

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 [Essaidi G Weinberg [1]]

Background. There are two main types of insurance markets that operate in the United States. The first is like Medicare Advantage, private insurance companies like Aetna or Blue Cross Blue Shield cover patients (and the government subsidizes them), but the government does not actually regulate entry into the market. The idea here is that a lot of competition should lead to lower prices of insurance plans, but what we actually wind up seeing is the counterintuitive opposite. Essentially, there are high barriers to entry in health insurance—setting up a network of doctors, staff to manage care, etc.—and when we examine the actually Medicare advantage markets, the majority of consumers in most states choose from three providers, enabling them to set higher prices as they control the market.

The other main kinda of insurance market is that created by e.g. an employer like BU. BU will negotiate with several different insurance plans, and come up with a way to select insurance plans so that they have better deals for their employees on the ones that they contract with. Then, within BU's insurance market, employees are free to select whichever insurance plan they like at the rates that BU has negotiated. Here, BU acts as the mechanism designer, or the market maker, determining which insurance providers get to enter the market and how.

The question posed to my coauthors and I by a health insurance expert was: if the government stepped in and regulated the Medicare market more like an employer, could they

bring down prices and lower barriers to entry in a way that would improve consumer welfare? Intuitively, it should lower prices, but it would also cut out options, so it is unclear what will have the best impact on consumer utility.

Setting. There are a bunch of sellers (e.g. insurance plans) who wish to enter the market and sell their product. Each plan i will simply post a price p_i that consumers can pay to purchase their product (plan).

A consumer has a value for each product which is drawn i.i.d. from some distribution F . Then a specific realization of a consumer has values (v_1, \dots, v_m) for products $1, \dots, m$.

A consumer will purchase the product that maximizes their utility, $i \in \operatorname{argmax} v_i - p_i$. They *must* purchase a product even if their utility is negative for all products. This is motivated by settings where purchase is required by law, or is prohibitively expensive for their utility not to purchase (e.g. health repercussions).

The mechanism designer can choose how best to regulate entry into this market, if at all. In particular, we compare two specific choices:

1. No regulation. This is the “free market” model that is used in markets like medicare where anyone can enter into the market.
2. Allow all but the most expensive seller/product to enter the market. This “limited entry” model corresponds to procurers who negotiate with the sellers and then choose a smaller subset of curated products. We start with the most general case.

We will study both settings and try to understand first how the agents act strategically *in equilibrium*, and specifically, we would like to study a unique symmetric equilibrium. The motivation for this is that (1) since consumers are a priori identical toward sellers, it makes sense that sellers would have the same strategies, and (2) we better only have one equilibrium to know for certain what to analyze next. Then, we use these strategies to analyze consumer welfare (consumers expected utility) under equilibrium strategies, and compare across the two different regulation strategies.

Notation. Let X_k^n denote the k^{th} highest value from n i.i.d. draws. That is, if we draw x_1, x_2, \dots, x_n i.i.d., then $\max\{x_1, \dots, x_n\} = X_1^n$.

Let V_k^n denote the expected value of the k^{th} highest value from n i.i.d. draws. That is, $V_k^n = \mathbb{E}[X_k^n]$.

Recall that $\varphi_F(v) = v - \frac{1-F(v)}{f(v)}$ is the Myerson virtual value. the inverse of this second term is called the *hazard rate*: $\frac{f(v)}{1-F(v)}$.

We will use $h_2^n(F)$ to denote the expected hazard rate of the second maximum out of n

draws from F , that is, $h_2^n(F) = \mathbb{E} \left[\frac{f(X_2^n)}{1-F(X_2^n)} \right]$.

We will use $H_1^n(F)$ to denote the expected inverse of the hazard rate of the maximum out of n draws from F , that is, $H_1^n(F) = \mathbb{E} \left[\frac{1-F(X_1^n)}{f(X_1^n)} \right]$.

The Limited Entry Setting. In this setting, one product is eliminated from the market. *Bertrand competition* here indicates that there is a unique equilibrium for seller prices, where all sellers set prices at $p_i = 0 \forall i$, and then a random seller is removed.

With all prices being identical and a seller being randomly removed, a consumer will simply choose the remaining product for which they have the highest value. Their value is precisely the realization of V_1^{n-1} , as there are $n - 1$ products remaining. That is, their expected equilibrium utility of value minus price is

$$V_1^{n-1} - 0.$$

The Free Market Setting. The difficulty here is to try to reason about how a seller will strategize *in response* to the other sellers' strategies. Our key idea will be to just set up the problem pretending that the seller is a monopolist, because we already know how to solve that problem.

