Introduction to Game Theory: Simple Decisions Models

John C.S. Lui

Department of Computer Science & Engineering
The Chinese University of Hong Kong

Outline

1. Optimization
2. Making Decisions
3. Modeling Rational Behavior

Example

Example 1

- Consider the function $f(x) = \sqrt{x}$ where $x \in [0, \infty)$. What is the maximum value of $f(x)$? What if $x \in [0, 4)$?
- Maximize $f(n) = 1 - n^2 + \frac{9}{2}n$ where n must take on integer value.
- Let a, b and c be positive constant. Let

$$f(x) = \begin{cases} \left(1 - \frac{x}{b}\right) & \text{if } 0 \leq x \leq b, \\ 0 & \text{otherwise} \end{cases}$$

  and

$$\mathcal{X}(x) = \begin{cases} 1 & \text{if } x > 0, \\ 0 & \text{if } x \leq 0. \end{cases}$$

  Maximize $g(x) = axf(x) - c\mathcal{X}(x)$.

Brief introduction to C. Optimization

Using *explicit function*

- Maximize $f(x, y) = x - y^2$ subject to $x - y = 0$.
- Rewrite $y = x$, then $f(x, x) = x - x^2$.
- Differentiating, $f^{'}(x, x) = 1 - 2x = 0$, we have $(x^*, y^*) = (1/2, 1/2)$ with $f^*(x^*, y^*) = 1/4$.

Lagrange Multiplier

Introduction

- Maximize $f(x,y)$ subject to $g(x,y) = 0$.
- Set up Lagrangian:

$$L(x,y) = f(x,y) - \lambda g(x,y)$$

- Perform *unconstrained maximization* on $L(x,y)$
- Another way is to view it is that it is an unconstrained optimization on $f(x,y)$ with an additional penalty for violating the constraint $g(x,y) = 0$.

Lagrangian Multiplier

Example

- $f(x, y) = x - y^2$ with $x - y = 0$.
- $L(x, y) = x - y^2 - \lambda(x - y)$.
- Differentiating, we have: $1 - \lambda = 0$; $-2y^* + \lambda = 0$; $x^* - y^* = 0$.
- Solving, $x^* = y^* = 1/2$.

LM with inequality constraints

Inequality constraints

Let say $g(x,y) \leq 0$. Define $L(x,y) = f(x,y) - \lambda g(x,y)$ but now require $\lambda > 0$. If $L(x^*, y^*)$ is an unconstrained maximum and $\lambda > 0$, then the maximum of $f(x,y)$ subject to $g(x,y) \leq 0$ occurs at (x^*, y^*).

Lagrangian Multiplier

Example

- $f(x, y) = 2x^2 + y^2$ subject to $x^2 + y^2 - 1 \leq 0$.
- Define $L(x, y) = 2x^2 + y^2 - \lambda(x^2 + y^2 - 1)$.
- Differentiating, we have $x(2 - \lambda) = 0$; $y(1 - \lambda) = 0$; $x^2 + y^2 = 1$.
- Solutions (a) $\lambda = 1, x = 0, y = \pm 1$; (b) $\lambda = 2, x = \pm 1, y = 0$;
- Both points: $(0, \pm 1)$ and $(\pm 1, 0)$ are extrema. Because $f(0, \pm 1) = 1$ while $f(\pm 1, 0) = 2$.
- So $(x^*, y^*) = (\pm 1, 0)$.

Some definitions

Definition

argmax is defined by the following equivalence:

$$x^* \in argmax_{x \in X} f(x) \iff f(x^*) = \max_{x \in X} f(x)$$

Note that we do not write $x^* = argmax_{x \in X} f(x)$ since *argmax* returns a set of values.

Definition

A choice of behavior in a single-decision problem is called an *action*. The set of alternative actions available will be denoted as $\boldsymbol{A}$. This will either be discrete set, e.g., $\{a_1, a_2, \ldots, \}$, or a continuous set, .e.g., the unit interval $[0, 1]$.

