“Game Theory"

Yves Zenou
Stockholm University
Office: A781
Email: yves.zenou@ne.su.se
Office hours: Tuesdays 16.00-17.00

Lecture Note 4

- Mixed-strategy Nash equilibrium
- Attitudes toward risk
- Applications

Main questions

- What is a mixed strategy game?
- Payoff functions and vNM preferences
- Definition mixed strategy Nash equilibrium
- Best responses
- What if it’s difficult to draw the best responses?
- Pure and mixed equilibria
- Formation of players’ beliefs
- Soccer penalties – the Neeskens effect

What is a mixed strategy game?

- ***Players***: as before
- ***Actions***: as before, but now each player assigns a probability of playing each action.
- ***Preferences*** are represented by the expected utility of choosing different actions.

vNM expected utility: $U=pu(c_0)+(1-p)u(c_1)$

  - ☐ Where p is the probability of receiving payoff c_0 and $u(c_0)$ is the utility you get if this state of the world occurs.

- Now cardinal utility is important! – that is, not only the ranking of payoffs is important, but also how different the payoffs are!

Attitudes toward risk

- Games in which players face uncertain choices.
- You have two people with exactly the same odds and payoffs, but they may behave differently.
- We need to figure out how to incorporate different attitudes into games that we analyze.

Attitudes toward risk

- Player's attitudes towards risks in games are manifested in their payoffs.
- Dominance arguments work the same, but now we are incorporating attitudes about risk.
- These are called **von Neumann-Morgenstern (VNM) payoffs**. Most of the time, attitude towards risk is built into the game, so usually VNM is assumed in the game.

Attitudes toward risk

- On the horizontal axis, we have money from $0 to $100. The horizontal axis has **Utility Payoff**.

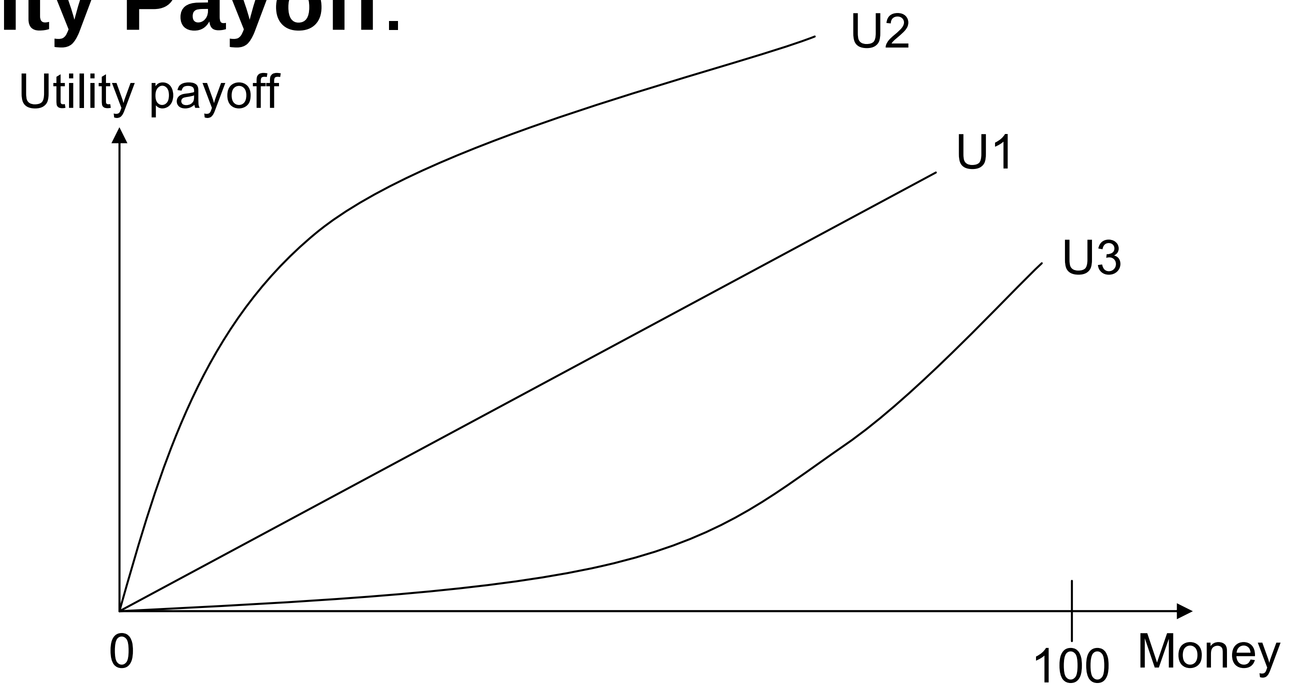

Attitudes toward risk

- **Utility function:**
- Assume player thinks more is better than less
- Incorporates player’s attitude towards risk
- Satisfies the von Neumann-Morgenstern (VNM) property

Example: Lottery Game (Property 2—Player's risk attitudes)

- Options:
- L1: win \$100 with probability ½, win \$0 with probability ½
- L2: win \$90 with probability 2/3, win \$30 with probability 1/3
- L3: win \$50 with probability 1 (win for sure)
- Which lottery would you choose?
- Many people prefer L2>L1
- Most people prefer L3>L1

Example: Lottery Game (Property 2—Player’s risk attitudes)

- If you played L1 a large number of times, your winnings would eventually average out to $50.
- However, L3 guarantees $50.
- In the long run, the amount of money you’d expect to walk away with is the same with L1 and L3.
- Whether you choose L1 or L3 tells us something about your willingness to run risk.

Example: Lottery Game Continued (Property 3—VNM)

- We need to attach utility to each outcome in a game. We do this by taking the **expected utility**— everything that could happen, weighted by the probability that it will happen.
- The utility of an outcome is equal to its expected utility

Example: Lottery Game Continued (Property 3—VNM)

- The following expected utility takes the utility of winning U(100) times the probability of winning plus the utility of losing U(0) times the probability of losing (Lottery L1).
- U(L1) = U(100)*1/2 + U(0)*1/2

Example: Lottery Game Continued (Property 3—VNM)

- We need to find a utility function that would satisfy this condition.
- In this game, because we are measuring utility by money, the **Expected Monetary Payoff** (EMP) is the same as the expected utility.
- EMP = 100*1/2 + 0*1/2

Example: Student's Risk Preference

- We're going to assess a student's preference towards risk, then look at the VNM property.

- Lottery (P) = \$100 with probability (p), \$0 with probability (1-p)

- *Question*: What if I offered you the lottery with winning probability 0.1, or \$1 for sure?
- *Answer:* \$1

- *Question*: What if I offered you the lottery or \$0.50 for sure?
- *Answer:* The lottery.

- *Question*: What would be the maximum amount you'd be willing to pay if you have a 10% chance of winning this lottery?
- *Answer:* I'd take \$0.75

Example: Student's Risk Preference

- So at $0.75, the student would be indifferent between playing the lottery with a 10% chance of winning and taking the $0.75 for sure.
- In this situation, $0.75 is the student's **Certainty Equivalent (CE)**—an amount in which a player will be indifferent between playing the lottery or taking the CE money for sure.

