Advanced Microeconomic Analysis, Lecture 1

Prof. Ronaldo CARPIO

March 6, 2017

Welcome to Advanced Microeconomic Analysis

- This course is an introduction to the foundations of microeconomic theory, that is, the analysis of the behavior of individual rational agents (consumers, firms, etc).
- The course will be taught entirely in English.
- Website: `http://rncarpio.com/teaching/AdvMicro`
- Announcements, slides, & homeworks will be posted on website

About Me: Ronaldo Carpio

- BS Electrical Engineering & CS, UC Berkeley
- Master's in Public Policy, UC Berkeley
- PhD Economics, UC Davis
- Joined School of Banking & Finance in 2012

Textbook

- Main textbook: *Advanced Microeconomic Theory, 3rd Ed* (2003) by Jehle & Reny
- We will tentatively cover Chapters 1-5 and 7-8. Material may be added or dropped, depending on time constraints.
- For those interested, a more mathematically rigorous textbook is *Microeconomic Theory* by Mas-Colell, Whinston, and Green.

Grading

- There will be around 5 graded homework assignments, due every 2 weeks. Assignments will be posted on the course website.
- Homeworks: 15%
- Midterm exam: 35%
- Final exam: 50%

Contacting Me

- Email: rncarpio@yahoo.com
- Office: 913 Boxue Bldg.
- Office Hours: Monday 16:00-17:00 or by appointment

Course Outline

- Preliminaries: Convexity and Constrained Optimization (Appendix A2.2 & A2.3)
- Consumer Theory (Chapter 1)
- Duality (Chapter 2.1)
- Risk and Uncertainty (Chapter 2.4)
- Theory of the Firm (Chapter 3)
- Partial Equilibrium (Chapter 4.1)
- General Equilibrium (Chapter 5.1, 5.2, 5.4)
- Game Theory (Chapter 7)
- Imperfect Competition (Chapter 4.2)
- Asymmetric Information (Chapter 8)

Convexity

- Let's begin by considering sets of points in $\mathbb{R}^n$.
- A set $S \subset \mathbb{R}^n$ is *convex* if: for any two points $\boldsymbol{x}^1, \boldsymbol{x}^2 \in S$, all points on the line segment joining $\boldsymbol{x}^1$ and $\boldsymbol{x}^2$ are also in S:

$$t\boldsymbol{x}^1 + (1-t)\boldsymbol{x}^2 \in S \qquad \text{for all } t \in [0,1]$$

- The weighted average $t\boldsymbol{x}^1 + (1-t)\boldsymbol{x}^2$, where the weights are positive and add up to 1, is called a *convex combination*.
- We can also have convex combinations of any number of points: $t_1\boldsymbol{x}^1 + t_2\boldsymbol{x}^2 + ... + t_I\boldsymbol{x}^I$, where $\sum_i^I t_i = 1$ and $t_i \geq 0$ for all i.
- An *extreme point* of S is a point that cannot be written as a convex combination of other points in S.

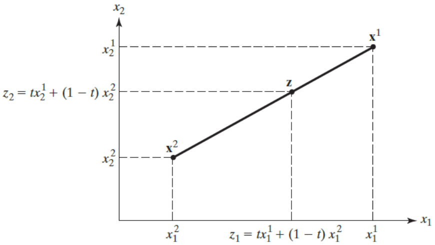

Figure A1. 4. Some convex combinations in $\mathbb{R}^2$.

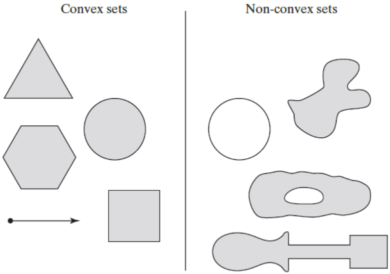

Figure A1. 5. Convex and non-convex sets in $\mathbb{R}^2$.

Properties of Convex Sets

- If S is convex, then the convex combination of any finite number of points in S is also in S.
- The intersection of any number of convex sets is convex.

