Economics of Networks
Introduction to Game Theory: Part 1

Evan Sadler

Massachusetts Institute of Technology

Agenda

Decision theory:

- Theory of rational choice
- Choice under uncertainty

Games

- Definitions
- Strategies and best responses

Pure strategy Nash Equilibrium

Examples

Suggested reading: EK chapter 6; Osborne chapters 1-3

Motivation

Choice is a big part of economics

- What choices do we observe and why?

- What information do I share?
- What information sources do I trust?
- Do I want to become friends?
- From whom do I buy stuff?
- What websites do I visit?
- How do I get to work in the morning?

How do we model this?

Rational Choice

Primitive in microeconomic models: the economic agent

- "[A] unit that responds to a scenario called a *choice problem*," (Rubinstein, 2007)

Typically think of an agent as a person

- Sometimes a firm, a nation, a family

Three step process of rational choice:

- What is desirable?
- What is feasible?
- Choose the most desirable of the feasible alternatives

Observations:

- Note desires precede recognition of feasible alternatives
- Economics says nothing about what agent should desire

Preferences

Let X denote set of all conceivable alternatives

- e.g. set of all undergraduate institutions to which one could be admitted

A rational agent has a *preference ordering* over X

- For each $x, y \in X$, either $x \succeq y$ or $x \preceq y$
- Ordering must be complete and transitive:

$$x \succeq y \text{ and } y \succeq z \implies x \succeq z$$

Choice problem $C \subseteq X$

- Choose $\succeq$-maximal element of C

Utility Representation

Associate a number $u(x)$ to each $x \in X$, *utility* of x

- $u(x) \geq u(y)$ if and only if $x \succeq y$
- Easier to work with utility functions than directly with preferences

Definition of rational choice uses only *ordinal* information

- Says nothing about preference intensity
- All monotone transformations of u are equivalent

With uncertainty, we need *cardinal* information

- How do I compare outcomes with different probabilities?

Choice Under Uncertainty

von Neumann and Morgenstern proposed a set of axioms for choice over "lotteries"

Formally:

- Let Y denote a set of possible outcomes
- Let $\mathcal{L}_Y$ denote the set of lotteries over Y
- Let L denote a particular lottery

Examples:

- An even money gamble: with probability $\frac{1}{2}$ you lose \$10, with probability $\frac{1}{2}$ you win \$10
- Getting to work: with probability $\frac{9}{10}$ it takes you 20 minutes, with probability $\frac{1}{10}$ there is road work and it takes an hour

Choice Under Uncertainty

A choice problem $C \subseteq \mathcal{L}_Y$ is a collection of lotteries over Y

- An action induces a lottery (e.g. place a bet or not, which route do I take to work?)

Need to define preferences over lotteries $L, M, N \in \mathcal{L}_Y$

The vNM axioms:

- Completeness: $L \succeq M$ or $L \preceq M$ for all L, M
- Transitivity: $L \preceq M$ and $M \preceq N$ implies $L \preceq N$
- Continuity: if $L \preceq M \preceq N$, there exists $p \in [0, 1]$ such that

$$pL + (1 - p)N \sim M$$

- Independence: if $L \prec M$, then for any N and $p \in (0, 1]$, we have

$$pL + (1 - p)N \prec pM + (1 - p)N$$

Choice Under Uncertainty

Theorem

Suppose preferences $\preceq$ satisfy the vNM axioms. There exists a utility function u on the set of outcomes Y such that $L \preceq M$ if and only if

$$\mathbb{E}[u(L)] \leq \mathbb{E}[u(M)].$$

Expected utility theory

Suppose action a induces distribution $F^a(y)$ over consequences, expected utility

$$U(a) = \int u(y) \mathrm{d}F^a(y)$$

Choice Under Uncertainty

Choose a over b if

$$U(a) = \int u(y) \mathrm{d}F^a(y) > \int u(y) \mathrm{d}F^b(y) = U(b)$$

Model of rationality is conceptually simple

- In practice, computation may be difficult

Notes:

- Objective versus Subjective uncertainty
- Savage (1954) axiomatizes subjective probability and subjective expected utility

From Single Agent to Multi-Agent Choice

In most social situations, the outcome an agent cares about depends on actions of others

- Benefit from working hard on group work depends on others' efforts
- Time I spend stuck in traffic depends on others' route choices
- Disutility from a rumor spreading depends on others' sharing

- How hard should I work?
- What route should I take?
- Should I tell a friend a secret?