Example: Student's Risk Preference

- Here are the student's Certainty Equivalents:

Probability	**CE**
0.1	0.75
0.2	5
0.4	15
0.6	20
0.8	30

Example: Student's Risk Preference

- Now we can assess her attitude towards risk, and draw a utility function that satisfies all three properties.
- Assigned Utilities for \$0 and \$100:
- $U(0) = 0$
- $U(100) = 10$

Example: Student's Risk Preference

Prob	**CE**	**Utility (L)**
0.1	0.75	U(100)*.1 + U(0)*.9 = 1
0.2	5	U(100)*.2 + U(0)*.8 = 2
0.4	15	U(100)*.4 + U(0)*.6 = 4
0.6	20	U(100)*.6 + U(0)*.4 = 6
0.8	30	U(100)*.8 + U(0)*.2 = 8

Example: Student's Risk Preference

- Plot her Certainty Equivalents (x-axis) against the Utility Payoffs (y-axis).

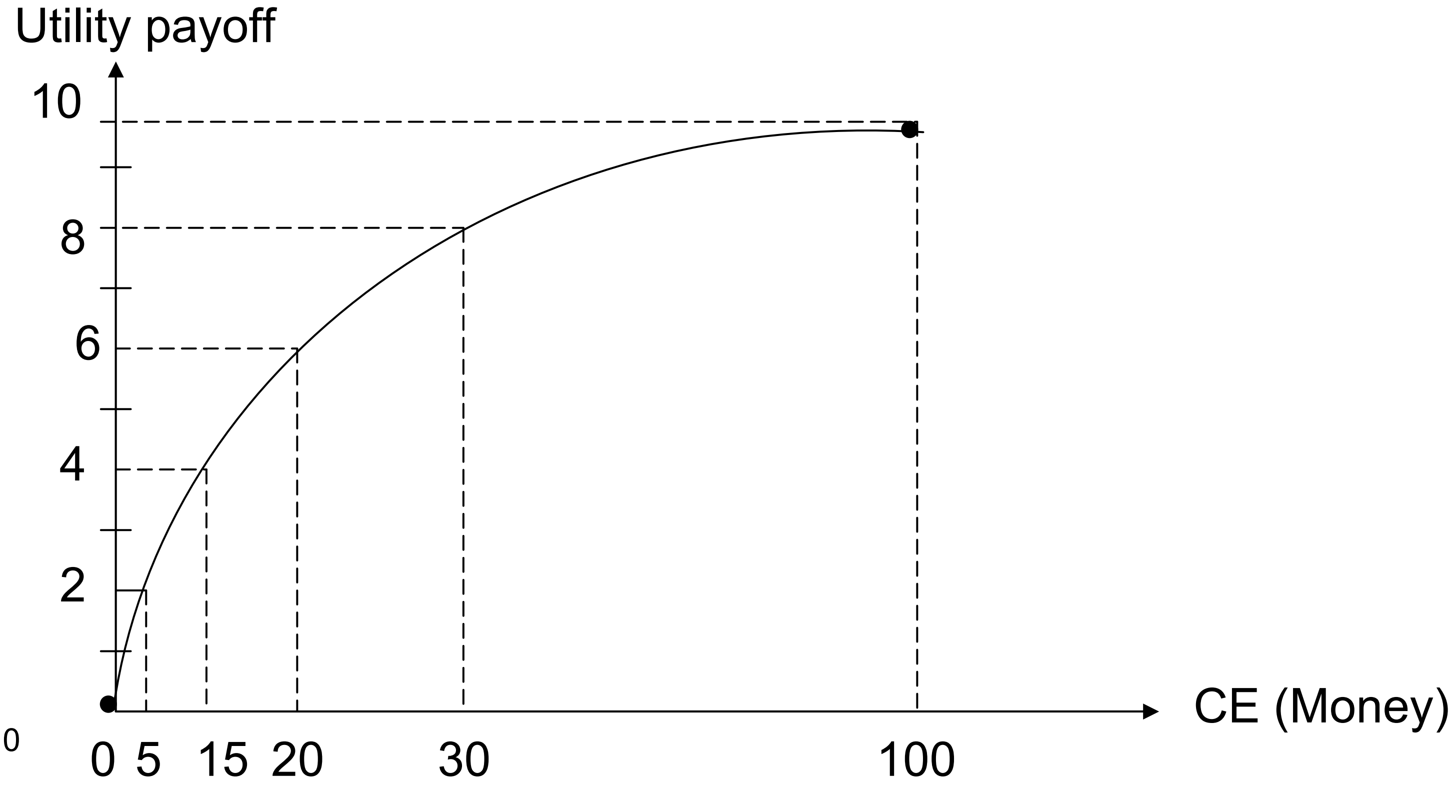

Attitudes toward risk

- There are three broad attitudes toward risk:
- **Risk Adverse**—people who do not like risk
- **Risk Acceptance**—people who like risk
- **Risk Neutral**—people who are indifferent

Attitudes toward risk

- L1 = $100 with prob ½, $0 with prob ½
- L3 = $50 for sure
- What if a player is **risk adverse**?
- So if a risk adverse person had a choice between L1 and L3, they will choose L3.
- What do we know about that person's Certainty Equivalent for L1? Is it equal to 50, bigger than 50, or less than 50? **It will be less than 50.**
- Think about someone's Certainty Equivalent as the largest amount of money the person would be willing to take *in place of* playing the lottery.

Attitudes toward risk

- So a Risk Adverse Player would have the following preference:
- CE(L1) < EMP(L1)
- This is because they would rather walk away with a smaller amount of money for sure, than risk gambling.
- EMP(L1) - CE(L1): the individual’s **risk premium** for the lottery L1

Attitudes toward risk

- A Risk Neutral Player will have this preference:
- CE(L1) = EMP(L1)
- She has no preference between the expected value of the gamble and taking the certainty equivalent.

Attitudes toward risk

- A Risk Lover Player will have this preference:
- CE(L1) > EMP(L1)
- This person is willing to take a lower CE because they are attracted to the upside of the game (winning $100), and thus would risk the gamble.

Attitudes toward risk

- On the horizontal axis, we have money from $0 to $100. The horizontal axis has **Utility Payoff**.

