# IEOR 265 – Lecture 18
# Inverse Decision-Making

## 1 Problem Framework

Suppose we have an individual (or system) making decisions by optimizing some utility (or cost) function that depends on inputs from the environment. In some problems, we have the following scenario: We get to observe the inputs and the corresponding decisions, and we would like to infer the utility function of the individual. This problem is actually a type of regression problem, but there are additional complications which mean that we have to use a more complex formulation in order to solve the problem.

The first question that needs to be addressed is the model of the individual, and the number of individuals. There are two cases we will consider.

### 1.1 Utility Maximizing Agent

Suppose that an agent makes decisions by solving the following optimization problem:

$$x_i^* = \arg\max\{J(x, u_i) \mid x \in \mathcal{X}(u_i)\},$$

where $u_i \in \mathbb{R}^q$ are inputs, $x_i^* \in \mathbb{R}^d$ are decisions, $J(x, u_i)$ is the utility function of the agent, and $\mathcal{X}(u_i)$ is a bounded feasible set (that depends on $u_i$). In this model, we observe $(u_i, x_i^*)$ for $i = 1, \ldots, n$ and would like to infer the function $J(x, u_i)$.

Note that this same model has an alternative interpretation. Suppose we have an expert (or controller) that is operating a system by minimizing a cost function, similar to MPC or LBMPC. And suppose that we observe the control actions and states of the system, but do not know the cost function. This scenario is captured by the above model by noting that the cost function is $-J(x, u_i)$, $u_i$ is the initial condition of the system (corresponding to $x_0$ in MPC), and $x_i^*$ is the control action chosen (corresponding to $u_0$ in MPC).

### 1.2 Multiple Strategic Agents

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. Perhaps the most well known notion is that of the Nash equilibrium, in which

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

However, there are other notions of rationality such as correlated equilibria. Similar to the case of the utility maximizing agent, we observe $(u_i, x_i^{*,1}, \ldots, x_i^{*,p})$ for $i = 1, \ldots n$ and would like to infer the utility function of each agent $J^k$.

## 1.3 Technical Difficulty

These problems are difficult to solve, and they are special cases of what are known as inverse optimization problems. To specifically understand the difficulty, consider the utility maximizing agent. And assume that we have a parametrization of the utility function, that is we have $\phi(x, u; \beta)$ and a bounded set $B$ such that there exists $\beta_0 \in B$ with $J(x, u) = \phi(x, u; \beta_0)$. Then, one way that the inverse decision making problem can be formulated is as

$$\begin{aligned}\hat{\beta} = \arg\min_{\beta}\ & 0 \\ \text{s.t. }\ & x_i^* = \arg\max_x \{\phi(x, u_i; \beta) \mid x \in \mathcal{X}(u_i)\} \\ & \beta \in B\end{aligned}$$

This feasibility problem is difficult to solve because it has an atypical constraint: The constraint that $x_i^*$ be the maximizer to some optimization problem cannot be directly handled by nonlinear programming techniques.

# 2 Utility Maximizing Agent

Recall the following abstract model: Suppose that an agent makes decisions by solving the following optimization problem:

$$x_i^* = \arg\max\{J(x, u_i) \mid x \in \mathcal{X}(u_i)\},$$

where $u_i \in \mathbb{R}^q$ are inputs, $x_i^* \in \mathbb{R}^d$ are decisions, $J(x, u_i)$ is the utility function of the agent, and $\mathcal{X}(u_i)$ is a bounded set (that depends on $u_i$). In this model, we observe $(u_i, x_i^*)$ for $i = 1, \ldots, n$ and would like to infer the function $J(x, u_i)$.

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

- The constraint set is described by linear equality and inequality constraints:

  $$\mathcal{X}(u) = \{x : Ax + Bu_i \leq c, Fx + Gu_i = h\},$$

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

- Assume that we have a parametrization of the utility function, that is we have $\phi(x, u; \beta)$ and a bounded set $\Gamma$ such that there exists $\beta_0 \in \Gamma$ with $J(x, u) = \phi(x, u; \beta_0)$.

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.

## 2.1 Technical Difficulty

Recall the feasibility problem formulation of the inverse decision-making problem for this single utility maximizing agent model:

$$\begin{aligned}\hat{\beta} = \arg\min_{\beta}\ & 0 \\ \text{s.t. } & x_i^* \in \arg\max_x \{\phi(x, u_i; \beta) \mid Ax + Bu_i \le c, Fx + Gu_i = h\} \\ & \beta \in \Gamma.\end{aligned}$$

This feasibility problem is difficult to solve because it has an atypical constraint: The constraint that $x_i^*$ be the minimizer to some optimization problem cannot be directly handled by nonlinear programming techniques. There are two reasons that this constraint presents challenges:

1. Depending on the value of $\beta$ there may be one or multiple maximizers. This means that in general we must treat the function

   $$P(u_i, \beta) = \arg\max_x \{\phi(x, u_i; \beta) \mid Ax + Bu_i \le c, Fx + Gu_i = h\}$$

   as a multi-valued function.

