Lecture 1

Intro, math and micro review

Ivan Rudik
AEM 6510

AEM 6510: Environmental and resource economics

Welcome to AEM 6510

AEM 6510: Environmental and resource economics

Welcome to AEM 6510

Class time: Tuesday and Thursday 2:45-4:00

Instructor: Ivan Rudik (ir229)

Office hours: Tuesday 4:00-5:00 Warren 462

Teaching assistant: Weiliang Tan (wt289)

TA office hours: Monday 3:00-4:00 Warren 372

AEM 6510: Environmental and resource economics

Textbooks:

- A Course in Environmental Economics by Phaneuf and Requate
- Causal Inference: The Mixtape by Cunningham

Formal Prerequisites: MATH 1110 or equivalent.

Highly recommended: Intermediate micro and masters-level econometrics

Lecture notes: https://github.com/irudik/aem6510, accessible through Canvas

What's the point of the course?

To learn:

Core environmental theory

What's the point of the course?

To learn:

Core environmental theory

Theory of regulation

What's the point of the course?

To learn:

Core environmental theory

Theory of regulation

Applied welfare analysis

What's the point of the course?

To learn:

Core environmental theory

Theory of regulation

Applied welfare analysis

How to do empirical analysis on the environment

What do you need to succeed in this class?

Attend lectures

What do you need to succeed in this class?

Attend lectures

Do your problem sets and prelims

What do you need to succeed in this class?

Attend lectures

Do your problem sets and prelims

Start your final project early

What do you need to succeed in this class?

Attend lectures

Do your problem sets and prelims

Start your final project early

Contact me or Weiliang if you have and questions or concerns

Here's how you'll be graded

Prelims (2): 30% and 25% (one theory, one empirical)

Problem sets (4): 20%

Literature review paper: 20%

Literature review presentation: 5%

A: 92-100; A-: 90-91

B+: 88-89; B: 82-87; B-: 80-81

C+: 78-79; C: 72-77; C-: 70-71

D+: 68-69; D: 62-67; D-: 60-61

F: < 60

Prelims

There are two, one on theory and one on empirical methods

Prelims

There are two, one on theory and one on empirical methods

Your highest-scoring prelim is 30% of your grade,
your lowest-scoring prelim is 25%

Prelims

There are two, one on theory and one on empirical methods

Your highest-scoring prelim is 30% of your grade,
your lowest-scoring prelim is 25%

The empirical prelim is takehome, and you are expected to complete it on your own

Prelims

There are two, one on theory and one on empirical methods

Your highest-scoring prelim is 30% of your grade,
your lowest-scoring prelim is 25%

The empirical prelim is takehome, and you are expected to complete it on your own

You have 48 hours to finish the empirical prelim

Prelims

There are two, one on theory and one on empirical methods

Your highest-scoring prelim is 30% of your grade,
your lowest-scoring prelim is 25%

The empirical prelim is takehome, and you are expected to complete it on your own

You have 48 hours to finish the empirical prelim

Late prelims automatically get 30% deducted: don't wait until the last minute to turn them in

Problem sets

There will be 4 problem sets worth 5% each

Problem sets

There will be 4 problem sets worth 5% each

Problem sets will be 3-4 questions long, typically from the Phaneuf and Requate book

Problem sets

There will be 4 problem sets worth 5% each

Problem sets will be 3-4 questions long, typically from the Phaneuf and Requate book

You will have at least 1 week to complete each problem set

Final project

There is a final project for the course aimed to get you moving toward research

Final project

There is a final project for the course aimed to get you moving toward research

You can choose one of two options:

1) Literature review

2) Data dive

Final project

There is a final project for the course aimed to get you moving toward research

You can choose one of two options:

1) Literature review

2) Data dive

TBD but depending on the number of students you may be able to work in pairs

Final project: lit review

Write a literature review on up to 3 papers all in a particular subfield of environmental and resource economics

Final project: lit review

Write a literature review on up to 3 papers all in a particular subfield of environmental and resource economics

- What are the findings? Common threads? Work yet to be done?
- You'll need to run the papers by me before committing

Final project: lit review

Write a literature review on up to 3 papers all in a particular subfield of environmental and resource economics