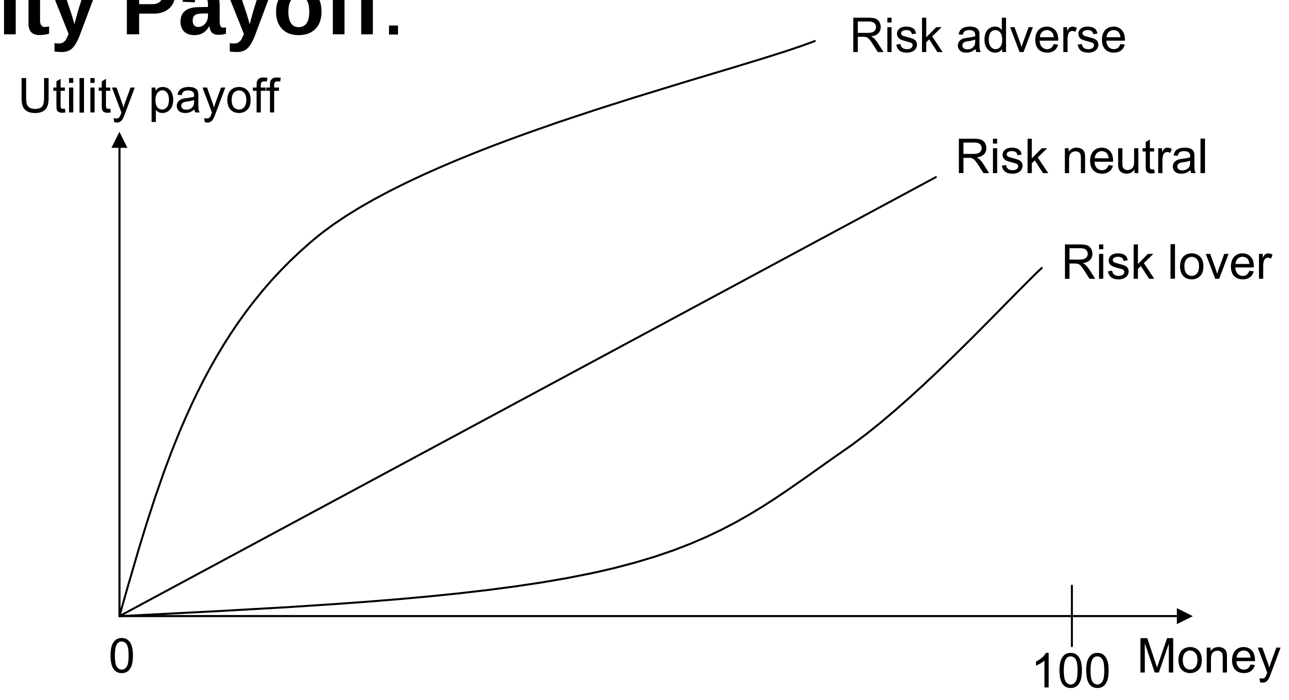

Payoff functions – what can and can't be done to them without changing the game.

- A payoff function over deterministic outcomes
  - The ranking of different outcomes will not be affected by a positive and monotone transformation of the original payoff function.
- A payoff function over deterministic outcomes with uncertainty
  - The ranking of different "lotteries" according to VNM expected utility (payoff) will not be affected by a positive, monotone AND linear transformation of the original payoff function.

Illustrations of transformed payoff functions I

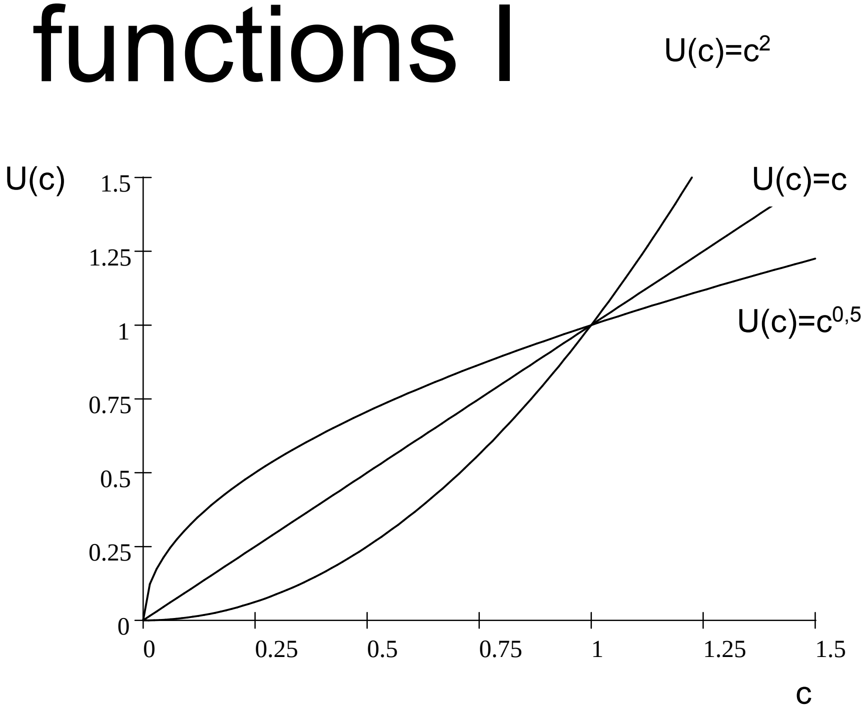

The ordinal ranking (over c) is not affected by any positive and monotone transformation (more is still always better).

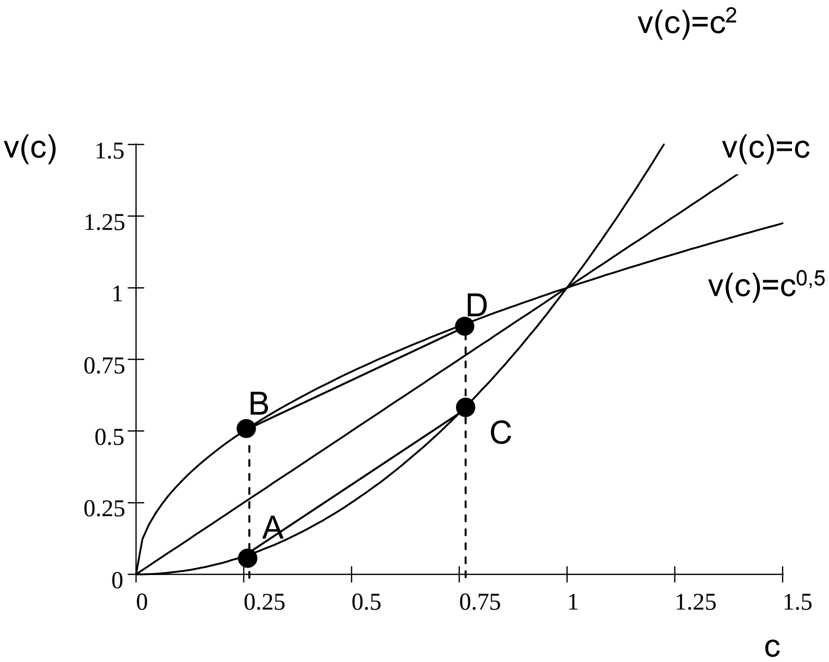

Compare a sure income of c=0,5 with a lottery which yields either 0,25 or 0,75 with 50% probability each. The ranking of the sure income and the lottery is affected by the same transformations.