More...

Definition

A *payoff* is a function $\pi : \boldsymbol{A} \rightarrow R$ that associates a numerical value with every action $a \in \boldsymbol{A}$.

Definition

An action a^* is an *optimal action* if

$$\pi(a^*) \geq \pi(a) \qquad \forall a \in \boldsymbol{A}.$$

or equivalently, $a^* \in argmax_{a \in \boldsymbol{A}} \pi(a)$.

Definition

An *affine transformation* changes payoff $\pi(a)$ into $\pi^{'}(a)$ as

$$\pi^{'}(a) = \alpha\pi(a) + \beta$$

where α, β are constants independent of a and $\alpha > 0$.

Theorem

The optimal action is unchanged if payoffs are altered by an affine transformation.

Proof

because $\alpha > 0$, we have

$$\begin{aligned} argmax_{a \in \boldsymbol{A}} \pi^{'}(a) &= argmax_{a \in \boldsymbol{A}} [\alpha\pi(a) + \beta] \\ &= argmax_{a \in \boldsymbol{A}} \pi(a). \end{aligned}$$

Example

The Convent Fields Soup Company needs to determine the price p. The demand function is:

$$Q(p) = \begin{cases} Q_0\left(1 - \frac{p}{p_0}\right) & \text{if } p < p_0, \\ 0 & \text{if } p \geq p_0. \end{cases}$$

The payoff is $\pi(p) = (p - c)Q(p)$ where c is the unit production cost.

- Solving, we have $p^* = \frac{1}{2}(p_0 + c)$.
- Now, let say we need to consider a fixed cost to build the factory, the payoff function is $\pi(p) = (p - c)Q(p) - B$, where B is a constant. What is p^*?

Uncertainty

Modeling uncertainty

- If uncertainty exists, we compare the expected outcome for each action.
- Let X be the set of states with $P(X = x)$.
- Payoff for adopting action a is:

$$\pi(a) = \sum_{x \in X} \pi(a|x) P(X = x)$$

- An optimal action is

$$a^* \in argmin_{a \in A} \sum_{xinX} \pi(a|x) P(X = x).$$

Example

- An investor has \$1000 to invest in one year. The available actions (1) put the money in the bank with 7% interest per year; (2) invest in stock which returns \$1500 if the stock market is good or returns \$600 if the stock market is bad.
- $P(Good) = P(Bad) = 0.5$.
- Expected payoff:
  1. $\pi(a_1) = \$1070$;
  2. $\pi(a_2) = 1500/2 + 600/2 = \$1050$.
- So a_1^* and we should put the money in the bank.

Definition

Let $\Omega = \{\omega_1, \omega_2, \ldots\}$ be the set of possible outcomes.

- We say $\omega_i \succ \omega_j$ if an individual *strictly prefers* outcome ω_i over ω_j.
- If the individual is indifferent: $\omega_i \sim \omega_j$.
- Either prefer or indifferent: $\omega_i \succeq \omega_j$.

Definition

An individual will be called rational under certainty if his preference for outcomes satisfy the following conditions:

- (Completeness) Either $\omega_i \succeq \omega_j$ or $\omega_j \succeq \omega_i$.
- (Transitivity) If $\omega_i \succeq \omega_j$ and $\omega_j \succeq \omega_k$, then $\omega_i \succeq \omega_k$.

Definition

A utility function is a function $u : \Omega \to \mathbf{R}$ such that

$$\begin{aligned} u(\omega_i) > u(\omega_j) &\iff \omega_i \succ \omega_j \\ u(\omega_i) = u(\omega_j) &\iff \omega_i \sim \omega_j \end{aligned}$$

The immediate consequence of this definition is an individual who is rational under certainty should seek to maximize his utility.

What happens when an action does not produce a definite outcome and instead, we allow each outcome to occur with a known probability?