- What are the findings? Common threads? Work yet to be done?
- You'll need to run the papers by me before committing

You will present your literature review to the class

Final project: lit review

Write a literature review on up to 3 papers all in a particular subfield of environmental and resource economics

- What are the findings? Common threads? Work yet to be done?
- You'll need to run the papers by me before committing

You will present your literature review to the class

More details later

Final project: data dive

Find a new data set that we will not use in class but is useful for environmental economics

Final project: data dive

Find a new data set that we will not use in class but is useful for environmental economics

- What are the data? How do you get them? How do you use them? What makes them relevant for economics?
- You'll need to run it by me

Final project: data dive

Find a new data set that we will not use in class but is useful for environmental economics

- What are the data? How do you get them? How do you use them? What makes them relevant for economics?
- You'll need to run it by me

Do some preliminary analyses on the data

Final project: data dive

Find a new data set that we will not use in class but is useful for environmental economics

- What are the data? How do you get them? How do you use them? What makes them relevant for economics?
- You'll need to run it by me

Do some preliminary analyses on the data

Present the dataset to the class

Final project: data dive

Find a new data set that we will not use in class but is useful for environmental economics

- What are the data? How do you get them? How do you use them? What makes them relevant for economics?
- You'll need to run it by me

Do some preliminary analyses on the data

Present the dataset to the class

More details later

Course schedule: Theory

1. Micro & math recap
2. Theory of externalities
3. Intro to the theory of environmental policy
4. Imperfect information
5. Competitive output markets
6. Non-competitive output markets
7. Pre-existing distortions
8. Theory of applied welfare economics
9. Revealed preference models
10. Hedonics (theory)
11. Theory prelim

Course schedule: Empirics

1. R and the tidyverse, causal inference
2. Deforestation, regression discontinuity
3. Hedonics (applied)
4. Environmental health, difference-in-differences, and event studies
5. Climate change science
6. Climate change and the Ricardian model
7. Climate change and two-way fixed effects
8. Climate change in space
9. Empirical prelim office hours

Part 1: Math and micro review

Before getting into the core material we need to review the mathematical foundations underpinning the economics

Part 1: Math and micro review

Before getting into the core material we need to review the mathematical foundations underpinning the economics

Constrained optimization is central in economic decisionmaking

Part 1: Math and micro review

Before getting into the core material we need to review the mathematical foundations underpinning the economics

Constrained optimization is central in economic decisionmaking

- We want to max utility subject to a budget

Part 1: Math and micro review

Before getting into the core material we need to review the mathematical foundations underpinning the economics

Constrained optimization is central in economic decisionmaking

- We want to max utility subject to a budget
- We want to max profits subject to a constraining regulation

Part 1: Math and micro review

Before getting into the core material we need to review the mathematical foundations underpinning the economics

Constrained optimization is central in economic decisionmaking

- We want to max utility subject to a budget
- We want to max profits subject to a constraining regulation
- We want to max social welfare subject to an environmental target

Part 1: Math and micro review

Before getting into the core material we need to review the mathematical foundations underpinning the economics

Constrained optimization is central in economic decisionmaking

- We want to max utility subject to a budget
- We want to max profits subject to a constraining regulation
- We want to max social welfare subject to an environmental target
- We want to minimize costs subject to meeting electricity demand

Part 1: Math and micro review

Before getting into the core material we need to review the mathematical foundations underpinning the economics

Constrained optimization is central in economic decisionmaking

- We want to max utility subject to a budget
- We want to max profits subject to a constraining regulation
- We want to max social welfare subject to an environmental target
- We want to minimize costs subject to meeting electricity demand
- Etc

Unconstrained optimization

Unconstrained optimization is just Calculus 101

Unconstrained optimization

Unconstrained optimization is just Calculus 101

Suppose we want to maximize utility from eating fruit:

Unconstrained optimization

Unconstrained optimization is just Calculus 101

Suppose we want to maximize utility from eating fruit:

$$\max_{x_1,\dots,x_N \in \mathbf{A}} U(\mathbf{x})$$

where $U(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}$

Unconstrained optimization

Unconstrained optimization is just Calculus 101

Suppose we want to maximize utility from eating fruit:

$$\max_{x_1,\ldots,x_N \in \mathbf{A}} U(\mathbf{x})$$

where $U(\mathbf{x}) : \mathbb{R}^N \to \mathbb{R}$

Theorem: Let U be once continuously-differentiable ($\mathbb{C}^1$) so that first partial derivatives exist and are continuous. If x^* is a local maximum of U on A then

Unconstrained optimization

Unconstrained optimization is just Calculus 101

Suppose we want to maximize utility from eating fruit:

$$\max_{x_1,\ldots,x_N \in \mathbf{A}} U(\mathbf{x})$$

where $U(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}$

Theorem: Let U be once continuously-differentiable ($\mathbb{C}^1$) so that first partial derivatives exist and are continuous. If x^* is a local maximum of U on A then

$$\frac{\partial U(\mathbf{x}^*)}{\partial x_i} = 0 \text{ for } i = 1, \ldots, N$$

Unconstrained optimization

What does this theorem tell us?

Unconstrained optimization

What does this theorem tell us?

That the first-order conditions give us the **necessary conditions** for the utility-maximizing amount of each type of fruit to eat

Unconstrained optimization

What does this theorem tell us?

That the first-order conditions give us the **necessary conditions** for the utility-maximizing amount of each type of fruit to eat

These conditions must be satisfied for any bundle that maximizes utility

Unconstrained optimization

What does this theorem tell us?

That the first-order conditions give us the **necessary conditions** for the utility-maximizing amount of each type of fruit to eat

These conditions **must** be satisfied for any bundle that maximizes utility

In a logical statement, *if X then Y*, Y is the necessary condition: Y **must** be true for X to be true

X is the **sufficient condition**

Unconstrained optimization

$$\frac{\partial U(\mathbf{x}^*)}{\partial x_i} = 0 \text{ for } i = 1, \ldots, N$$

Our necessary conditions tell us that the marginal utility of each fruit equals zero at the optimal choice

If x^* is a maximum of U on A, then the tangent line/plane of U at $(U(x^*), x^*)$ must be horizontal