Work or Shirk?

You are paired with a friend for a group project

- You can each work hard or shirk
- Working has an effort cost, but leads to better grades for both

	Work	Shirk
Work	$(2, 2)$	$(-1, 1)$
Shirk	(1,-1)	$(0, 0)$

Matrix game, entries represent payoffs

- Not necessarily monetary

Do you work or shirk?

Bertrand Competition

Two firms produce identical goods at a fixed marginal cost $c > 0$

Consumers buy from the lowest priced firm, split evenly if they charge the same price

Total demand at price p is $D(p) = 1 - p$

Firms simultaneously set prices, consumers then make purchases

What price do you set?

Normal Form Games

In a normal form game, players make one choice, simultaneously

Elements of a game

- Set of players
- Set of actions or strategies
- Payoffs

Player order, multiple moves, and information sets captured in extensive form games

- To come later

Normal Form Games

Definition (Normal Form Game)

A normal form game is a triple $(N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ such that

- $N = \{1, 2, ..., n\}$ is a set of players
- S_i is the set of actions available to player i
- $u_i : S \to \mathbb{R}$ is the payoff of player i, where $S = \prod_{i \in N} S_i$ is the set of all action profiles

Some notation:

- s_{-i}: vector of actions for all players except i
- S_{-i}: set of action profiles for all players except i
- $s = (s_i, s_{-i}) \in S$ is an action profile, or outcome

Choice Under Uncertainty

Theorem

Suppose preferences $\preceq$ satisfy the vNM axioms. There exists a utility function u on the set of outcomes Y such that $L \preceq M$ if and only if

$$\mathbb{E}[u(L)] \leq \mathbb{E}[u(M)].$$

Expected utility theory

Choice Under Uncertainty

Choose a over b if

$$U(a) = \int u(y)\mathrm{d}F^a(y) > \int u(y)\mathrm{d}F^b(y) = U(b)$$

Model of rationality is conceptually simple

- In practice, computation may be difficult

From Single Agent to Multi-Agent Choice

In most social situations, the outcome an agent cares about depends on actions of others

- Benefit from working hard on group work depends on others' efforts
- Time I spend stuck in traffic depends on others' route choices
- Disutility from a rumor spreading depends on others' sharing

Work or Shirk?

You are paired with a friend for a group project

- You can each work hard or shirk
- Working has an effort cost, but leads to better grades for both

	Work	Shirk
Work	$(2, 2)$	$(-1, 1)$
Shirk	(1,-1)	$(0, 0)$

Matrix game, entries represent payoffs

- Not necessarily monetary

Bertrand Competition

Two firms produce identical goods at a fixed marginal cost $c > 0$

Consumers buy from the lowest priced firm, split evenly if they charge the same price

Total demand at price p is $D(p) = 1 - p$

Firms simultaneously set prices, consumers then make purchases

Normal Form Games

In a normal form game, players make one choice, simultaneously

Elements of a game

- Set of players
- Set of actions or strategies
- Payoffs

Normal Form Games

Definition (Normal Form Game)

A normal form game is a triple $(N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N})$ such that

- $N = \{1, 2, ..., n\}$ is a set of players
- S_i is the set of actions available to player i
- $u_i : S \to \mathbb{R}$ is the payoff of player i, where $S = \prod_{i\in N} S_i$ is the set of all action profiles

Strategies versus Actions

Often use strategy and action interchangably, though there is a formal distinction

In an extensive form game (i.e. with multiple moves in sequence, think chess), a strategy means a complete contingent plan of action at every possible realization of play

Extensive form game is equivalent in a sense to a normal form game

- The action in the normal form game is the selection of a strategy in the extensive form

Distinction also exists in normal form games when we talk about mixed strategies

The "Solution" of a Game

How do people play the game?

How should people play the game?