2. The function $P(u_i, \beta)$ has a complex form, because it is defined as a set of maximizers. This means that in general we cannot even hope for continuity of $P(u_i, \beta)$ (cf. the Berge Maximum Theorem, which says that for continuous $\phi$ we can only expect upper-hemicontinuity of $P(u_i, \beta)$), much less differentiablity.

## 2.2 KKT Reformulation

Since $P(u_i, \beta)$ is a multi-valued function, we can make the problem more tractable by imposing additional conditions on our model so that instead $P(u_i, \beta)$ is a single-valued (and hence continuous by the Berge Maximum Theorem) function. In particular, suppose that for all fixed values of $\beta \in \Gamma$ the function $\phi(x, u_i; \beta)$ is strictly concave in $(x, u_i)$. Then the corresponding optimization problem has a single maximizer, and so this additional condition fixes our first difficulty.

The second difficulty regarding the complex form of $P(u_i, \beta)$ still remains. However, since our constraints are linear, we have linearity constraint qualification, and so the unique maximizer $x_i^* = P(u_i, \beta)$ satisfies the KKT conditions: There exist row-vectors $\lambda_i$ and $\mu_i$ such that

$$\begin{aligned}-\nabla_x \phi(x_i^*, u_i; \beta) + \lambda_i A + \mu_i F &= 0 \\ Ax_i^* + Bu_i &\le c \\ Fx_i^* + Gu_i &= h \\ \lambda_i^j &\ge 0 \\ \lambda_i^j &= 0 \text{ if } A_j x_i^* + B_j u_i < c_j,\end{aligned}$$

where $A_j, B_j, c_j$ denote the $j$-th row of $A, B, c$ respectively. As a result, we can now pose our feasibility problem as

$$\begin{aligned}
\hat{\beta} = \arg\min_{\beta}\ & 0 \\
\text{s.t. } & -\nabla_x \phi(x_i^*, u_i; \beta) + \lambda_i A + \mu_i F = 0 \\
& \lambda_i^j \geq 0 \\
& \lambda_i^j = 0 \text{ if } A_j x_i^* + B_j u_i < c_j \\
& \beta \in \Gamma.
\end{aligned}$$

Note that because $(u_i, x_i^*)$ are measured, they are constant in our feasibility formulation and in the KKT conditions. Therefore, the conditional statement "if $A_j x_i^* + B_j u_i < c_j$" is computed before we solve the feasibility problem. In other words, we decide to either include or exclude the constraint $\lambda_i^j = 0$ in our feasibility problem based on a precomputed conditional.

This problem can still be difficult to solve, because this reformulated problem may not be convex. Consider the constraint

$$-\nabla_x \phi(x_i^*, u_i; \beta) + \lambda_i A + \mu_i F = 0,$$

and note that it is an equality constraint. However, a standard result is that an equality constraint $Q(\beta)$ is convex if and only $Q$ is an affine function (meaning that it can be written as $Q = M\beta + k$ where $M$ is a matrix and $k$ is a constant vector). As a result, our feasibility problem to estimate the parameters $\beta$ of our utility function is convex if and only if $Q(\beta) = -\nabla_x \phi(x_i^*, u_i; \beta)$ is an affine function. Stated in another way, our formulation is convex if and only if the gradient of $\phi$ with respect to $x$ is affine in $\beta$ when the gradient is evaluated at $x_i^*$ and $u_i$.

# 3 Optimality-Conditions Feasibility Formulation

Recall the feasibility formulation when we can parameterize the utility function by $\phi(x, u; \beta)$, where this function is strictly concave in $(x, u)$ for every fixed value of $\beta \in \Gamma$, with a gradient that is affine in $\beta$ for every fixed value of $(x, u)$:

$$\begin{aligned}
\hat{\beta} = \arg\min_{\beta}\ & 0 \\
\text{s.t. } & -\nabla_x \phi(x_i^*, u_i; \beta) + \lambda_i A + \mu_i F = 0 \\
& \lambda_i^j \geq 0 \\
& \lambda_i^j = 0 \text{ if } A_j x_i^* + B_j u_i < c_j \\
& \beta \in \Gamma.
\end{aligned}$$

Note that we will assume that $x \in \mathbb{R}^d$ and $u \in \mathbb{R}^q$.

# 4 Examples

This might seem like a restrictive formulation (in particular the requirement that the gradient is affine in $\beta$), but it can capture many useful situations. A few examples are described here.

## 4.1 Quadratic Utilities

Consider a quadratic utility given by

$$\phi(x, u; \beta) = -x'Qx + u'F'x + k'x,$$

where $Q \in \mathbb{R}^{d \times d} : Q \succeq 0$, $F \in \mathbb{R}^{d \times q}$ are matrices, and $k \in \mathbb{R}^d$ is a vector. Note that its gradient

$$\nabla_x \phi(x, u; \beta) = -2Qx + Fu + k.$$

is an affine function of the parameters $Q, F, k$.