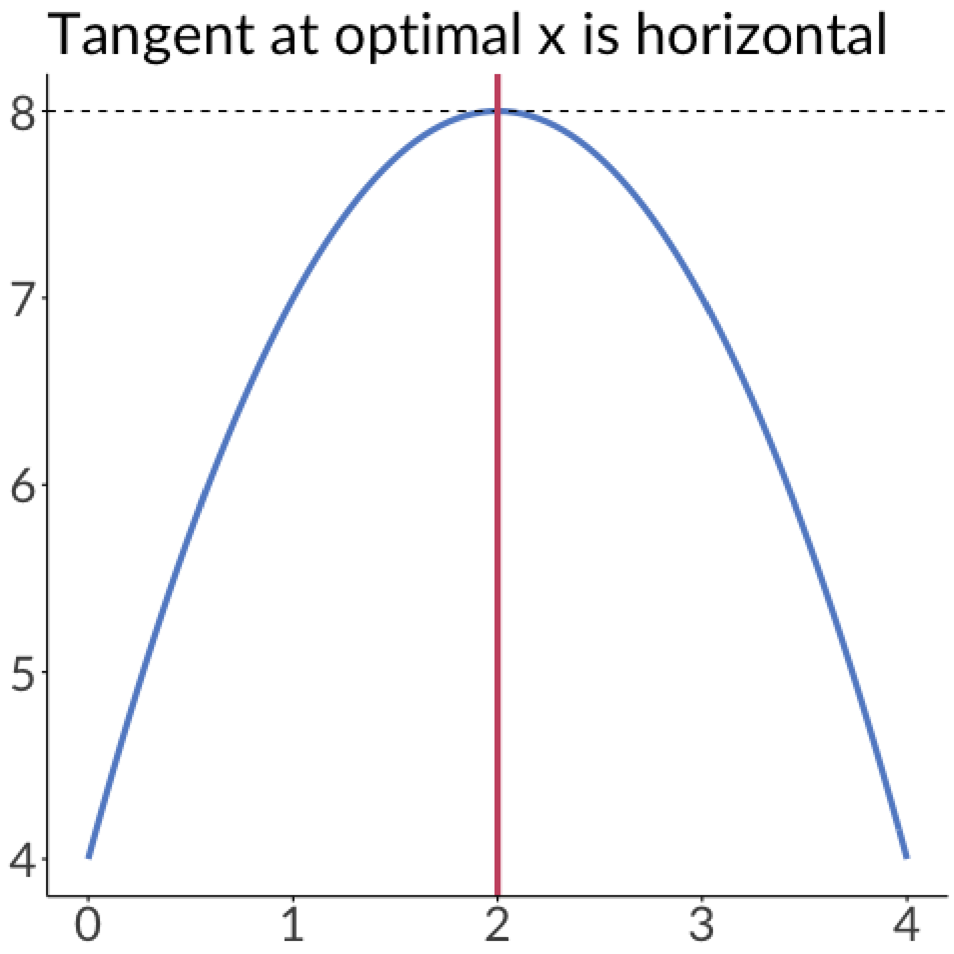

Unconstrained optimization

Do necessary conditions guarantee that we have found the utility-maximizing bundle of fruit?

Unconstrained optimization

Do necessary conditions guarantee that we have found the utility-maximizing bundle of fruit?

NO!

Unconstrained optimization

Do necessary conditions guarantee that we have found the utility-maximizing bundle of fruit?

NO!

They are necessary but not sufficient: you can come up with counter-examples where a function's derivative equals zero but it is not maximized (e.g. x^3)

Unconstrained optimization

Do necessary conditions guarantee that we have found the utility-maximizing bundle of fruit?

NO!

They are necessary but not sufficient: you can come up with counter-examples where a function's derivative equals zero but it is not maximized (e.g. x^3)

We need additional conditions to impose sufficiency

Unconstrained optimization

These are the second-order sufficient conditions

Unconstrained optimization

These are the second-order sufficient conditions

Theorem: Let U be twice continuously differentiable. Suppose that x^* is a critical point of U on A and its Hessian matrix is negative (semi-)definite at x^* ($D^2U(x^*)$ is negative (semi-)definite). Then x^* is a strict (weak) local maximum.

Unconstrained optimization

These are the second-order sufficient conditions

Theorem: Let U be twice continuously differentiable. Suppose that x^* is a critical point of U on A and its Hessian matrix is negative (semi-)definite at x^* ($D^2U(x^*)$ is negative (semi-)definite). Then x^* is a strict (weak) local maximum.

Negative semi-definite is basically just saying that a multivariate function is weakly concave

Unconstrained optimization

Words Details | Math Details

We need second-order conditions to guarantee we have found a maximum

We need marginal utility to be decreasing (i.e. utility be concave) at x^*

This guarantees x^* maximizes U if x^* satisfies the first-order necessary conditions

Constrained optimization

Before we were able to eat as much fruit as we wanted

Constrained optimization

Before we were able to eat as much fruit as we wanted

What if our actions are now constrained?

Constrained optimization

Before we were able to eat as much fruit as we wanted

What if our actions are now constrained?

$$\max_{x_1,\ldots,x_N \in \mathbf{A}} U(\mathbf{x}), \quad A = \{\mathbf{x} \in \mathbb{R}^n : h_1(\mathbf{x}) = a_1, \ldots, h_m(\mathbf{x}) = a_m\}$$

where $m < n$

Constrained optimization

A simple and familiar example is utility maximization:

$$\max_{x_1, x_2} U(x_1, x_2) \text{ subject to } h(x_1, x_2) = a$$

We want to be on the **highest indifference curve** subject to our budget constraint