Hyperplanes

- A *hyperplane* is a generalization of a line in 2D to N dimensions:

$$\{\boldsymbol{x} \in \mathbb{R}^n | \boldsymbol{a} \cdot \boldsymbol{x} = b\}$$

- Here, $\boldsymbol{a}$ is a *normal vector* to the hyperplane: it is perpendicular to any vector that lies in the hyperplane.
- For example, in 2D, the equation $ax + by = c$ defines a hyperplane, and (a, b) is a normal vector to this hyperplane.
- An equivalent equation defining a hyperplane with normal vector $\boldsymbol{a}$, passing through a point $\boldsymbol{x}^{\boldsymbol{0}}$ is: $\boldsymbol{a} \cdot (\boldsymbol{x} - \boldsymbol{x}^0) = 0$.
- A hyperplane defines two *half-spaces*:
  - the half-space *above* the hyperplane, $\{\boldsymbol{x} | \boldsymbol{a} \cdot \boldsymbol{x} \geq b\}$
  - the half-space *below* the hyperplane, $\{\boldsymbol{x} | \boldsymbol{a} \cdot \boldsymbol{x} \leq b\}$
- In economics, the most common hyperplane is the *budget line*: if there are n goods with prices $p_1 ... p_n$ and wealth w, then the equation $p_1 x_1 + ... p_n x_n = w$ defines a hyperplane.

Half-Spaces

- A half-space is a convex set.
- Therefore, the (non-empty) intersection of any number of half-spaces is also a convex set.
- In the standard consumer problem with two goods, the feasible set (that is, the set of possible bundles we want to maximizer over) is the intersection of three half-spaces:

$$x_1 \geq 0, \qquad x_2 \geq 0, \qquad p_1 x_1 + p_2 x_2 \leq w$$

- A polygon with n faces is the intersection of n half-spaces.

Separating Hyperplane

- Suppose X and Y are closed, nonempty, and disjoint convex subsets of $\mathbb{R}^n$. If either X or Y are compact (i.e. closed and bounded), then there exists a hyperplane that *separates* X and Y.
- That is, there exists a nonzero vector $\boldsymbol{a} \in \mathbb{R}^n$ and scalar b such that:
  - $\boldsymbol{a} \cdot \boldsymbol{x} > b$ for all $x \in X$
  - $\boldsymbol{a} \cdot \boldsymbol{y} < b$ for all $y \in Y$
- This theorem is be used to prove the existence of general equilibrium and existence of solutions in game theory.

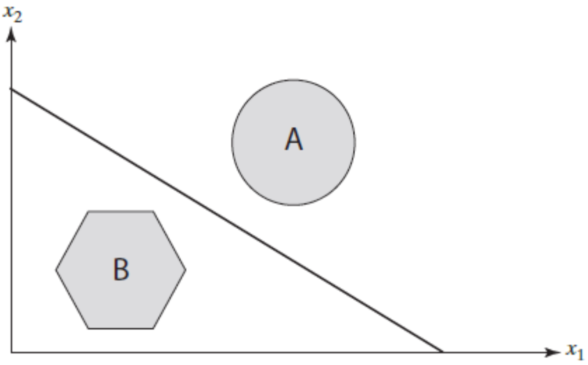

Figure A2.14. Separating convex sets.

Convex and Concave Functions

- Consider a function $f : X \to \mathbb{R}$, where X is a convex set.
- Let $\boldsymbol{x}^1, \boldsymbol{x}^2$ be any two points in X, and let $\boldsymbol{x}^t$ denote their convex combination: $\boldsymbol{x}^t = tx^1 + (1-t)\boldsymbol{x}^2$, for $t \in [0,1]$.
- f is *concave* if $f(\boldsymbol{x}^t) \geq tf(\boldsymbol{x}^1) + (1-t)f(\boldsymbol{x}^2)$, for all $t \in [0,1]$.
- f is *strictly concave* if the inequality is strict.
- f is *convex* if $f(\boldsymbol{x}^t) \leq tf(\boldsymbol{x}^1) + (1-t)f(\boldsymbol{x}^2)$, for all $t \in [0,1]$.
- f is *strictly convex* if the inequality is strict.

