by

$$W_i = v_i - \max_{j \neq i}(v_j - p).$$

The best competing offer is absorbed into W_i. For example, if $v_i = 10$ and the best competing utility is 7, then seller i can charge at most 3 and still win (up to tie-breaking). The consumer's value of 10 alone does not determine demand.

Let F_p^* be the CDF of W_i. This distribution depends on the competitors' price p and can assign probability to negative values. Holding p fixed, seller i's expected revenue per consumer is

$$\begin{aligned} q\Pr[\text{consumer buys from } i] &= q\Pr[W_i \geq q] \\ &= q(1 - F_p^*(q)). \end{aligned}$$

Thus a best response solves

$$\max_{q \geq 0} q(1 - F_p^*(q)).$$

This is the monopoly-pricing problem against a new distribution. The reduction does not yet solve the game: at a symmetric equilibrium, posting $q = p$ must itself maximize this expression when the other sellers post p.

Returning to the domain question

Result (stated without proof). Under the paper's MHR$^+$ and non-increasing-density assumptions on F, the free market has a unique symmetric equilibrium, and limited entry gives at least as much expected consumer utility. Uniform and exponential distributions satisfy these assumptions.

These are sufficient conditions on the distribution, not a guarantee for every market. Within this model, the reduction in prices can compensate consumers for the loss of an option.

Here MHR means that the hazard rate $h_F(x) = f(x)/(1 - F(x))$ is nondecreasing. The stronger MHR$^+$ condition requires a constant $c \geq 0$ with $cf(x) \geq -f'(x)$ and $h_F(x) \geq c$ throughout the support. Non-increasing density means f is non-increasing. The result uses the paper's smoothness and finiteness conditions; we do not derive its equilibrium characterization here.

The conclusion answers the original tradeoff within the model: fewer plans need not mean lower consumer utility once sellers adjust their prices. It does not say that fewer plans always help, or that the model captures all features of insurance markets.

References

Meryem Essaidi, Kira Goldner, and S. Matthew Weinberg. *To Regulate or Not to Regulate: Using Revenue Maximization Tools to Maximize Consumer Utility.* SAGT, 2024. Preprint: *When to Limit Market Entry under Mandatory Purchase*, arXiv:2002.06326.