Illustrations of transformed utlity functions II

Security dilemma

		State 2	
		Disarm	*Arm*
State 1	*Disarm*	3, 3	0, 2
	Arm	2, 0	1, 1

A) Payoffs to the square (positive, monotone transformation)

		State 2	
		Disarm	*Arm*
State 1	*Disarm*	9, 9	0, 4
	Arm	4, 0	1, 1

B) Payoffs multiplied by 0,5 (positive, monotone & linear transformation)

		State 2	
		Disarm	*Arm*
State 1	*Disarm*	1,5, 1,5	0, 1
	Arm	1, 0	0,5, 0,5

Exercise: Expected utility

- Suppose you have $10,000 to invest. A broker phones you with the information you requested on certain junk bonds.
- If the company issuing the bonds posts a **profit** this year, it will pay you a **40% interest rate** on the bond.
- If the company files for **bankruptcy** you will **lose** all you invested
- If the company **breaks even**, you will earn a **10% interest rate**.
- Your broker tells you there is a **50% chance that they will break even** and a **20% chance that the company will file for bankruptcy**.

Exercise: Expected utility

- Your other option is to invest in a risk-free government bond that will guarantee 8% interest for a year.

Question (a)

- (a) What is the expected interest rate for the junk bond investment?

Answer to question (a)

- 50% break even (10% int. rate), 20% bankruptcy (0% int. rate), 30% profit (40% int. rate).
- Expected Return =

  0.5 ($10,000+$1,000) + 0.2 ($0)

  + 0.3 ($10,000+$4,000) = 9700
- Therefore, expected interest rate

  = (9700-10,000)/10,000 = -0.03; i.e., -3%.

Question (b)

- (b) What investment will you choose if your utility function were given by $U(M) = M^2$?
- Observe: $U(M) = M^2$ means risk lover.

Answer to question (b)

- Junk bond: 50% break even (10% int. rate), 20% bankruptcy (0% int. rate), 30% profit (40% int. rate).
- Expected Utility =

$$0.5\,(\$10{,}000+\$1{,}000)^2+0.2\,(\$0)^2$$

$$+\,0.3\,(\$10{,}000+\$4{,}000)^2=\$119{,}300{,}000$$

- Risk-free government bond (8% int. rate)
- Expected Utility =

$$(\$10{,}000+\$800)^2=116{,}640{,}000$$

- Hence, you would chose to invest in the junk bond

Question (c)

- (c) What is the certainty equivalent of the chosen investment?
- Obs: the individual’s **Certainty Equivalent (CE)** is the amount for which the individual is indifferent between investing her money in the junk bond or taking the CE money for sure.

Answer to question (c)

- Denote the certainty equivalent of the junk bond by CE.
- Then,

  EU (M from junk bond) = U (CE)
- This is equivalent to:

$$\$119,300,000 = (\mathrm{CE})^2$$

- which implies

$$\mathrm{CE} = \sqrt{\$119,300,000} = \$10,922.5$$

Answer to question (c)

- So at $10,925.5, the individual would be indifferent between investing in a bond (with a 20% chance of losing everything) and taking the $10,925.5 for sure.

- Observe that: Expected Return =

0.5 ($10,000+$1,000) + 0.2 ($0)

+ 0.3 ($10,000+$4,000) = $9700

- CE = $10,925.5 > ER = $9700 (risk lover)

Question (d)

- (d) Which investment would you choose if your utility function were given by $U(M) = \sqrt{M}$?
- Observe $U(M) = \sqrt{M}$ means risk aversion

Answer to question (d)

- Junk bond: 50% break even (10% int. rate), 20% bankruptcy (0% int. rate), 30% profit (40% int. rate).
- Expected Utility =

$$0.5\sqrt{\$10{,}000+\$1{,}000}+0.2\sqrt{\$0}$$

$$+0.3\sqrt{\$10{,}000+\$4{,}000}=\$87.94$$

- Risk-free government bond (8% int. rate)
- Expected Utility =

$$\sqrt{\$10{,}000+\$800}=\$103.92$$

- Hence, you would chose to invest in Risk-free government bond.

Some reminders on probabilities, random variables and distributions

Definition: A *probability space* is a pair (S, P), where S is a finite set, called the sample space, and P is a function that assigns to every subset A of S a real number between 0 and 1. If we let $S = \{s_1, s_2, ..., s_n\}$ and $p_i = \text{Pr}(\{s_i\})$ (the probability of a singleton $\{s_i\}$, the the function P satisfies the following properties:

(1) $p_i \geq 0$ for each i and $\sum\limits_{i=1}^{i=n} p_i = 1$;

(2) If A is a subset of S, the $\text{Pr}(A) = \sum\limits_{s_i \in A} p_i$.

In particular, we have:

$$\text{Pr}(\emptyset) = 0 \text{ and } \text{Pr}(S) = 1$$

The subsets of the sample set S are called *events*.

A *random variable* is a function $X : S \to \mathbb{R}$

Definition (Distribution of a random variable): If $X : S \to \mathbb{R}$ is a random variable, then the (cumulative) distribution function of X is the function $F : \mathbb{R} \to \mathbb{R}$ defined by:

$$F(x) = \Pr\{X \leq x\}$$

for each real number x.

Theorem (Properties if Distributions): If $F : \mathbb{R} \to \mathbb{R}$ is the distribution function of a random variable X, then:

(i) $0 \le F(x) \le 1$ for each x;

(ii) F is increasing, that is, $F(x) \le F(y)$ if $x < y$;

(iii) $\lim_{x \to -\infty} F(x) = 0$ and $\lim_{x \to +\infty} F(x) = 1$;

(iv) $\Pr\{a < X \le b\} = F(b) - F(a)$ if $a < b$;

(v) $\Pr\{X = a\} =$ the jump of the distribution function $F(x)$ at $x = a$.

Let us first consider the case when random variables can take on a *finite* or countably infinite set values (Discrete random variables)

Definition (Expectation, Variance for Discrete Distribution of Random Variables): If $X : S \rightarrow \mathbb{R}$ is a random variable, then

(i) The expected value $\mathbb{E}(X)$ of X is the real number:

$$\mathbb{E}(X) = \sum_{i=1}^{i=n} p_i X(s_i)$$

(ii)The variance $Var(X)$ of X is the real number:

$$Var(X) = \sum_{i=1}^{i=n} p_i \left[X(s_i) - \mathbb{E}(X)\right]^2$$

Let us now consider the case when random variables can take any value of a real line or certain subset of the real line.

Definition (density function): A distribution function F of a random variable is said to have a *density function* if there exists a non-negative function $f : \mathbb{R} \to \mathbb{R}$ (called the density function) such that:

$$F(x) = \int\limits_{-\infty}^{x} f(t)dt$$

holds for all x. The density function is positive an satisfies the property $\int\limits_{-\infty}^{+\infty} f(t)dt = 1$.