Epigraphs and Hypographs of a Function

- The *epigraph* of a function f is the set of all points that are *above* the graph:
$$\{(\mathbf{x}, r) \in \mathbb{R}^{n+1} | r \geq f(\mathbf{x})\}$$
- Similarly, the *hypograph* of a function f is the set of all points that are *below* the graph:
$$\{(\mathbf{x}, r) \in \mathbb{R}^{n+1} | r \leq f(\mathbf{x})\}$$

Epigraph

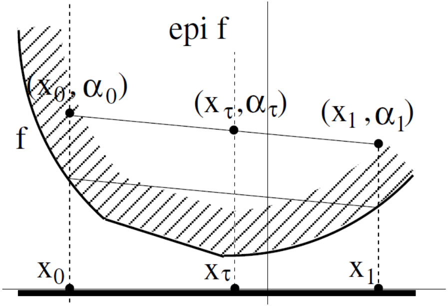

Convex and Concave Functions

- A function f is concave iff its hypograph (the set of points *below* the graph) is convex.
- A function f is convex iff its epigraph (the set of points *above* the graph) is convex.

Calculus Conditions for Convexity and Concavity

- A twice-differentiable function f is concave iff its Hessian matrix (i.e. the matrix of second derivatives) is negative semidefinite at all points of its domain.
- f is strictly concave if its Hessian is negative definite at all points on its domain.
- f is (strictly) convex if its Hessian is (positive definite) positive semidefinite at all points on its domain.

Quasiconvex and Quasiconcave Functions

- The upper level set of a function f at a value r is the set $\{\boldsymbol{x} \in X | f(x) \geq r\}$.
- Similarly, the lower level set is $\{\boldsymbol{x} \in X | f(x) \geq r\}$.
- A function f is *quasiconcave* if its upper level sets are convex for every $r \in \mathbb{R}$.
- Similarly, f is *quasiconvex* if its lower level sets are convex for every $r \in \mathbb{R}$.
- Suppose that $f : X \to \mathbb{R}$ is quasiconcave and X is convex. The the set of maximizers of f over X is convex.

A Quasiconcave Function

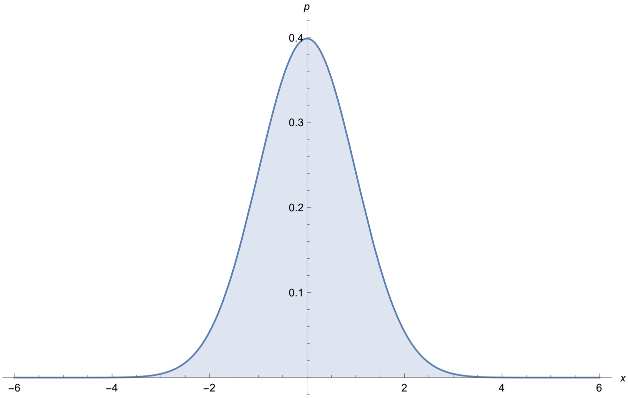

- For example, a "bell curve" or normal distribution is not concave, but it is quasiconcave.
- For any value of p, the upper level set (that is, the set of all x such that $f(x) \geq p$) is a convex set.

Unconstrained Optimization

- Suppose $f : X \to \mathbb{R}$ is twice continuously differentiable.
- Necessary conditions for $f(\boldsymbol{x})$ to be a local maximum are:
  - First-order: $\nabla f(\boldsymbol{x}) = 0$
  - Second-order: The Hessian of f at $\boldsymbol{x}$ is negative semidefinite.
- If f is concave, then all local maxima are also global.
- This is the reason why concavity is so desirable for optimization problems: to find the *global* maximum, you only need to check *local* information about the function.
- Otherwise, to guarantee a *global* maximum, you might have to compare *all* local maxima.
- If f is strictly concave, then there is a unique maximum.