In the classic setting with a single revenue-maximizing seller where $v \sim F$,

$$\begin{aligned} \mathbb{E}[\text{revenue}] &= \Pr[\text{customer willing to pay } q] \cdot q \\ &= [1 - F(q)] \cdot q \end{aligned}$$

Then to maximize revenue, the seller would set the Myerson price of $p = \varphi_F^{-1}(0)$. This is the solution to $\text{argmax}_p \{p \cdot (1 - F(p))\}$.

Our method will be to define a new distribution F_p^* where

$$1 - F_p^*(q) = \Pr[\text{patient purchases from } i \mid i \text{ posts } q, \text{ others post } p].$$

Our goal is to calculate

$$\begin{aligned} \text{Provider } i\text{'s } \mathbb{E}[\text{payoff}] &= \Pr[\text{patient purchases from } i] \cdot p_i \\ &= [1 - F_p^*(q)] \cdot q \end{aligned}$$

Then seller i 's *best response* is to set the Myerson price of this distribution $q = \varphi_{F_p^*}^{-1}(0)$.

Back in the classic setting, when is this solution *unique*? The answer is that MHR of F is a sufficient condition to imply uniqueness. Observe that $\varphi_F^{-1}(0)$ is the set of all v such that $v = \frac{1}{h_F(v)}$, since we can rewrite $\varphi_F(v) = v - \frac{1-F(v)}{f(v)} = v - \frac{1}{h_F(v)}$. Well, v is strictly increasing in v , and if F is MHR, by definition, then $\frac{1}{h_F(v)}$ is weakly decreasing in v , hence $\varphi_F(v)$ is

strictly increasing in v . Therefore, $v = 1/h_F(v)$ cannot have multiple solutions, and thus is unique.

In our setting, we thus want F_p^* to be MHR (as a sufficient condition) for this equilibrium to be unique. But then what condition do we need on the original distribution F ? The result is a restriction of MHR that we call “MHR⁺” as well as decreasing density. What does this mean exactly?

To have a *symmetric* equilibrium, the best response of the seller must be setting $q = p$. Such an equilibrium exists so long as F_p^* is regular (which is more general than MHR).

Then we do have a unique symmetric equilibrium with $p = \frac{1}{h_2^n(F)}$ so long as F is MHR⁺. Since all seller prices are identical, then a consumer will just choose the product for which they have the highest value, and so their expected equilibrium utility will be value minus price:

$$V_1^n(F) - \frac{1}{h_2^n(F)}.$$

Comparing Free Market and Limited Entry. Then the limited entry setting has higher expected utility in equilibrium for consumers over the free market setting precisely when

$$V_1^{n-1}(F) \geq V_1^n(F) - \frac{1}{h_2^n(F)}.$$

Instead of directly sampling $n - 1$ values and taking the max, as the left-hand side now suggests, we instead sample n values and remove one uniformly at random to get down to $n - 1$ values. We now consider the maximum of these $n - 1$ values. If the sample removed was any but the original highest, so for $n - 1$ out of n of the possible choices of a sample to remove, the maximum is the original maximum V_1^n . If the original maximum was removed (with probability $\frac{1}{n}$, then the maximum of $n - 1$ is V_2^n . We replace the left-hand side accordingly.

$$\frac{n-1}{n}V_1^n(F) + \frac{1}{n}V_2^n(F) \geq V_1^n(F) - \frac{1}{h_2^n(F)}.$$

Now, we draw on Myerson again, which tells us for i.i.d. bidders that the second-price auction with a reserve price is revenue-optimal. That is,

$$V_2^n(F) = \mathbb{E}[\varphi_F(X_1^n(F))] = V_1^n(F) - H_1^n(F).$$

Plugging this into what we already had and grouping like terms, we get

$$\frac{1}{n}(V_1^n(F) - H_1^n(F)) \geq \frac{1}{n}V_1^n(F) - \frac{1}{h_2^n(F)}.$$

which implies the Limit Entry Condition:

$$H_1^n(F) \leq \frac{1}{h_2^n(F)}.$$

This is now just a condition about hazard rates, so about the tails of distributions, and so we can essentially look at our population and say, if it satisfies these conditions, then we know that limiting entry will improve consumer welfare.

Specifically, F is MHR (or MHR^+), then limiting entry is best. This corresponds in essence to having an exponential tail or lighter, and includes distributions such as the exponential, uniform, and Gaussian distributions. We require F to be MHR^+ and have decreasing density in order to have a symmetric unique equilibrium to analyze, so these conditions together imply that limiting entry is best.

References

- [1] Meryem Essaidi, Kira Goldner, and S Matthew Weinberg. When to Limit Market Entry under Mandatory Purchase. *arXiv preprint arXiv:2002.06326*, 2020.