Economics 696: Lecture Note 11

Treatment Assignment as a Statistical Decision Problem

Note: this is based on Manski (2004) and Dehejia (2005); 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 $\theta \in \Theta$ is a parameter vector describing the distribution.

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

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

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

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

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.

Social welfare: let $U(y, t, x)$ be the utility of getting treatment t and outcome y for an individual with characteristic x. Then the expected utility of treatment t for an individual with characteristic x is

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

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

$$\delta(x) \cdot E_\theta[U(Y(1), 1, x)] + (1 - \delta(x)) \cdot E_\theta[U(Y(0), 0, x)]$$

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

$$\int \Big\{ \delta(x) \cdot E_\theta[U(Y(1), 1, x)] + (1 - \delta(x)) \cdot E_\theta[U(Y(0), 0, x)] \Big\} dF_X(x|\theta).$$

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:

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

If $E_\theta[U(Y(1), 1, x)] = E_\theta[U(Y(0), 0, x)]$, then any value for $\delta^*(x) \in [0, 1]$ 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

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

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

$$W(\theta, \delta) := \int \int \Big\{ \delta(x, z) \cdot E_\theta[U(Y(1), 1, x)] + (1 - \delta(x, z)) \cdot E_\theta[U(Y(0), 0, x)] \Big\} dF_X(x|\theta) dQ(z|\theta).$$

This can be simplified to:

$$\begin{aligned} W(\theta, \delta) &= \int \{E_\theta[\delta(x, Z)] \cdot E_\theta[U(Y(1), 1, x)] + (1 - E_\theta[\delta(x, Z)]) \cdot E_\theta[U(Y(0), 0, x)]\} \, dF_X(x|\theta) \\ &= \int \{\beta(\delta, x, \theta) \cdot E_\theta[U(Y(1), 1, x)] + (1 - \beta(\delta, x, \theta)) \cdot E_\theta[U(Y(0), 0, x)]\} \, dF_X(x|\theta) \end{aligned}$$

Ordering Statistical Decision Rules: Bayes Welfare

First consider a Bayesian approach. Suppose we place a prior Π on the parameter θ, with density function $\pi(\theta)$. Then we can define the "Bayes welfare"

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

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 maximization can be simplified as follows. Suppose that $Q(z|\theta)$ has a density function $q(z|\theta)$. (This is the likelihood function of the data z.) Then we can write:

$$w_\pi(\delta) = \int_\Theta \int_\mathcal{Z} \int_\mathcal{X} \{\delta(x,z) E_\theta[U(Y(1),1,x)] + (1-\delta(x,z)) E_\theta[U(Y(0),0,x)]\} \, dF_X(x|\theta) q(z|\theta) dz \pi(\theta) d\theta.$$

Change the order of integration:

$$w_\pi(\delta) = \int_\mathcal{Z} \int_\mathcal{X} \int_\Theta \{\delta(x,z) E_\theta[U(Y(1),1,x)] + (1-\delta(x,z)) E_\theta[U(Y(0),0,x)]\} \, q(z|\theta) \pi(\theta) d\theta dF_X(x|\theta) dz.$$

Then, for each value of (x,z), we can choose δ to minimize the inner term

$$\int_\Theta \{\delta E_\theta[U(Y(1),1,x)] + (1-\delta) E_\theta[U(Y(0),0,x)]\} \, q(z|\theta)\pi(\theta) d\theta.$$

Doing so (for any relevant value of (x,z)) will solve the Bayes problem. However, since the posterior density satisfies

$$p(\theta|z) \propto q(z|\theta)\pi(\theta),$$

we can equivalently minimize

$$\int_\Theta \{\delta E_\theta[U(Y(1),1,x)] + (1-\delta) E_\theta[U(Y(0),0,x)]\} \, p(\theta|z) d\theta.$$

This is the posterior expected welfare given data z and individual characteristics x. Since this can also be written as

$$\delta \int_\Theta E_\theta[U(Y(1),1,x)] p(\theta|z) d\theta + (1-\delta) \int_\Theta E_\theta[U(Y(0),0,x)] p(\theta|z) d\theta,$$

it is easy to see that the solution will typically have either $\delta = 0$ or $\delta = 1$, depending on whether treatment 1 or treatment 0 gives higher expected welfare.