For example,

$$\mathsf{Pr}\left\{a \leq X \leq b\right\} = \int\limits_{a}^{b} f(t)dt$$

Definition (Expectation, Variance for Continuous Distribution of Random Variables): If $X : S \to \mathbb{R}$ is a random variable, then

(i) The expected value $\mathbb{E}(X)$ of X is the real number:

$$\mathbb{E}(X) = \int\limits_{-\infty}^{+\infty} t\, f(t) dt$$

(ii)The variance $Var(X)$ of X is the real number:

$$Var(X) = \int\limits_{-\infty}^{+\infty} [t - \mathbb{E}(X)]^2\, f(t) dt$$

The Uniform Distribution

Definition (Uniform Distribution): A random variable X is said to have a uniform distribution F over a finite closed interval $[a, b]$ if F has the density function:

$$f(x) = \begin{cases} 1/(b-a) & \text{if} \quad a < x < b \\ 0 & \text{if} \quad x \leq a \text{ or } x \geq b \end{cases}$$

The (cumulative) distribution function is given by:

$$F(x) = \int_{-\infty}^{x} t\, f(t)dt = \begin{cases} 0 & \text{if} \quad x \leq a \\ (x-a)/(b-a) & \text{if} \quad a < x < b \\ 1 & \quad\quad x \geq b \end{cases}$$

Also:

$$\mathbb{E}(X) = \frac{a+b}{2} \text{ and } Var(X) = \frac{(b-a)^2}{12}$$

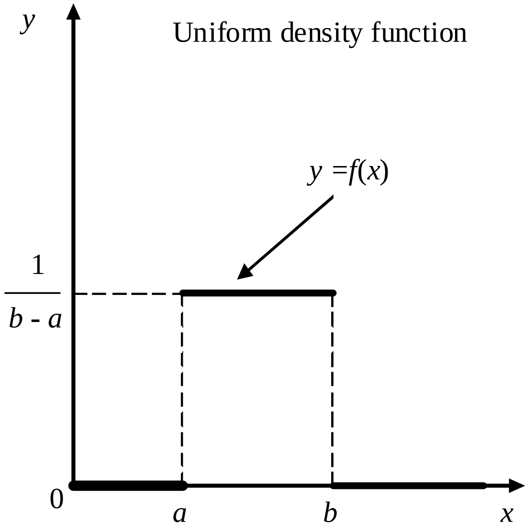
y
Uniform density function
y =f(x)
1
b - a
0
a
b
x

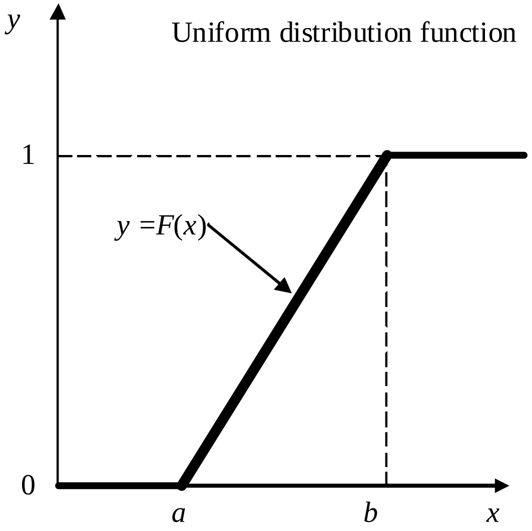
y
Uniform distribution function
1
y =F(x)
0
a
b
x

The Normal Distribution

A random variable X is said to have a Normal Distribution with parameters m and σ^2 if the distribution has the density function:

$$f(x) = \frac{1}{\sigma 2\pi} \exp\left[-\frac{(x-m)^2}{2\sigma^2}\right]$$

That is, the (cumulative) distribution function of a normally distributed random variable X with the parameters m and σ^2 is given by the formula:

$$F(x) = \frac{1}{\sigma 2\pi} \int_{-\infty}^{x} \exp\left[-\frac{(t-m)^2}{2\sigma^2}\right] dt$$

Also:

$$\mathbb{E}(X) = m \text{ and } Var(X) = \sigma^2$$

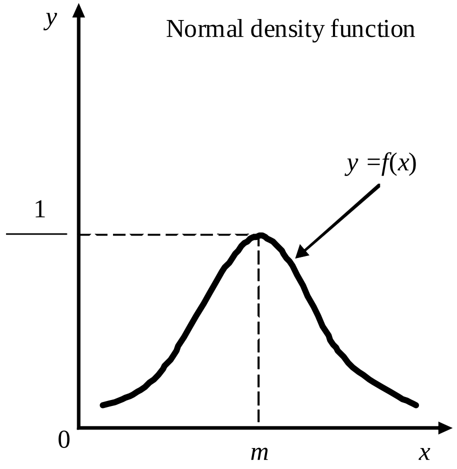
y
Normal density function
y =f(x)
1
0
m
x

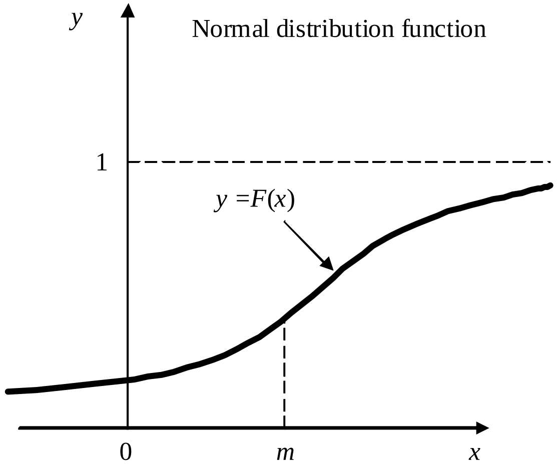
y
Normal distribution function
1
y =F(x)
0
m
x

A random variable that is normally distributed with parameter $m = 0$ and $\sigma = 1$ is said to have the standard normal distribution:

$$F_S(x) = \frac{1}{2\pi} \int\limits_{-\infty}^{x} \exp\left[-\frac{t^2}{2}\right] dt$$

In general, if X is a normally distributed random variable with parameters m and σ^2, then:

$$F(x) = F_S(\frac{x - m}{\sigma})$$

Lotteries

We denote a **lottery** L by the collectionof pairs $L = \{(x_i, p_i) : i = 1, ..., n\}$, where x_i corresponds to wealth of individual i.

The lottery $L = \{(x_i, p_i) : i = 1, ..., n\}$ has n possible alternative outcomes and each outcome i occurs with probability i.

Consider two different lotteries: $L_1 = \{(x_i, p_i) : i = 1, ..., n\}$ and $L_2 = \left\{\left(x'_i, p'_i\right) : i = 1, ..., n\right\}$. We define a **compound lottery** as $pL_1 + (1-p)\, L_2$, which is:

$$pL_1 + (1-p)\, L_2 = \left\{\left(p\, x_i + (1-p)\ x'_i\ ,\ p\, p_i + (1-p)\, p'_i\right) : i = 1, ..., n\right\}$$

Definition (The Independence Axiom): An individual's choice over lotteries satifies the *Independence Axiom* whenever a lottery L_1 is preferred to another lottery L_2, then for each $0 < p < 1$, the compound lottery $pL_1 + (1 - p)\, L_3$ is preferred to the compound lottery $pL_2 + (1 - p)\, L_3$ for all lotteries L_3.