Large number of "solution concepts" based on different assumptions about

- What players know about each others' plans
- How smart players are
- How smart players think others are

Dominant Strategies

Example: The Prisoner's Dilemma

	Confess	Silence
Confess	$(-3,-3)$	$(0,-4)$
Silence	(-4,0)	$(-1,-1)$

What should the outcome be?

- "Confess" is always the better choice

"Confess" dominates "silence"

Dominant Strategy Equilibrium

A fairly compelling solution concept: everyone plays a dominant strategy

- Play a strategy that is obviously good

Definition

A strategy $s_i^* \in S_i$ is *dominant* for player i if for all $s_i \in S_i$ and all $s_{-i} \in S_{-i}$

$$u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$$

A strategy profile s^* is a *dominant strategy equilibrium* if s_i^* is a dominant strategy for each i.

Issue: this rarely exists

Dominated Strategies

Conversely, we might think to eliminate strategies that are dominated

- Don't play a strategy that is obviously bad

Definition

A strategy $s_i \in S_i$ is *strictly dominated* if there exists $s_i' \in S_i$ such that for all $s_{-i} \in S_{-i}$

$$u_i(s_i', s_{-i}) > u_i(s_i, s_{-i})$$

A strategy $s_i \in S_i$ is *weakly dominated* if there exists $s_i' \in S_i$ such that for all $s_{-i} \in S_{-i}$

$$u_i(s_i', s_{-i}) \geq u_i(s_i, s_{-i})$$

Iterated Deletion

No one should play a dominated strategy

- Common knowledge of payoffs and rationality implies iterated elimination of dominated strategies

Example:

	Confess	Silence	Suicide
Confess	$(-3,-3)$	$(0,-4)$	$(-3,-10)$
Silence	(-4,0)	$(-1,-1)$	$(-1,-10)$
Suicide	$(-10,-3)$	$(-10,-1)$	$(-10,-10)$

No dominant strategy

- Still dominance solvable

Iterated Deletion

More formally, define the iterative procedure:

- Step 0: Define $S_i^0 = S_i$ for each i
- Step k $>$ 0: Define for each i

$$S_i^k = \left\{ s_i \in S_i^{k-1} \,|\, \nexists\, s_i' \in S_i^{k-1} \text{ that dominates } s_i \right\}$$

- Step ∞: $S_i^\infty = \bigcap_{k=0}^\infty S_i^k$

The set S^∞ of strategy profiles is what survives

Iterated Deletion

Theorem

Suppose that either:

- S_i *is finite for each* i*, or*
- $u_i(s_i, s_{-i})$ *is continuous and* S_i *is compact for each* i*.*

Then S_i^∞ *is nonempty for each* i*.*

First part trivial, second part homework

May not yield a unique prediction

Best Responses

Another approach: suppose players make conjectures about each other

- Beliefs about what other players will do

Formally, a conjecture for player i is a lottery over S_{-i}

Given a conjecture μ, choose

$$s_i \in BR(\mu) = \arg\max_{s_i} \int u_i(s_i, s_{-i}) \mathrm{d}\mu(s_{-i})$$

Should play a best reply to *some* conjecture

- What conjectures should a player entertain?

Pure Strategy Nash Equilibrium

Nash Equilibrium: conjectures are correct

Definition

A *pure strategy Nash Equilibrium* is a strategy profile $s^* \in S$ such that for all $i \in N$ we have

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*)$$

for all $s_i \in S_i$.

Every player is playing a best response to what others actually do

Why might this be reasonable?

Interpretation of Nash Equilibrium

Two main justifications:

- Introspection: try to be consistent, assuming everyone else is also smart
- Learning: can think of Nash Equilibrium as the steady state of a learning process

Another idea: ex-ante agreement among the players

Nash Equilibrium is a standard workhorse in economic models

- Might not be reasonable in all contexts

Example: Bertrand Competition

Recall our two competiting firms

- Marginal cost $c > 0$
- Total demand $D(p) = 1 - p$
- Firms choose what prices to charge

Can $p_1 \geq p_2 > c$ be an equilibrium?

- No. Firm 1 should charge $p_2 - \epsilon$ and steal the market

Would a firm ever charge less than c?