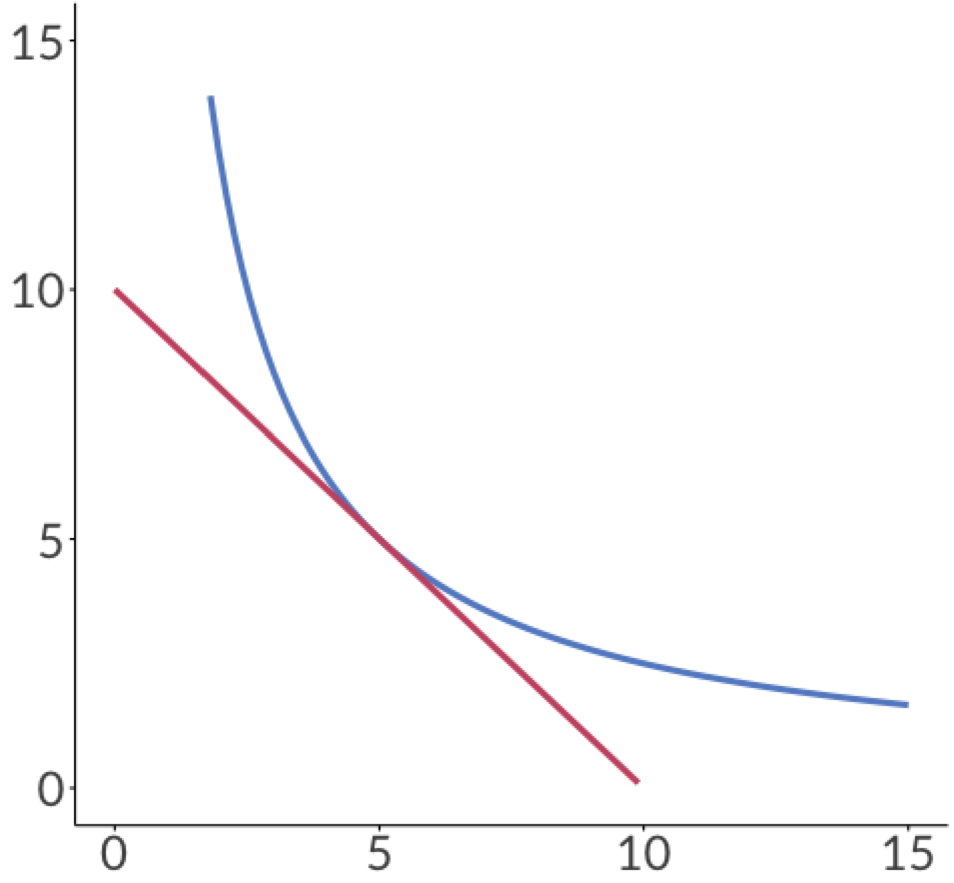

Constrained optimization

Because of our budget constraint, marginal utility may not equal zero at the utility-maximizing bundle!

Constrained optimization

Because of our budget constraint, marginal utility may not equal zero at the utility-maximizing bundle!

No proof, but for some λ we can characterize this bundle by

Constrained optimization

Because of our budget constraint, marginal utility may not equal zero at the utility-maximizing bundle!

No proof, but for some λ we can characterize this bundle by

$$DU(x_1^*, x_2^*) - \lambda Dh(x_1^*, x_2^*) = (0, 0)$$

and

Constrained optimization

Because of our budget constraint, marginal utility may not equal zero at the utility-maximizing bundle!

No proof, but for some λ we can characterize this bundle by

$$DU(x_1^*, x_2^*) - \lambda Dh(x_1^*, x_2^*) = (0, 0)$$

and

$$h(x_1^*, x_2^*) = a$$

Constrained optimization

Because of our budget constraint, marginal utility may not equal zero at the utility-maximizing bundle!

No proof, but for some λ we can characterize this bundle by

$$DU(x_1^*, x_2^*) - \lambda Dh(x_1^*, x_2^*) = (0, 0)$$

and

$$h(x_1^*, x_2^*) = a$$

We now have 3 equations, and 3 unknowns $(x_1^*, x_2^*, \lambda^*)$, we just transformed our constrained problem into an unconstrained problem

Constrained optimization

How does this magic happen?

Constrained optimization

How does this magic happen?

Intuition: We can reframe a constrained optimization problem in an alternative fashion where we allow the agent to pick whatever (x_1, x_2) they want, **but**, they must pay a fine λ per unit violation of the restriction

Constrained optimization

How does this magic happen?

Intuition: We can reframe a constrained optimization problem in an alternative fashion where we allow the agent to pick whatever (x_1, x_2) they want, **but**, they must pay a fine λ per unit violation of the restriction

This gives us an alternative, unconstrained problem, known as a *Lagrangian*:

$$\mathcal{L}(x_1, x_2, \lambda) = \max_{x_1, x_2, \lambda} U(x_1, x_2) - \lambda[h(x_1, x_2) - a]$$

Constrained optimization

How does this magic happen?

