

IEOR 265 – Lecture 19

Estimating Multiple Utilities

1 Multiple Strategic Agents

Recall the following abstract model: Suppose that we have $p \geq 2$ agents, and the k -th agent is making decisions $x^{*,k} \in \mathcal{X}^k(u_i)$ to maximize their utility function while also taking into account the strategic behavior of each other agent. In this model, we need to also specify our notion of strategic behavior. Here, we will restrict our attention to the case of a Nash equilibrium, in which

$$x_i^{*,k} \in \arg \max \{J^k(x_i^{*,1}, \dots, x_i^{*,k-1}, x^k, x_i^{*,k+1}, \dots, x_i^{*,p}, u_i) \mid x^k \in \mathcal{X}^k(u_i)\}.$$

We observe $(u_i, x_i^{*,1}, \dots, x_i^{*,p})$ for $i = 1, \dots, n$ and would like to infer the utility function of each agent J^k .

To make this model more concrete, we will specify a specific instantiation of this problem. In particular, suppose that

- The constraint set $\mathcal{X}^k(u_i)$ is described by linear equality and inequality constraints:

$$\mathcal{X}^k(u_i) = \{x : x \geq 0, F^k x + G^k u_i = h^k\},$$

where (A^k, b^k) and (F^k, h^k) are suitably defined matrices and vectors.

- Assume that we have a parametrization of the utility function, that is we have

$$\phi(x^1, \dots, x^p, u; \beta),$$

and a compact set Γ such that there exists $\beta_k \in \Gamma$ with

$$J^k(x^1, \dots, x^p, u) = \phi(x^1, \dots, x^p, u; \beta_k).$$

Though these two conditions make the problem more specific, we will still impose additional conditions on the model formulation to make the problem computationally tractable.

2 Key Technical Difficulty

Consider the following feasibility problem formulation of the inverse decision-making problem for this multiple strategic agents model:

$$\begin{aligned} \hat{\beta} &= \arg \min_{\beta} 0 \\ \text{s.t. } &x_i^{*,k} \in \arg \max \{ \phi(x_i^{*,1}, \dots, x_i^{*,k-1}, x^k, x_i^{*,k+1}, \dots, x_i^{*,p}, u_i; \beta) \mid \\ &x \geq 0, F^k x + G^k u_i = h^k \} \\ &\beta \in \Gamma. \end{aligned}$$

This feasibility problem is difficult to solve because it has an atypical constraint: The constraint that x_i^* be the equilibrium point of a game-theoretic model cannot be directly handled by nonlinear programming techniques. There are two reasons that this constraint presents challenges:

1. Depending on the value of β there may be one or multiple equilibria to the game.
2. The constraints have a complex form because it corresponds to equilibrium points to a game, and so we cannot in general hope for continuity of the equilibrium with respect to β , much less differentiability.

3 Variational Inequalities

Similar to how KKT conditions represent optimality conditions for a nonlinear programming problem, there is an alternative representation for the Nash equilibrium of a game. Let

$$X = \begin{bmatrix} x^1 \\ \vdots \\ x^p \end{bmatrix}$$

and

$$\mathcal{X}(u) = \mathcal{X}^1(u) \times \cdots \times \mathcal{X}^p.$$

Assume that the utilities for each agent are differentiable, and define

$$D(X, u; \beta) = \begin{bmatrix} -\nabla_{x^1} \phi(X, u; \beta_1) \\ \vdots \\ -\nabla_{x^p} \phi(X, u; \beta_p) \end{bmatrix}.$$

Then the Nash equilibrium is also characterized by the following variational inequality:

$$D(X_i^*, u_i; \beta)'(X^* - X) \geq 0, \forall X \in \mathcal{X}(u_i).$$

Unfortunately, these are semi-infinite constraints because they must hold for all points $X \in \mathcal{X}(u_i)$. Because our constraints $\mathcal{X}^k(u_i)$ are such that they are a subset of $x \geq 0$, standard techniques from semi-infinite optimization and from robust optimization can be used to show that satisfaction of the variational inequality is equivalent to the existence of y_i^k such that

$$F^{k'} y_i^k \leq -\nabla_{x^k} \phi(X_i^*, u_i; \beta_k), \forall k = 1, \dots, p$$

$$D(X_i^*, u_i; \beta)' X_i^* - \sum_{k=1}^p (h^k - G^k u_i)' y_i^k \leq 0$$

This characterization is better because it is only finite-dimensional (as opposed to semi-infinite).