Constrained Optimization

- Suppose we want to solve a maximization problem with a single *equality* constraint:

$$\max_{x_1,x_2} f(x_1, x_2) \qquad \text{subject to } g(x_1, x_2) = 0$$

- For example, the standard consumer problem with 2 goods is

$$\max_{x_1,x_2} u(x_1, x_2) \qquad \text{subject to } p_1x_1 + p_2x_2 - w = 0$$

- Lagrange's method is to form a new function, the Lagrangian, with an additional variable λ:

$$L(x_1, x_2, \lambda) = f(x_1, x_2) - \lambda g(x_1, x_2)$$

- Then, we find an inflection point of this function, by setting its gradient to 0.

$$\nabla L = \begin{pmatrix} \frac{\partial L}{\partial x_1} \\ \frac{\partial L}{\partial x_2} \\ \frac{\partial L}{\partial \lambda} \end{pmatrix} = \begin{pmatrix} \frac{\partial f(x_1,x_2)}{\partial x_1} - \lambda \frac{\partial g(x_1,x_2)}{\partial x_1} \\ \frac{\partial f(x_1,x_2)}{\partial x_2} - \lambda \frac{\partial g(x_1,x_2)}{\partial x_2} \\ g(x_1, x_2) \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

Example: Find the Closest Point on a Line

- Suppose we are given a point (x_0, y_0) and a line $a_1 x + a_2 y = b$.
- The (squared) Euclidean distance between two points (x_1, y_1) and (x_2, y_2) is $(x_2 - x_1)^2 + (y_2 - y_1)^2$.
- We want to find the point on the line that is closest to (x_0, y_0).
- We can write this as a constrained maximization problem:

$$\max_{x,y} -(x - x_0)^2 - (y - y_0)^2 \qquad \text{s.t. } a_1 x + a_2 y - b = 0$$

- The (concave) objective function is $f(x, y) = -(x - x_0)^2 - (y - y_0)^2$.
- There is one equality constraint: $g(x, y) = a_1 x + a_2 y - b = 0$
- The Lagrangian is:

$$L(x, y, \lambda) = -(x - x_0)^2 - (y - y_0)^2 - \lambda(a_1 x + a_2 y - b)$$

$$L(x,y,\lambda) = -(x-x_0)^2 - (y-y_0)^2 - \lambda(a_1 x + a_2 y - b)$$

$$\frac{\partial L}{\partial x} = -2(x-x_0) - \lambda a_1 = 0$$

$$\frac{\partial L}{\partial y} = -2(y-y_0) - \lambda a_2 = 0$$

$$\frac{\partial L}{\partial \lambda} = a_1 x + a_2 y - b = 0$$

Divide eqn 1 by eqn 2 to eliminate λ:

$$\frac{x-x_0}{y-y_0} = \frac{a_1}{a_2} \Rightarrow a_2(x-x_0) = a_1(y-y_0)$$

Now we have 2 unknowns and 2 equations. Solve to get:

$$x = \frac{a_1 b + a_2^2 x_0 + a_1 a_2 y_0}{a_1^2 + a_2^2}, y = \frac{a_2 b - a_1 a_2 x_0 + a_1^2 y_0}{a_1^2 + a_2^2}$$

Lagrange's Theorem (A2.16)

Lagrange's Theorem for multiple equality constraints:

- Let $f(\boldsymbol{x}), g^1(\boldsymbol{x}), ..., g^m(\boldsymbol{x})$ be continuously differentiable functions over some domain D.
- Let $\boldsymbol{x}^*$ be an interior point of D, and suppose that $\boldsymbol{x}^*$ is an optimum of $f(\cdot)$ subject to the constraints $g^1(\boldsymbol{x}^*) = 0, ..., g^m(\boldsymbol{x}^*) = 0$.
- If the gradient vectors $\nabla g^1(\boldsymbol{x}^*), ..., \nabla g^m(\boldsymbol{x}^*)$ are linearly independent, then there exist m unique numbers $\lambda_1^*, ..., \lambda_m^*$ such that:

$$\frac{\partial L(\boldsymbol{x}^*, \boldsymbol{\lambda}^*)}{\partial x_i} = \frac{\partial f(\boldsymbol{x}^*)}{\partial x_i} - \sum_{j=1}^{m} \lambda_j^* \frac{\partial g(\boldsymbol{x}^*)}{\partial x_i} \qquad \text{for } i = 1, ..., n$$