If someone prefers lottery L_1 to lottery L_2, then this axiom means that the ranking of L_1 and L_2 remains uchanged if we mix a third lottery L_3 with both L_1 and L_2.

Definition (The Continuity Axiom): An individual's choice over lotteries satifies the *Continuity Axiom* whenever a sequence $\{p_n\}$ of probabilities (i.e. $0 \leq p_n \leq 1$ holds for each n) converges to p, that is $p_n \to p$, and the lottery $p_n L_1 + (1 - p_n) L_2$ is preferred to a lottery L_3 for all n, then $pL_1 + (1 - p) L_2$ is preferred to L_3.

An individual's choice over lotteries satisfies the continuity axiom if small changes in the probabilities with which lotteries are chosen change the rank over the lotteries only slightly.

Let $\mathcal{L}$ denote the set of all lotteries.

Theorem (Expected Utility): If an individual's utility function $U : \mathcal{L} \to \mathbb{R}$ over the set of lotteries satisfies the *Independence Axiom* and the *Continuity Axiom*, then there is a von Neumann-Morgenstern utility function u over wealth such that

$$U(L) \equiv \mathbb{E}\left[u(x)\right] = \sum_{i=1}^{i=n} p_i u(x_i)$$

for every lottery $L = \{(x_i, p_i) : i = 1, ..., n\}$.

If the lottery is described by a continuous distribution over an interval $X = [a, b]$ with cumulative distribution function $F : X \to [0, 1]$, then the von Neumann-Morgenstern utility function u over wealth is such that:

$$U(L) \equiv \mathbb{E}\left[u(x)\right] = \int_a^b u(x) f(x) dx$$

Exercise: Choosing the optimal portfolio

Suppose an individual has \$10,000 to invest between a stock and a bond. The stock is a financial asset which has a variable return that is *uniformly distributed* with an average return of 8.525% and a standard deviation of 3.767%. The bond returns 8.5% with certainty. Denote by $0 \leq s \leq 1$ be the proportion of the \$10,000 invested in the stock.

The individual is risk averse and her utility function is:

$$u(x) = (x)^{1/2}$$

Given that the investor maximizes her expected utility, what will be the proportion s of the \$10,000 that will be invested in the stock?

The investor will choose the portfolio that *maximizes her expected utility*.

First, for a uniform distribution, we have:

$$\mathbb{E}\left[u(x), s\right] = \int\limits_{a}^{b} u(x) f(x) dx = \frac{1}{(b-a)} \int\limits_{a}^{b} u(x) dx = \frac{1}{(b-a)} \int\limits_{a}^{b} x dx$$

Let $0 \leq s \leq 1$ be the proportion of the \$10,000 invested in the stock. In this case, her expected utility is:

$$\begin{aligned}
\mathbb{E}\left[u(x), s\right] &= \frac{1}{(b-a)} \int_a^b \left[10,000\, s(1+r) + 10,000(1-s)\,(1+0.085)\right]^{1/2} dr \\
&= \frac{100}{(b-a)} \int_a^b \left[s(1+r) + (1-s)\,(1+0.085)\right]^{1/2} dr \\
&= \frac{100}{(b-a)} \int_a^b \left[(r-0.085)\,s + 1.085\right]^{1/2} dr
\end{aligned}$$

where $[a, b]$ is the interval which gives the possible returns from the stock and r represents the return from the stock.

Since the mean of the uniform distribution is 0.0825 and the standard deviation 0.03767, we can calculate a and b because:

$$\mathbb{E}(X) = \frac{a+b}{2} = 0.0825 \text{ and } \sigma^2 = \frac{(b-a)^2}{12} = (0.03767)^2$$

This gives

$$a + b = 0.1705 \text{ and } b - a = \frac{6 \times 0.03767}{3}$$

By combining these two equations, we obtain:

$$a = 0.02 \text{ and } b = 0.1505$$

Therefore,

$$\mathbb{E}\left[u(x), s\right] = 766.2\,8 \int_{0.02}^{0.1505} \left[(r - 0.085)\, s + 1.085\right]^{1/2} dr$$

Choosing s that maximizes $\mathbb{E}\left[u(x), s\right]$ leads to:

$$\frac{\partial \mathbb{E}\left[u(x), s\right]}{\partial s} = 766.2\,8 \int_{0.02}^{0.1505} \frac{(r - 0.085)}{\left[(r - 0.085)\, s + 1.085\right]^{1/2}} dr = 0$$

and we obtain

$$s^{*} \approx 0.9996$$

As a result, the investor will put 99.96% of her $10,000 in stocks and the rest (0.04%) in bonds.

Mixed strategy Nash equilibrium I

- A mixed strategy for a player in a strategic game is a probability distribution over his/her available actions.

Formally;

$\alpha_i(a_i) =$ the probability assigned by player i's mixed strategy α_i to her action a_i.

- For a complete mixed strategy, probabilities need to be assigned to all possible actions.

Mixed strategy Nash equilibrium II

- *Mixed strategy Nash equilibrium in a strategic game with vNM preferences*

For each player i;

$U_i(\alpha^*) \geq U_i(\alpha_i, \alpha_{-i}{}^*)$ for all α_i of player i,

where $U_i(\alpha)$ is player i's vNM expected payoff of the mixed strategy profile α.

Denote a game as $\Gamma = [I, \{\Delta S_i\}, \{u_i(.)\}]$ where $I = \{1, 2, ..., n\}$ is the set of players, S_i is the set of strategies of player i and $u_i(.)$ is the utility function of player i. $\{\Delta S_i\}$ indicates that we allow for the possibility of *randomized* choices by the players. $\{S_i\}$ indicates that we only consider *pure strategies*.

Theorem 0.1 *Every game* $\Gamma = [I, \{\Delta S_i\}, \{u_i(.)\}]$ *in which the sets* S_1, *...,* S_n *have a finite number of elements has a mixed strategy Nash equilibrium.*

This theorem is true only for games with finite strategy sets. However, in economic applications, we frequently encounter games in which players have strategies modeled as continuous variables (e.g. Cournot duopoly).

Infinite Strategy Sets (Myerson)

So far, we dealt with finite games because easier.

Consider now games in which players have *infinitely* many strategies.

We want to include the set of pures strategies for a player which may be a bounded interval on the real number line, such as $[0, 1]$. This set includes all rational and irrational number between 0 and 1, so it an *uncountably infinite set*.

To state results in greater generality, we need to assume that, for each player i, the strategy set S_i is a *compact metric space*.

A compact metric space is a general mathematical structure for representing infinite sets that can well be approximated by large finite sets.