The first thing to note is that this utility is equivalent to the following:

$$\tilde{\phi}(x, u; \tilde{\beta}) = -\begin{bmatrix} x \\ u \end{bmatrix}' \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{12}' & Q_{22} \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix} + \begin{bmatrix} k_1 \\ k_2 \end{bmatrix}' \begin{bmatrix} x \\ u \end{bmatrix},$$

where $\begin{bmatrix} Q_{11} & Q_{12} \\ Q_{12}' & Q_{22} \end{bmatrix} \succeq 0$ is a block matrix that is appropriately sized, and $\begin{bmatrix} k_1 \\ k_2 \end{bmatrix}$ is an appropriately sized block vector. The equivalence of this utility can be seen by noting that

$$\tilde{\phi}(x, u; \tilde{\beta}) = -\begin{bmatrix} x \\ u \end{bmatrix}' \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{12}' & Q_{22} \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix} + \begin{bmatrix} k_1 \\ k_2 \end{bmatrix}' \begin{bmatrix} x \\ u \end{bmatrix} = -x'Q_{11}x - u'Q_{22}u + 2u'Q_{12}'x + k_1'x + k_2'u.$$

Since the utility maximizing agent optimizes over $x$ for a fixed value of $u$, this means that the minimizer of this second problem will be equivalent to the first if $Q = Q_{11}$, $F = 2Q_{12}'$, and $k = k_1$.

The second thing to note is that there is a problem with the corresponding feasibility formulation

$$\begin{aligned} \hat{\beta} = \arg\inf_{\beta}\ & 0 \\ \text{s.t. } & 2Qx_i^* - Fu_i - k + \lambda_i A + \mu_i F = 0 \\ & \lambda_i^j \geq 0 \\ & \lambda_i^j = 0 \text{ if } A_j x_i^* + B_j u_i < c_j \\ & Q \succ 0. \end{aligned}$$

The following $\beta$ is a feasible point of the above problem: $Q = 0$, $F = 0$, $k = 0$,$\lambda_i = 0$,and $\mu_i = 0$. This problem is a manifestation of the fact that there are an infinite number of utility functions that can lead to an observed set of decisions. To fix this problem, we must ensure that the formulation is properly normalized. One approach is to change the formulation to

$$\begin{aligned}\hat{\beta} = \arg\min_{\beta}\ & 0 \\ \text{s.t. }\ & 2Qx_i^* - Fu_i - k + \lambda_i A + \mu_i F = 0 \\ & \lambda_i^j \geq 0 \\ & \lambda_i^j = 0 \text{ if } A_j x_i^* + B_j u_i < c_j \\ & Q \succeq \mathbb{I}.\end{aligned}$$

## 4.2 Nonparametric Utilities

Instead of a parametric form of the utility, we can also define a nonparametric utility (essentially meaning an infinite number of parameters). For instance, we could have

$$\phi(x, u; \beta) = \sum_{i=0}^{\infty} k_i f_i(x, u),$$

where $f_i(x, u) : \mathbb{R}^d \times \mathbb{R}^q \to \mathbb{R}$ is a differentiable nonlinear function, and the $\beta$ are the $k_i$ parameters. In this case, the gradient is given by

$$\phi(x, u; \beta) = \sum_{i=0}^{\infty} k_i \nabla_x f_i(x, u),$$

which is affine in the $\beta$. Note that in general we will face a normalization issue, and so we would have to include an appropriate constraint in our feasibility problem to deal with this.
An example of the above is a finite polynomial expansion:

$$\phi(x, u; \beta) = k_1 x + k_2 x^2 + k_3 xu + k_4 x^2 u + k_5 x u^2,$$

where we the inputs are such that $u > 0$. In this case, the feasibility problem with (one-potential) normalization is given by

$$\begin{aligned}\hat{\beta} = \arg\min_{\beta}\ & 0 \\ \text{s.t. }\ & -k_1 - 2k_2 x_i^* - k_3 u_i - 2k_4 x_i^* u_i - k_5 u_i^* 2 + \lambda_i A + \mu_i F = 0 \\ & \lambda_i^j \geq 0 \\ & \lambda_i^j = 0 \text{ if } A_j x_i^* + B_j u_i < c_j \\ & k_2 \geq 0, k_4 \geq 0 \\ & k_2 \geq 1.\end{aligned}$$

Here, we have chosen the normalization $k_2 \geq 1$. Note that we could have chosen other normalization constraints, such as $k_1 \geq 1$.

# 5 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 $\epsilon_i$ 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 optimal values.) This introduces a new problem because now our optimality conditions will not be true. To overcome this difficulty, we define the new feasibility problem:

$$
\begin{aligned}
\hat{\beta} = \arg\min_{\beta} \ & \sum_i \|r_{i,s}\|_2^2 + \|r_{i,c}\|_2^2 \\
\text{s.t.} \ & -\nabla_x \phi(x_i^*, u_i; \beta) + \lambda_i A + \mu_i F = r_{i,s} \\
& \lambda_i^j \geq 0 \\
& \lambda_i^j = r_{i,c}^j \text{ if } A_j x_i^* + B_j u_i < c_j \\
& \beta \in \Gamma.
\end{aligned}
$$

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