- Alternatively:

$$\nabla f(\boldsymbol{x}^*) = \sum_{j=1}^{m} \lambda_j^* \nabla g^j(\boldsymbol{x}^*), \qquad g^j(\boldsymbol{x}^*) = 0 \;\text{ for } j = 1, ..., m$$

Lagrange's Theorem (A2.16)

- Note that these conditions are *necessary*, but not *sufficient*.
- If an optimum exists, then $\lambda_1^*, ..., \lambda_m^*$ exist, but not vice versa.
- This does not guarantee that the optimum exists (for example, if there is no point that satisfies all the equality constraints simultaneously).
- Only first-order conditions are used, which means that if there is an optimum, it may be a minimum or a maximum.
- Second-order conditions (e.g. the Hessian) are necessary to prove whether it is a minimum or a maximum.
- In this course, we will only deal with concave or quasiconcave functions, so it will be a maximum.

Why does the Lagrange method work?

- Recall that the gradient of a function is the direction of *steepest increase*.
- Suppose we have a smooth function f, with gradient $\nabla f(\boldsymbol{x})$.
- At point $\boldsymbol{x}$, if we move at a direction not perpendicular to $\nabla f(\boldsymbol{x})$, the value of f changes.
- If we move in a direction perpendicular to $\nabla f(\boldsymbol{x})$, the value of f is constant.
- Suppose $\boldsymbol{x}$ satisfies the constraint $g(\boldsymbol{x}) = 0$. If we are at a constrained maximum, it is not possible to move along the constraint while increasing f.
- This is only possible if $\nabla f(\boldsymbol{x})$ is perpendicular to g, i.e. if $\nabla f(\boldsymbol{x})$ is parallel to $\nabla g(\boldsymbol{x})$.
- This is an alternative way of writing the Lagrangian:

$$\nabla f(\boldsymbol{x}) = \lambda \nabla g(\boldsymbol{x}), \qquad g(\boldsymbol{x}) = 0$$

Inequality Constraints

- Suppose we want to have *inequality* constraints as well as *equality* constraints.
- A simple example: suppose we want to maximize $f(x)$ subject to $x \geq 0$.
- Assume $f(\cdot)$ is concave.
- There are three possible ways in which the constraint is satisfied at the optimum.

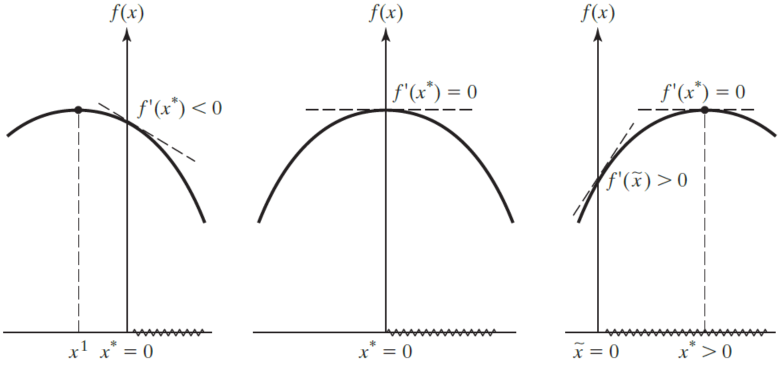

(a) Case 1 (b) Case 2 (c) Case 3

- Case 1: the global maximum is at $x^1 < 0$, which violates the constraint $x \geq 0$.
- Therefore, the constrained maximum must occur exactly at $x^* = 0$, since $f(\cdot)$ is decreasing at this point.
- We say that the constraint is "binding" or "active" at the optimum.
- If we removed this constraint, the optimizer would change.