4 Tractable Formulation

To make the problem more tractable, we will impose additional conditions on our model to ensure that the feasibility problem is convex. In particular, recall that the feasibility formulation is now given by

$$\begin{aligned} \hat{\beta} &= \arg \min_{\beta} 0 \\ \text{s.t. } F^{k'} y_i^k &\leq -\nabla_{x^k} \phi(X_i^*, u_i; \beta_k), \forall k = 1, \dots, p \\ D(X_i^*, u_i; \beta)' X_i^* - \sum_{k=1}^p (h^k - G^k u_i)' y_i^k &\leq 0 \\ \beta &\in \Gamma. \end{aligned}$$

To ensure that these constraints are convex, we require that $\nabla_{x^k} \phi(X_i^*, u_i; \beta_k)$ is affine in β_k , for all k .

5 Optimality-Conditions Feasibility Formulation

Let

$$X = \begin{bmatrix} x^1 \\ \vdots \\ x^p \end{bmatrix}$$

and

$$\mathcal{X}(u) = \mathcal{X}^1(u) \times \dots \times \mathcal{X}^p,$$

and assume that we can parameterize the utility function of the k -th agent by $\phi(X, u; \beta_k)$ with a gradient that is affine in β_k for every fixed value of (X, u) . Assuming that the parameterization is differentiable, define

$$D(X, u; \beta) = \begin{bmatrix} -\nabla_{x^1} \phi(X, u; \beta_1) \\ \vdots \\ -\nabla_{x^p} \phi(X, u; \beta_p) \end{bmatrix},$$

and recall that the feasibility formulation of estimating parameters in the agents' utilities is given by:

$$\begin{aligned} \hat{\beta} &= \arg \min_{\beta} 0 \\ \text{s.t. } F^{k'} y_i^k &\leq -\nabla_{x^k} \phi(X_i^*, u_i; \beta_k), \forall k = 1, \dots, p \\ D(X_i^*, u_i; \beta)' X_i^* - \sum_{k=1}^p (h^k - G^k u_i)' y_i^k &\leq 0 \\ \beta &\in \Gamma. \end{aligned}$$

6 Examples

This might seem like a restrictive formulation (in particular the requirement that the gradient is affine in β), but it can capture many useful situations. The same examples of utilities discussed for the case of a single utility maximizing agent are equally applicable in this case. Instead, here we will focus on one particular game-theoretic model of interest to engineers.

6.1 Vehicle Traffic

We define a road network by a directed graph $G = (N, L)$, where $N = \{1, \dots, n\}$ is a set of nodes representing junctions between roads, and

$$L = \{(i, j, m) : \text{the } m\text{-th road going between node } i \text{ to node } j\}$$

is a set of links representing roads between junctions. Note that roads going in opposite directions are distinct links, meaning that $(i, j, \cdot) \neq (j, i, \cdot)$, and having a road (i, j, \cdot) does not guarantee having a road in the opposite direction (j, i, \cdot) . We do not consider roads that loop back to the original node (i, i, \cdot) .

In this model, there are w origin-destination pairs $\{o_1, d_1\}, \dots, \{o_w, d_w\}$. We associate a finite and positive rate r_i for each pair $\{o_i, d_i\}$, and this rate denotes the number of vehicles entering and exiting, when the network is at equilibrium. The amount of flow on the link (i, j, m) is given by $f_{(i,j,m)}$.