Definition

A simple lottery, λ, is a set of probabilities for the occurrence of every $\omega \in \Omega$. The probability that outcome ω occurs is $p(\omega|\lambda)$. The set of all possible lotteries is denoted as $\boldsymbol{A}$.

Definition

A compound lottery is a linear combination of simple lotteries (from the same set $\boldsymbol{A}$). For example, $q\lambda_1 + (1-q)\lambda_2$ with $0 \leq q \leq 1$ is a compound lottery.

Theorem

Expected Utility Theorem: *If an individual is rational, then we can define a utility function $u : \Omega \to \mathbf{R}$ and the individual will maximize the payoff function $\pi(a)$ (or the* expected utility*) given by*

$$\pi(a) = \sum_{\omega \in \Omega} p(w|\lambda(a))u(\omega)$$

Definition

An individual whose utility function satisfies

- $E(u(w)) < u(E(w))$, it is said to be risk averse,
- $E(u(w)) > u(E(w))$, it is said to be risk prone,
- $E(u(w)) = u(E(w))$, it is said to be risk neutral.

Example

Someone flip a coin. If it is head (tail), you get $1 ($1M). Your utility function can be:

- $u(x) = x$,
- $u(x) = x^2$,
- $u(x) = e^{-x}$.

Classify the above as risk averse, risk prone and risk neutral utility function.

Homework

Consider an individual whose utility function of wealth, w, is given by $u(w) = 1 - e^{-kw}$ where $k > 0$. Assume that wealth is distributed by a normal distribution $N(\mu, \sigma^2)$. Show

- Individual's expected utility can be represented as a trade-off between μ and σ^2.
- Classify the utility function.

Up to now, we assume finding an optimal action a^* from a given set $\boldsymbol{A}$. But the selection can be *randomized*. Does this allow one to achieve a higher payoff?

Definition

We specify a general behavior β by giving a list of probabilities with which each available action is chosen. We denote the probability that action a is chosen by $p(a)$ and $\sum_{a \in \boldsymbol{A}} p(a) = 1$. The set of all randomizing behavior is denoted by $\boldsymbol{B}$. The payoff of using behavior β is

$$\pi(\beta) = \sum_{a \in \boldsymbol{A}} p(a)\pi(a).$$

An optimal behavior β^* is one for which

$$\pi(\beta^*) \geq \pi(\beta) \quad \forall \beta \in \boldsymbol{B}.$$

or $\beta^* \in argmax_{\beta \in \boldsymbol{B}} \pi(\beta)$.

Definition

The support of a behavior β is the set $\boldsymbol{A}(\beta) \subseteq \boldsymbol{A}$ of all the actions for which β specifies $p(a) > 0$.

Theorem

Let β^* be an optimal behavior with support $\boldsymbol{A}^*$. Then

$$\pi(a) = \pi(\beta^*) \quad \forall a \in \boldsymbol{A}^*.$$

The consequence of the above theorem is that if a randomized behavior is optimal, then two or more actions are optimal as well. So randomization is not necessary but it may be used to break a tie.

Example

A firm may make one of the marketing actions $\{a_1, a_2, a_3\}$. The profit for each action depends on the state of the economy $\boldsymbol{X} = \{x_1, x_2, x_3\}$:

	x_1	x_2	x_3
a_1	6	5	3
a_2	3	5	4
a_3	5	9	1

If $P(X = x_1) = 1/2$, $P(X = x_2) = P(X = x_3) = 1/4$. What *are* the optimal behaviors?

Answer

Because $\pi(a_1) = \pi(a_3) = 5$ and $\pi(a_2) = 3.75$, optimal randomizing behaviors have support $\boldsymbol{A}^* = \{a_1, a_3\}$ with $p(a_1) = p$ and $p(a_3) = 1 - p$ $(0 < p < 1)$. Using either a_1 or a_3 with probability 1 is also an optimal behavior.