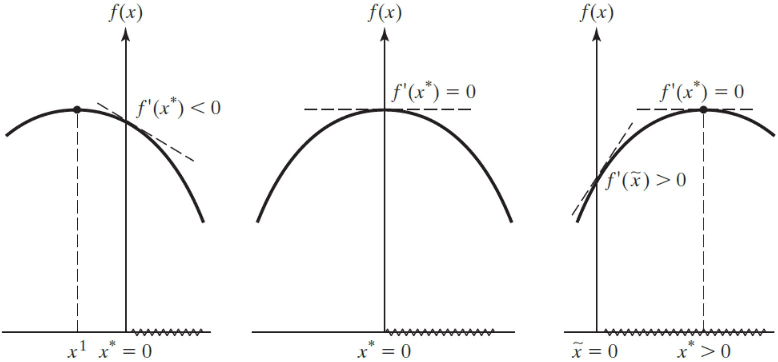

(a) Case 1 (b) Case 2 (c) Case 3

- Case 2: the global maximum is at $x^* = 0$, which satisfies the constraint $x \geq 0$ with equality.
- We say that the constraint is "binding" or "active", but *irrelevant*.
- If we removed this constraint, the optimizer would not change.
- The unconstrained FOC conditions are satisfied: $f'(x) = 0$.

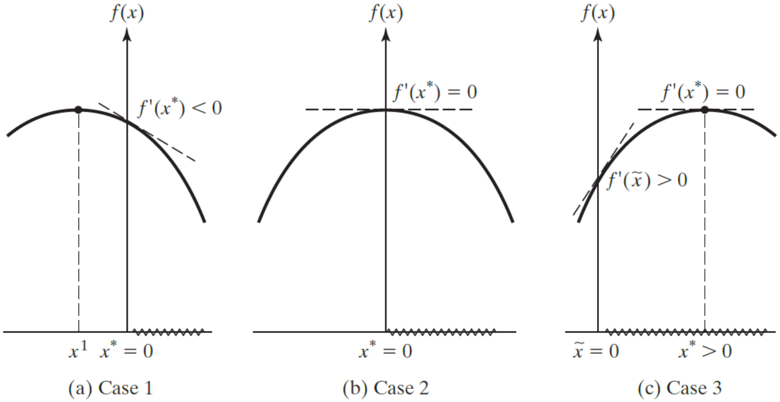

- Case 3: the global maximum is at $x^* > 0$, which satisfies the constraint $x \geq 0$ with *strict inequality*.
- We say that the constraint is "not binding" or "not active".
- If we removed this constraint, the optimizer would not change.
- The unconstrained FOC conditions are satisfied: $f'(x) = 0$.

- We can apply the method of Lagrange multipliers to inequality constraints, with some additional conditions.
- As before, each constraint has its own Lagrange multiplier λ_i.
- If a constraint is not binding, or binding and irrelevant, then $\lambda_i = 0$.
- If a constraint is binding and not irrelevant, then $\lambda_i > 0$.
- This leads to a more general version of Lagrange's theorem, called the Karush-Kuhn-Tucker or KKT conditions.

- **KKT Necessary Conditions for Optmality (A2.20)**
- Consider the maximization problem with m inequality constraints:

$$\max_{\boldsymbol{x}} f(\boldsymbol{x}) \qquad \text{s.t. } g^1(\boldsymbol{x}) \leq 0, ..., g^m(\boldsymbol{x}) \leq 0$$

- Assume $f(\cdot), g^1(\cdot), ..., g^m(\cdot)$ are continuously differentiable.
- Suppose that $\boldsymbol{x}^*$ is an solution to this problem (i.e. $\boldsymbol{x}^*$ is a maximizer subject to the constraints).
- If the gradient vectors of the *binding* constraints at $\boldsymbol{x}^*$ are linearly independent, then:
- There is a unique $\boldsymbol{\lambda}^* = (\lambda_1^*, ..., \lambda_m^*)$ such that

$$\frac{\partial L(\boldsymbol{x}^*, \boldsymbol{\lambda}^*)}{\partial x_1} = \frac{\partial f(\boldsymbol{x}^*)}{\partial x_i} - \sum_{j=1}^{m} \lambda_j^* \frac{\partial g^j(\boldsymbol{x}^*)}{\partial x_i} = 0 \text{ for } i = 1, ..., n$$