One important fact is that, in a compact metric space, any infinite sequence has a convergence subsequence.

Any closed and bounded subset of a finite-dimensional vector space is an example of a compact metric space.

More specifically, any closed and bounded interval of real numbers is an example of a compact metric space, where the distance between two points x and y is $|x - y|$.

When there are infinitely many actions in the set S_i, a randomized strategy for player i can no longer be described by just listing the probability of each individual action.

For example, suppose that $S_i = [0, 1]$. If player i selects her action from a uniform distribution over $[0, 1]$, then each individual action in $[0, 1]$ would have zero probability.

To describe a distribution probability distribution over S_i, we must list the probabilities of subsets of S_i.

For technical reasons, it may be mathematical impossible to consistently assign probabilities to all subsets of an infinite set, so weak restriction is needed on the class of subsets whose probabilities can be meaningfully defined. These are called the *measurable sets*.

Theorem 0.2 *A Nash equilibrium exists in game* $\Gamma = [I, \{S_i\}, \{u_i(.)\}]$ *if for all* $i = 1, ..., n$,

(i) S_i *is a nonempty, convex, and compact subset of some Euclidean space* $\mathbb{R}^M$;

(ii) $u_i(s_1, ..., s_n)$ *is continuous in* $(s_1, ..., s_n)$ *and quasi-concave in* s_i.

This Theorem provides a significant result whose requirements are satisfied in a wide range of applications.

The convexity of strategy sets and the nature of payoff functions help to smooth out the structure of the model, allowing us to achieve a pure strategy equilibrium.

Of course, these results do not mean that we cannot have an equilibrium if the conditions of these existence results do not hold. Rather, we just cannot be assured that there is one.

How think about a mixed NE?

1. Each member in a population chooses particular actions according to the probabilities in the Nash equilibrium.
2. Each member in a population chooses only one action. The proportion of the population that chooses a particular action corresponds to the probability that this action will be played.

Best response functions in a strategic game with vNM preferences

$B_i(\alpha_{-i})$ is the set with best response mixed strategies for player i when the other players play their mixed strategies α_{-i}.

Example, best response

Matching Pennies

			Player 2	
			Head	*Tail*
			q	*(1-q)*
Player 1	*Head*	*p*	1, -1	-1, 1
	Tail	*(1-p)*	-1, 1	1, -1

- How calculate the best response?
- How illustrate this in a graph?
- What is the Nash equilibrium, if there is one?

Best responses: Matching pennies Pure Strategies

- If player 2 plays ”Head” , the *BR* for player 1 is:

$$BR_1(H) = H$$

- If player 2 plays ”Tail” , the *BR* for player 1 is:

$$BR_1(T) = T$$

- If player 1 plays ”Head” , the *BR* for player 2 is:

$$BR_2(H) = T$$

- If player 1 plays ”Tail” , the *BR* for player 2 is:

$$BR_2(T) = H$$

Matching pennies Pure Strategies: No NE (no fixed point)

■ : BR for player 2

● : BR for player 1

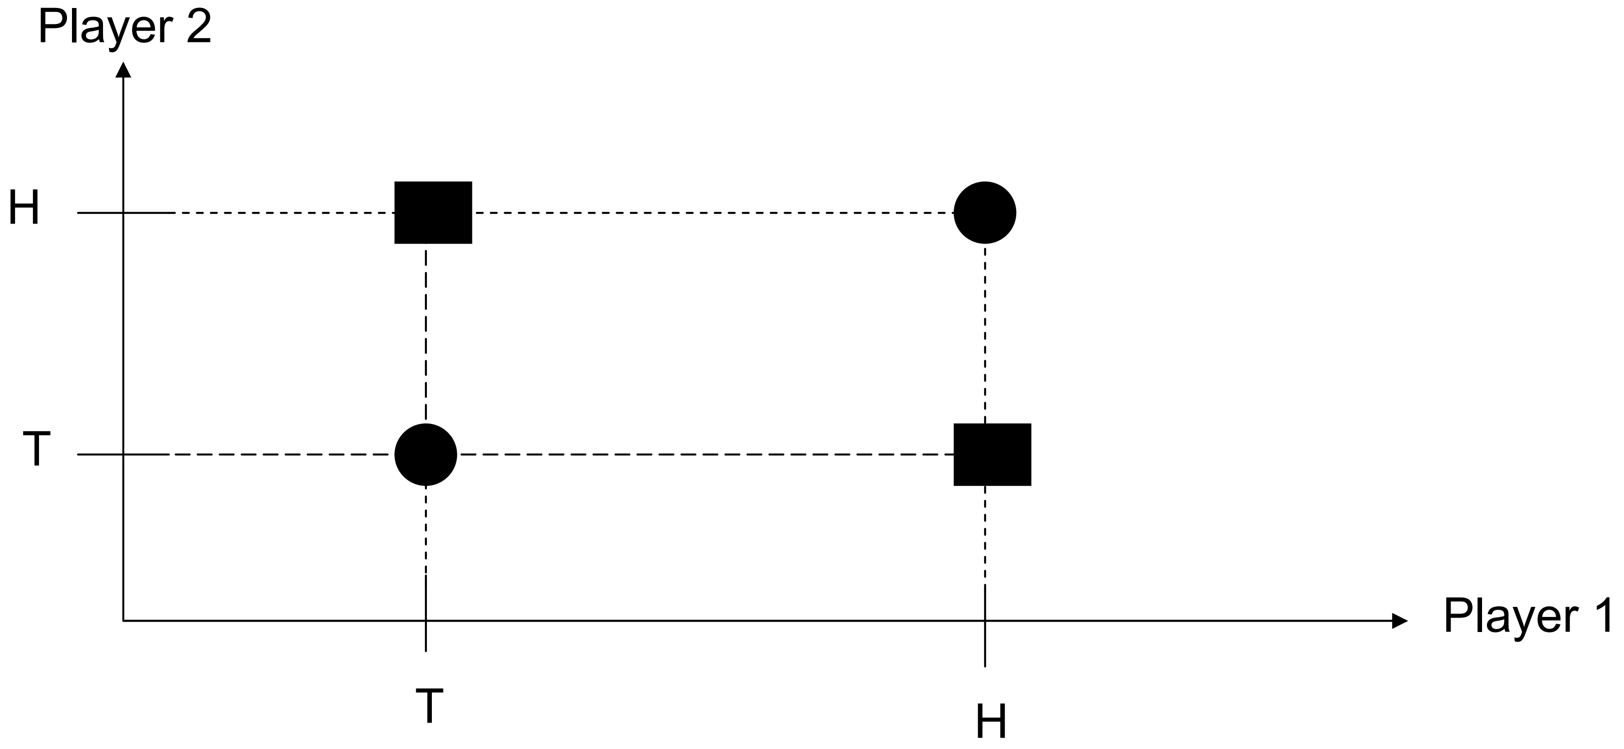

Best responses: Matching pennies Mixed Strategies

- Player 1 believes that player 2 will play ”Head” with proba q and ”Tail” with proba ”1-q”.
- Player 2 believes that player 1 will play ”Head” with proba p and ”Tail” with proba ”1-p”.

