Economics 696F: Lecture Note 15

Treatment Assignment as a Statistical Decision Problem

Note: this is based on Manski (2004); however, I have changed the notation considerably.

Basic Setup

We imagine a social planner who must assign an individual to one of two possible treatments. The planner observes some background characteristics of the individual and can base the treatment assignment on these characteristics.

$\mathcal{T} := \{0, 1\}$: set of possible treatments. (Note: could extend to multivalued treatments.)

$Y(0), Y(1)$: potential outcomes for the individual.

X: background characteristics.

Let $F(x, y(0), y(1)|\theta)$ denote the joint distribution of $(X, Y(0), Y(1))$. Here θ is a parameter describing the distribution. Set of possible parameter values is Θ.

Denote marginal distribution of X as $F_X(\cdot|\theta)$, and conditional distributions of $Y(0)$ given X and $Y(1)$ given X as $F_0(\cdot|x, \theta)$ and $F_1(\cdot|x, \theta)$. So

$$X \sim F_X(\cdot|\theta),$$

$$Y(0)|X = x \sim F_0(\cdot|x, \theta),$$
$$Y(1)|X = x \sim F_1(\cdot|x, \theta).$$

For simplicity, assume X takes on a finite set of possible values $\{x_1, \ldots, x_J\}$. Then denote

$$f_X(x_j|\theta) := Pr(X = x_j|\theta).$$

Treatment Assignment Rule and Social Welfare

A treatment assignment rule is a function that selects treatment based on X. We want to allow for randomized rules, so think of a rule as giving, for any value of X, a probability of assigning to treatment 1:

$$\delta(x) := Pr(\text{assign } T = 1|X = x).$$

Therefore

$$Pr(\text{assign } T = 0|X = x) = 1 - \delta(x).$$

We can think of a given rule $\delta(\cdot)$ as a member of some set Δ of possible/feasible treatment assignment rules.

In the case where X takes on J possible values, any rule is completely characterized by the J values

$$\delta(x_j) = Pr(\text{assign } T = 1 | X = x_j).$$

Thus, we can think of a rule δ as a $J-$ vector, with each component being some number between 0 and 1.

Social welfare: let $U(y,t,x)$ be the utility of getting treatment t and outcome y for an individual with characteristic x. For simplicity, let's just work with the special case where Y is measured in utils:

$$U(y,t,x) = y.$$

Then the expected utility of treatment t for an individual with characteristic x is

$$E_\theta[U(Y(t),t,x)] = E_\theta[Y(t)|X = x].$$

Here the θ subscript means take the expectation with respect to the distribution $F(\cdot|\theta)$.

So the mean utility of a rule $\delta(\cdot)$ is

$$Pr(T = 1|X = x) \cdot E[U(Y(1), 1, x)] + Pr(T = 0|X = x) \cdot E[U(Y(0), 0, x)]$$

$$= \delta(x) \cdot E_\theta[Y(1)|X = x] + (1 - \delta(x)) \cdot E_\theta[Y(0)|X = x].$$

We can think of applying the rule repeatedly over the population, yielding an overall mean utility of

$$\sum_{j=1}^{J} f_X(x_j|\theta)\Big\{\delta(x_j) \cdot E_\theta[Y(1)|X = x_j] + (1 - \delta(x_j)) \cdot E_\theta[Y(0)|X = x_j]\Big\}.$$

This is a utilitarian measure of social welfare: we are counting individuals "equally" and simply adding up their utilities.

We want to choose a rule $\delta(\cdot)$ to maximize social welfare.

It is easy to see that the optimal rule would set, for each j:

$$\delta^*(x_j) = \begin{cases} 1 & \text{if } E_\theta[Y(1)|X = x_j] > E_\theta[Y(0)|X = x_j], \\ 0 & \text{if } E_\theta[Y(1)|X = x_j] < E_\theta[Y(0)|X = x_j]. \end{cases}$$

If $E_\theta[Y(1)|X = x_j] = E_\theta[Y(0)|X = x_j]$, then any value for $\delta^*(x_j) \in [0, 1]$ is optimal.

Of course, in order to calculate this rule, we would need to know θ. For different values of θ, different rules might be optimal.

Since we usually do not know θ (at least not perfectly), it is not possible to perfectly implement the optimal rule δ^*.

But, we might have some data, say from a randomized experiment, that is informative about θ. So the question is how to use this past data to inform our policy rule.

Statistical Treatment Rule

Suppose that before making our treatment assignment decision, we observe some data Z, distributed as $Q(z|\theta)$. We assume Z is independent of the future individual we are making a decision about.

The idea is that Z is informative about θ, hence about whether treatment 1 or 0 is better for different types of individuals. We can base our treatment assignment on the value of Z as well as the individual's characteristics x.

Note the timing:

1. We run a randomized experiment or some other study which gives us Z.
2. Then we take a new individual, and observe their X.
3. We assign this individual to treatment based on her own X as well as the data of others collected in Z.