Intuition: We can reframe a constrained optimization problem in an alternative fashion where we allow the agent to pick whatever (x_1, x_2) they want, **but**, they must pay a fine λ per unit violation of the restriction

This gives us an alternative, unconstrained problem, known as a *Lagrangian*:

$$\mathcal{L}(x_1, x_2, \lambda) = \max_{x_1, x_2, \lambda} U(x_1, x_2) - \lambda[h(x_1, x_2) - a]$$

Maximize utility, but also pay a fine λ for each unit difference between $h(x_1, x_2)$ and a

Constrained optimization

$$\max_{x_1, x_2, \lambda} U(x_1, x_2) - \lambda[h(x_1, x_2) - a]$$

with first-order conditions:

Constrained optimization

$$\max_{x_1,x_2,\lambda} U(x_1,x_2) - \lambda[h(x_1,x_2) - a]$$

with first-order conditions:

$$\frac{\partial \mathcal{L}}{\partial x_i} = \frac{\partial U}{\partial x_i} - \lambda \frac{\partial h}{\partial x_i} = 0 \text{ for } i = 1,2$$

Constrained optimization

$$\max_{x_1, x_2, \lambda} U(x_1, x_2) - \lambda[h(x_1, x_2) - a]$$

with first-order conditions:

$$\frac{\partial \mathcal{L}}{\partial x_i} = \frac{\partial U}{\partial x_i} - \lambda \frac{\partial h}{\partial x_i} = 0 \text{ for } i = 1, 2$$

Notice marginal utility is no longer zero at the optimum allocation unless $\lambda = 0$

Constrained optimization

We have one remaining issue: some arbitrary λ won't guarantee that the solution of our Lagrangian will match the actual constrained problem

Constrained optimization

We have one remaining issue: some arbitrary λ won't guarantee that the solution of our Lagrangian will match the actual constrained problem

We need to find a way to pick the **correct** λ^* that incentivizes the agent to exactly satisfy the constraint $h(x_1, x_2) = a$

Constrained optimization

We have one remaining issue: some arbitrary λ won't guarantee that the solution of our Lagrangian will match the actual constrained problem

We need to find a way to pick the **correct** λ^* that incentivizes the agent to exactly satisfy the constraint $h(x_1, x_2) = a$

In other words, we need a third condition to pin down λ^*, this is the feasibility condition:

$$\frac{\partial \mathcal{L}}{\partial \lambda} = h(x_1, x_2) - a = 0$$

Which is just the constraint!

Constrained optimization

We can translate a constrained optimization problem into an unconstrained one (the Lagrangian) by:

Constrained optimization

We can translate a constrained optimization problem into an unconstrained one (the Lagrangian) by:

1. Adding a new variable λ, which is called the **Lagrange Multiplier** or the **shadow value** of the constraint

Constrained optimization

We can translate a constrained optimization problem into an unconstrained one (the Lagrangian) by:

1. Adding a new variable λ, which is called the **Lagrange Multiplier** or the **shadow value** of the constraint
2. Adding another first-order condition (feasibility condition) and modifying the original first-order conditions to reflect the constraint

Constrained optimization

We can translate a constrained optimization problem into an unconstrained one (the Lagrangian) by:

1. Adding a new variable λ, which is called the **Lagrange Multiplier** or the **shadow value** of the constraint
2. Adding another first-order condition (feasibility condition) and modifying the original first-order conditions to reflect the constraint

The intuition for λ is that it is the marginal value of loosening (or tightening) the constraint

Constrained optimization

We can translate a constrained optimization problem into an unconstrained one (the Lagrangian) by:

1. Adding a new variable λ, which is called the **Lagrange Multiplier** or the **shadow value** of the constraint
2. Adding another first-order condition (feasibility condition) and modifying the original first-order conditions to reflect the constraint

The intuition for λ is that it is the marginal value of loosening (or tightening) the constraint

In constrained utility maximization, λ is the marginal utility of income