What about $p_1 = p_2 = c$?

- Yes! Both firms earn zero profit, no way to improve.

Example: Cournot Competition

What if the firms choose what quantity to produce instead?

- Both face the market price $p = 1 - q_1 - q_2$

Given q_1, firm 2 chooses q_2 to maximize

$$q_2(1 - q_1 - q_2 - c)$$

First order condition implies $q_2^* = \frac{1-q_1-c}{2}$

Similarly, $q_1^* = \frac{1-q_2-c}{2}$

Unique Nash Equilibrium:

$$q_1^* = q_2^* = \frac{2(1-c)}{3}$$

Example: Cournot Competition

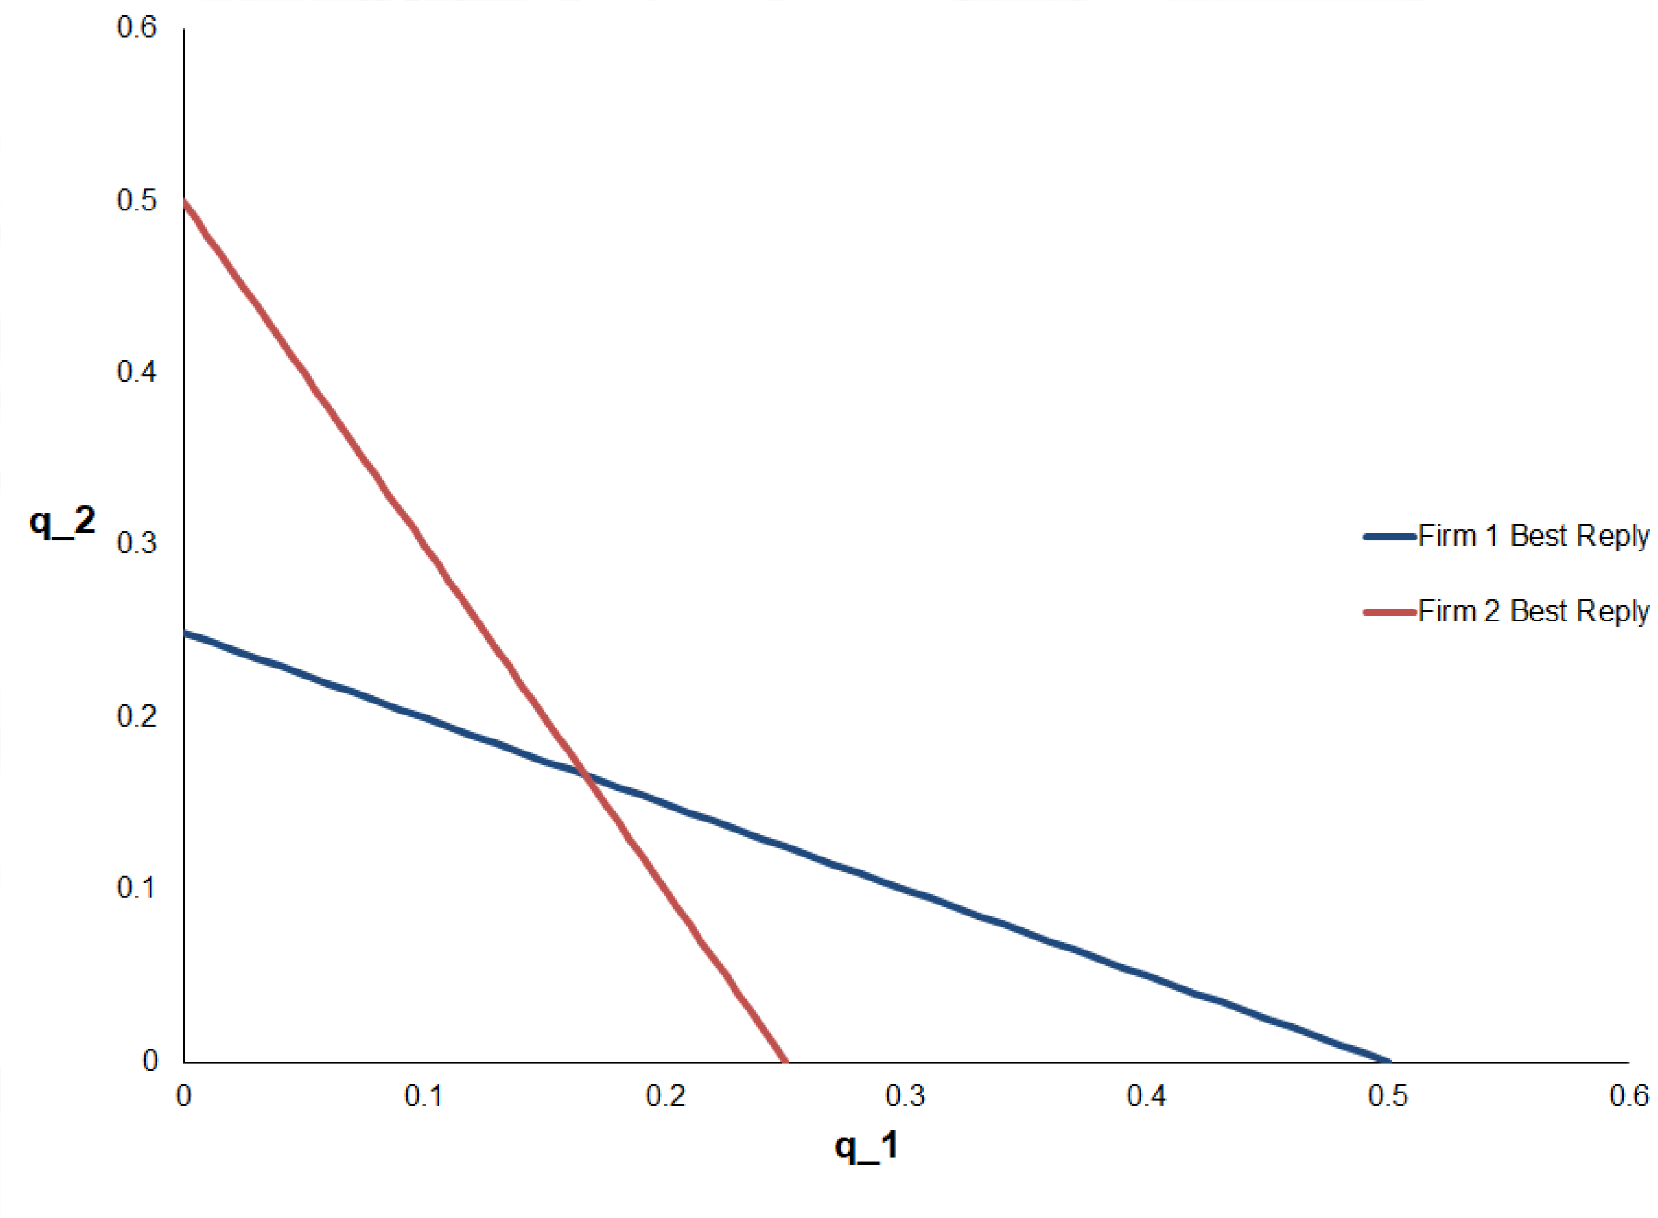

Example: The Partnership Game

Recall our earlier choice to work or shirk:

	Work	Shirk
Work	$(2, 2)$	$(-1, 1)$
Shirk	(1,-1)	$(0, 0)$

No dominant or dominated strategies

Best reply to work is work, best reply to shirk is shirk

- Two pure strategy Nash Equilibria
- Outcome depends on conjectures

Multiple Equilibria

What do we do with multiple equilibria?

- Model lacks a unique prediction

Two approaches

- Acceptance: we make set valued predictions, certain outcomes are possible, and we still rule out a lot of alternatives
- Refinement: come up with an argument why one equilibrium is more realistic than another

Refinement is hard

Example: Matching Pennies

Consider the following game:

	Heads	Tails
Heads	$(-1, 1)$	$(1, -1)$
Tails	(1,-1)	$(-1, 1)$

What is the best response to heads?

What is the best response to tails?

No pure strategy Nash Equilibrium exists

- Next time: mixed strategy equilibrium