$$\lambda_j^* \geq 0, \quad g^j(\boldsymbol{x}^*) \leq 0, \quad \lambda_j^* g^j(\boldsymbol{x}^*) = 0, \quad \text{for } j = 1, ..., m$$

$$\frac{\partial L(\boldsymbol{x}^*, \boldsymbol{\lambda}^*)}{\partial x_1} = \frac{\partial f(\boldsymbol{x}^*)}{\partial x_i} - \sum_{j=1}^{m} \lambda_j^* \frac{\partial g^j(\boldsymbol{x}^*)}{\partial x_i} = 0 \text{ for } i = 1, \ldots, n$$

$$\lambda_j^* \geq 0, \;\; g^j(\boldsymbol{x}^*) \leq 0, \;\; \lambda_j^* g^j(\boldsymbol{x}^*) = 0, \;\; \text{for } j = 1, \ldots, m$$

- Let's understand what these conditions mean.
- Suppose that constraint j is not binding, or binding and irrelevant (i.e. the optimizer would not change if we removed this constraint).
- Then its associated Lagrange multiplier, λ_j^*, is zero.
- Constraints that are binding and not irrelevant (i.e. the optimizer *would* change if the constraint were removed) have a strictly positive Lagrange multiplier: $\lambda_j^* > 0$.
- For such a constraint, the condition $g^j(\boldsymbol{x}^*) = 0$ holds with equality.
- For constraints that are nonbinding or irrelevant, the constraint may hold without equality.
- The above conditions are a compact way of writing this down.

- Consider the following maximization problem with two constraints:

$$\max_x \ -x^2 \qquad \text{s.t.} \ -1 \leq x \leq 1$$

- We can rewrite the constraints as: $g^1(x) = -x - 1 \leq 0, g^2(x) = x - 1 \leq 0$.
- Clearly, the maximizer is $x^* = 0$, and both constraints are non-binding.
- Therefore, the Lagrange multipliers for both constraints are zero: $\lambda_1^* = \lambda_2^* = 0$.
- Both constraints are satisfied with strict inequality.
- The first-order conditions for the Lagrangian simply reduce to $f'(x^*) = 0$, since both Lagrange multipliers are zero.

- Now, suppose we change the problem so that one constraint becomes binding:

$$\max_x \; -x^2 \qquad \text{s.t. } 1 \le x \le 2$$

- We can rewrite the constraints as: $g^1(x) = -x + 1 \le 0, g^2(x) = x - 2 \le 0$.
- Since $-x^2$ is decreasing as $x > 0$, the maximizer is where the first constraint becomes binding, so $x^* = 1$.
- The Lagrange multiplier for the first, binding constraint is positive: $\lambda_1^* > 0$.
- The Lagrange multiplier for the second, non-binding constraint is zero: $\lambda_2^* = 0$.
- Only the second constraint is satisfied with strict inequality.
- The Lagrangian becomes: $L(x, \lambda_1, \lambda_2) = -x^2 - \lambda_1(-x + 1)$.
- When finding the solution, we treat the binding constraint as an equality, and ignore the non-binding constraint.

- The KKT conditions don't tell us which constraints will be binding or not. We have to use additional information.
- For example, in consumer theory, we will frequently have a constrained maximization problem with two goods, x_1 and x_2, and three constraints:
  - $x_1 \geq 0$
  - $x_2 \geq 0$
  - a budget constraint: $p_1 x_1 + p_2 x_2 \leq y$.
- If we had no additional information, then the solution might be in the interior (all constraints nonbinding), or at one of the corners (two constraints binding).
- However, we will typically assume the objective function is strictly increasing. Then the maximum must be on the budget constraint.
- Furthermore, we may assume the objective function goes to negative infinity at $x_1 = 0$ or $x_2 = 0$. Then the first two constraints will never be binding.
- Taken together, this tells us that only the budget constraint will be binding, and we can treat it as a problem with a single equality constraint.

Next Week

- For next week, please begin reading Chapter 1 of the textbook.
- We will begin the study of consumer theory.