A statistical treatment rule is a function

$$\delta(x, z) = Pr(\text{assign } T = 1|X = x, Z = z).$$

This gives an ex ante probability of assigning individuals with $X = x$ to treatment as

$$E_\theta[\delta(x, Z)] = \int \delta(x, z) dQ_\theta(z).$$

You could think of this as the overall probability if we were to re-run Step 1 many times, generating new data sets Z and basing our decision on the realized Z.

We can think of an ex-ante expected social welfare for a given rule δ, defined as

$$W(\theta, \delta) := \int \sum_{j=1}^{J} f_X(x_j|\theta) \Big\{ \delta(x_j) \cdot E_\theta[Y(1)|X = x_j] + (1 - \delta(x_j)) \cdot E_\theta[Y(0)|X = x_j] \Big\} dQ(z|\theta).$$

This can be simplified to:

$$W(\theta, \delta) = \sum_{j=1}^{J} f_X(x_j|\theta) E_\theta[\delta(x_j, Z)] \cdot E_\theta[Y(1)|X = x_j] + (1 - E_\theta[\delta(x_j, Z)]) \cdot E_\theta[Y(0)|X = x_j].$$

Ordering Statistical Decision Rules

Introducing the historical data Z doesn't really resolve the issue. Some rules might work better for some values of θ, while other rules work better for other values of θ.

In order to make progress, we need to somehow construct an overall ranking of rules that doesn't depend on θ. There are two ways to do this:

One approach is to average $W(\theta, \delta)$ over $\theta \in \Theta$ in some way. Suppose we have a function $\pi(\theta)$ that 'weights' the different values of θ, and we define

$$w_\pi(\delta) := \int_\Theta W(\theta, \delta)\pi(\theta) d\theta.$$

This measures the performance of δ by taking a weighted average of its performance at different possible values of θ. We can then try to select the rule that mazimizes this measure of performance:

$$\delta_b := \arg\max_\delta w_\pi(\delta).$$

The function w_π is called the Bayes welfare, and the rule δ_b is called a Bayes rule. The weighting function π is also often called a "prior density."

For this approach, we need to choose π, and different choices will lead to different optimal rules, in general.

Another approach is to measure the performance of a rule by its worst-case welfare

$$\min_{\theta \in \Theta} W(\theta, \delta).$$

We could then look for a rule that is best in this sense:

$$\delta_m := \arg\max_\delta \min_{\theta \in \Theta} W(\theta, \delta).$$

A rule δ_m defined as above is called a maxmin rule.

(To be technically correct, we should replace "max" with "sup" and "min" with "inf". We still call the rule maxmin by tradition.)

The maxmin approach is often very conservative. We may be led to choose a strange rule in order to avoid bad performance at certain very extreme parameter values.

An alternative that is usually less conservative is based on the notion of regret. Basically, we compare the performance of our rule to the performance of the infeasible optimal rule δ^* defined above.

The regret of a rule δ is defined as

$$W(\theta, \delta) - W(\theta, \delta^*).$$

Then a maxmin-regret rule satisfies

$$\delta_{mr} := \arg\max_{\delta} \min_{\theta \in \Theta} [W(\theta, \delta) - W(\theta, \delta^*)].$$

It is generally difficult to solve the maxmin and maxmin-regret problems. Consider the maximization over δ. If Z is continuous then the space of possible $\delta(x, z)$ functions is infinite-dimensional. Moreover, we do not necessarily want to impose smoothness on the δ functions.

(The Bayes rule is actually easier to compute, using numerical methods. Unfortunately we don't have time to go into Bayesian methodology, but there is a rich literature and a number of specialized numerical tools that can be used.)

Special Case: Stratified Randomized Experiment

One example is when Z represents data from a randomized experiment.

For each possible value of x_j, draw N_j units randomly from the subpopulation with $X = x_j$, and assign N_j^1 to treatment 1, $N_j^0 = N_j - N_j^1$ to treatment 0. So we observe, for each j,

$$Y_{ji}, T_{ji}, \quad i = 1, \ldots, N_j$$

with

$$\sum_{i=1}^{N_j} T_{ji} = N_j^1, \quad \sum_{i=1}^{N_j} (1 - T_{ji}) = N_j^0.$$

So the overall data we have is

$$Z := \{(Y_{ji}, T_{ji}), \quad j = 1, \ldots, J, \quad i = 1, \ldots, N_j\}.$$

Conditional Empirical Success Rule

Suppose we have data from a randomized experiment. Then a natural treatment assignment rule is to calculate the average outcome under the two treatments for individuals with a given value of X, and assign future individual to whichever treatment appears to do better. Formally, let

$$\hat{\beta}_j := \frac{1}{N_j^1} \sum_{i=1}^{N_j} T_{ji} Y_{ji} - \frac{1}{N_j^0} \sum_{i=1}^{N_j} (1 - T_{ji}) Y_{ji}.$$

Then define

$$\hat{\delta}(x_j) = 1(\hat{\beta}_j > 0).$$

Manski calls this a conditional empirical success (CES) rule.

Under the condition that the range of Y is bounded, he develops bounds on $W(\theta, \hat{\delta})$ and $W(\theta, \hat{\delta}) - W(\theta, \delta^*)$.