Furthermore, we define a delay function for each link $D_{(i,j,m)} : f_{(i,j,m)} \rightarrow \mathbb{R}_+$. With these definitions, it can be shown that the Nash equilibrium corresponds to a flow in which each driver tries to minimize their own travel time. Such a flow is sometimes called a User Equilibrium flow or a Wardrop equilibrium. This flow is given by the minimizer to the following optimization problem:

$$\begin{aligned} \min \quad & \sum_{(i,j,m) \in L} \int_0^{f_{(i,j,m)}} D_{(i,j,m)}(f) df \\ \text{s.t.} \quad & f_{(i,j,m)} \geq 0 \\ & r_i + \sum_{j,m} f_{(j,o_i,m)} = \sum_{j,m} f_{(o_i,j,m)}, \forall \{o_i, d_i\} \\ & \sum_{j,m} f_{(j,d_i,m)} = r_i + \sum_{j,m} f_{(d_i,j,m)}, \forall \{o_i, d_i\} \\ & \sum_{j,m} f_{(j,k,m)} = \sum_{j,m} f_{(k,j,m)}, \forall k \neq o_i, d_i, \forall i \end{aligned}$$

Given this model, we can ask the inverse question: Suppose we observe the flows on a road network, can we infer the driver belief regarding the delay function of each road? It turns out that

we can use a modified version of the formulation discussed to address this question, with the only difference being that the constraints $\mathcal{X}^k(u_i)$ do not exist; rather, there is a single $\mathcal{X}(u)$ constraint since all of the decision variables enter constraints with each other. The corresponding feasibility problem can be written, but the main question is whether solving that problem is tractable. It turns out that the terms $\nabla_{x^k}\phi$ in the feasibility formulation are equivalent to $D_{(i,j,m)}(f)$ in this vehicle traffic model.

So suppose that we have some parameterization of the vehicle traffic, say $D_{(i,j,m)}(f) = \gamma(f, \beta_{(i,j,m)})$. Then we require that γ be affine with respect to $\beta_{(i,j,m)}$ in order for our feasibility problem to be convex. For example, we may try a model such as

$$D_{(i,j,m)}(f) = 1 + k_{(i,j,m),1}f + k_{(i,j,m),2}f^2 + k_{(i,j,m),3}f^3 + k_{(i,j,m),4}f^4.$$

This parametric form is clearly affine with respect to β . It may be reasonable to try this form because a common model in the traffic engineering community is to use $D(f) = 1 + 0.15t^4$.

There is another subtle issue with the formulation. If we have different parameters for each road, then we will have more unknown parameters than data points. This is a special case of a non-identifiable model. One way to fix this model is to reduce the number of parameters. The easiest approach is to assume that the parameters are identical for each road. Another approach would be to assume that there are m classes of roads (highway, two-laned, four-laned, etc.), and each class of road has a single set of parameters.

7 Suboptimal or Noisy Points

So far, we have assumed that the points (u_i, X_i^*) are measured without noise. Suppose instead that we measure $(u_i, X_i^* + \epsilon_i)$ where ϵ_i is some i.i.d. noise. (An alternative model is that the measured points (u_i, X_i) are suboptimal, meaning that they are close to the Nash equilibrium.) This introduces a new problem because now our variational inequality will not be true. To overcome this difficulty, we define the new feasibility problem:

$$\begin{aligned} \hat{\beta} &= \arg \min_{\beta} \|\mu\|^2 \\ \text{s.t. } & F^{k'} y_i^k \leq -\nabla_{x^k}\phi(X_i^*, u_i; \beta_k), \forall k = 1, \dots, p \\ & D(X_i^*, u_i; \beta)' X_i^* - \sum_{k=1}^p (h^k - G^k u_i)' y_i^k \leq \mu_i \\ & \beta \in \Gamma. \end{aligned}$$

The idea is that we allow for residuals in the appropriate inequality constraints that would be identically zero for optimal points, to take into account that a measured point may be suboptimal.

8 Further Details

More details about these concepts can be found in the paper *Data-Driven Estimation in Equilibrium Using Inverse Optimization* by Bertsimas, Gupta, and Paschalidis, from which much of the above material is found.