Dehejia (2005)

Dehejia examines the GAIN experiment, a randomized evaluation of a job training program in California. The GAIN program was an alternative to standard AFDC (welfare), which involved basic education, job training, and job search. He has data from the Alameda County portion of the GAIN experiment, and observes for each individual a number of background characteristics, whether they received the GAIN intervention or standard AFDC, and earnings over a number of quarters.

Since many individuals had zero earnings, he uses a Tobit model. Let earnings for individual i in quarter t after the experiment be denoted as y_{it}. The model specifies that individuals are independent, with

$$y_{it}^* = w_{it}'\beta_1 + d_i \cdot w_{it}'\beta_2 + r_{it}'\beta_3 + \epsilon_{it}.$$

Here w_{it} is a vector of background covariates, including dummies for quarter t, d_i is an indicator for receiving the GAIN intervention, and r_{it} are additional variables such as calendar time. We assume the ϵ_{it} are IID $N(0, \sigma^2)$ conditional on the regressors, and observed earnings are given by

$$y_{it} = 1(y_{it}^* > 0) y_{it}^*.$$

Dehejia then considered the decision problem facing a counselor who had to assign individuals to either AFDC or GAIN based on their covariates, with the goal of maximizing social welfare.

To put this in our framework, we define the parameter vector as $\theta = (\beta, \sigma^2)$. The data are:

$$Z = \{(w_{it}, y_{it}) : i = 1, \ldots, n, t = 1, \ldots, T\}.$$

The likelihood function $q(z|\theta)$ is the Tobit conditional likelihood. (We implicitly make the partial likelihood assumptions so that working with the conditional likelihood is appropriate.) Then the posterior distribution $p(\theta|z)$ can be simulated using the DA/Gibbs sampler.

We assume that the counselor is dealing with a new individual (person $n+1$), whose covariates are observed and whose earnings will follow the same distribution as the experimental subjects $i = 1, \ldots, n$. Thus, X are the covariates observed for individual $n+1$ in period t:

$$X = (w_{n+1,t}, r_{n+1,t}).$$

Earnings under AFDC is distributed according to the Tobit model with $d_i = 0$:

$$y_{n+1,t,0}^* = w_{n+1,t}' \beta_1 + r_{n+1,t}' \beta_3 + \epsilon_{n+1,t}.$$

$$Y(0) = 1(y_{n+1,t,0}^* > 0) y_{n+1,t,0}^*.$$

Similarly, earnings under GAIN can be written as:

$$y_{n+1,t,1}^* = w_{n+1,t}' (\beta_1 + \beta_2) + r_{n+1,t}' \beta_3 + \epsilon_{n+1,t}.$$

$$Y(1) = 1(y_{n+1,t,1}^* > 0) y_{n+1,t,1}^*.$$

As we saw above, the Bayes solution can be obtained by calculating the posterior expected utilities from treatment 1 and 0:

$$E[U(Y(1), 1, x)|z] = \int_\Theta E_\theta[U(Y(1), 1, x)] p(\theta|z) d\theta$$

and

$$E[U(Y(0), 0, x)|z] = \int_\Theta E_\theta[U(Y(0), 0, x)] p(\theta|z) d\theta.$$

For example, let us assume that $U(Y, d, x) = \log(Y)$, log utility. Then

$$\begin{aligned} E[U(Y(1), 1, x)|z] &= \int_{\Theta} E_\theta[U(Y(1), 1, x)]p(\theta|z)d\theta \\ &= \int_{\Theta}\int_0^\infty \log(y)f_1(y|x,\theta)p(\theta|z)dyd\theta \\ &= \int_0^\infty \log(y)f_{1,p}(y|x,z)dy. \end{aligned}$$

Here, f_1 is the conditional density of $Y(1)$ given by the Tobit model, and $f_{1,p}$ is the predictive density

$$f_{1,p} = \int f_1(y|x,\theta)p(\theta|z)d\theta.$$

A similar argument works for outcome under treatment 0. Hence, it is key to be able to simulate from the predictive distribution of earnings given covariate and treatment choice, integrating over the posterior distribution.

Ordering Decision Rules: Minmax and Minmax Regret

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.

Manski (2004): Conditional Empirical Success Rule

Suppose we have data from a randomized experiment, and the covariate X is discrete, taking on possible values $\{x_j : j = 1, \ldots, k\}$. Also, suppose that we are interested in the average outcome:

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

We suppose that data are obtained from a stratified 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

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

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 $W(\theta, \hat{\delta}) - W(\theta, \delta^*)$. Further work by Stoye (2006) and Schlag (2006) develop procedures that are minmax-regret optimal.