Best responses: Matching pennies Mixed Strategies

$$EU_1(H) = q \times 1 + (1-q)(-1) = 2q - 1$$

$$EU_1(T) = q(-1) + (1-q) \times 1 = 1 - 2q$$

- Thus player 1 will always plays H (i.e. *p*=1) iff

$$EU_1(H) > EU_1(T) \Leftrightarrow q > 0.5$$

- Thus player 2 will always plays T (i.e. *p*=0) iff

$$EU_1(H) < EU_1(T) \Leftrightarrow q < 0.5$$

Otherwise she is indifferent between the two strategies (i.e. for *q*=0.5)

Best responses: Matching pennies Mixed Strategies

- Best-reply function of player 1:

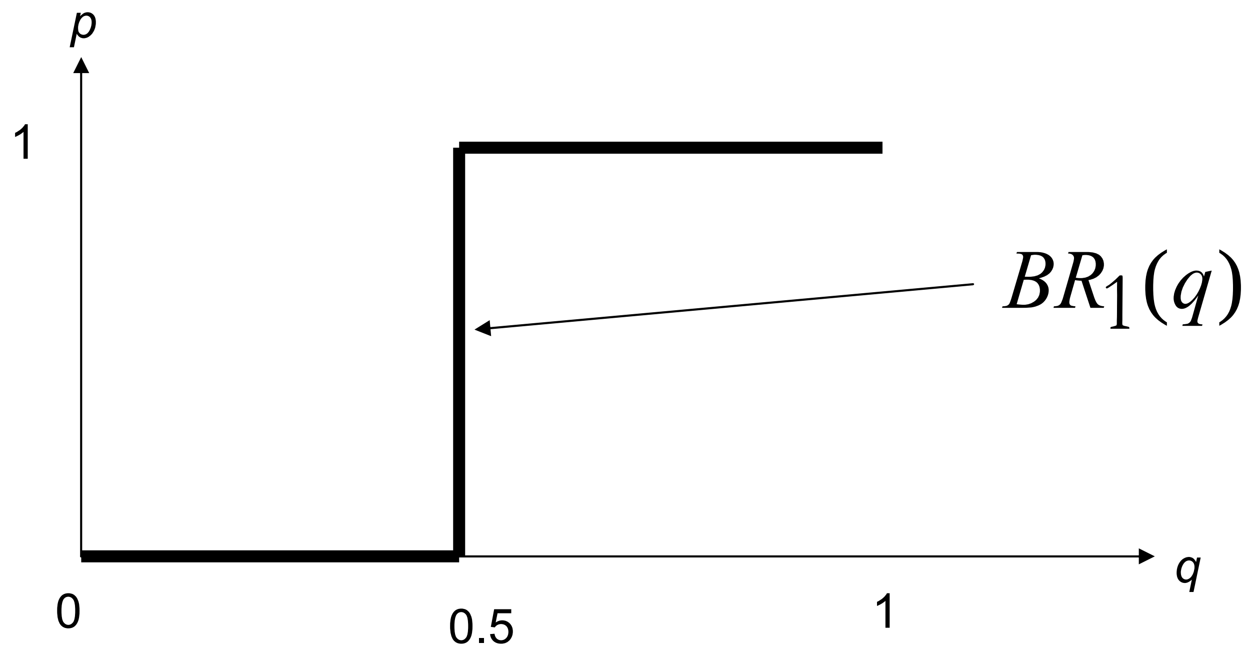

Best responses: Matching pennies Mixed Strategies

$$EU_2(H) = p(-1) + (1 - p) \times 1 = 1 - 2p$$

$$EU_2(T) = p \times 1 + (1 - p)(-1) = 2p - 1$$

- Thus player 1 will always plays H (i.e. *q*=1) iff

$$EU_2(H) > EU_2(T) \Leftrightarrow p < 0.5$$

- Thus player 2 will always plays T (i.e. *q*=0) iff

$$EU_2(H) < EU_2(T) \Leftrightarrow p > 0.5$$

Otherwise she is indifferent between the two strategies (i.e. for *p*=0.5)

Best responses: Matching pennies Mixed Strategies

- Best-reply function of player 2:

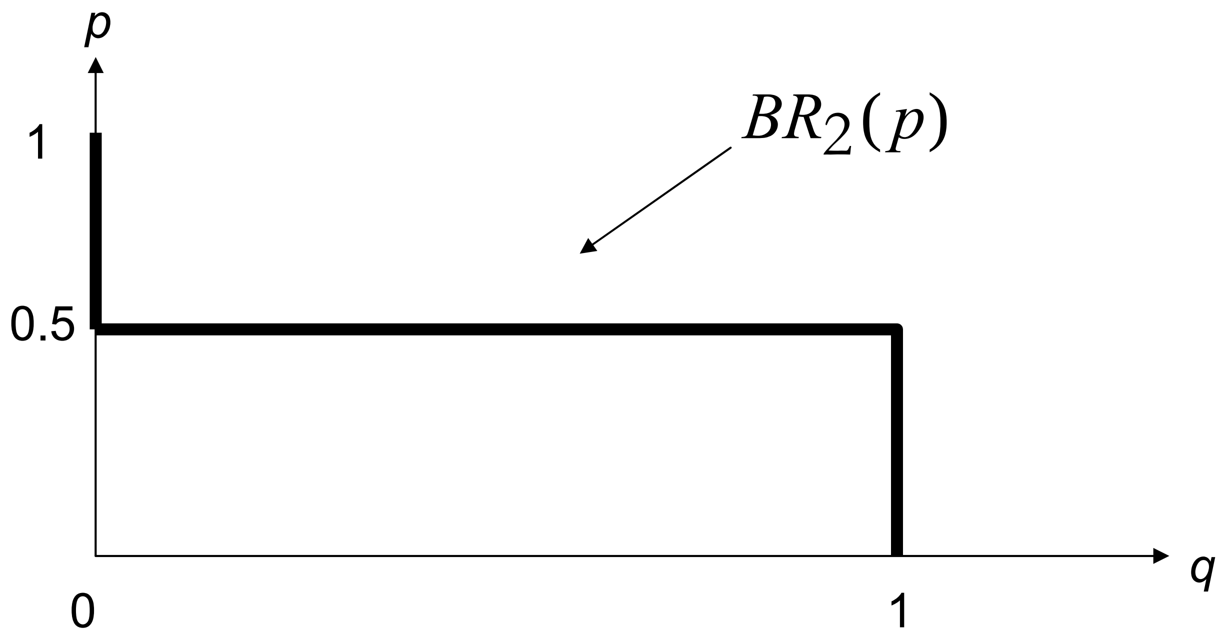

Nash equilibrium: Fixed point NE: (1/2,1/2)

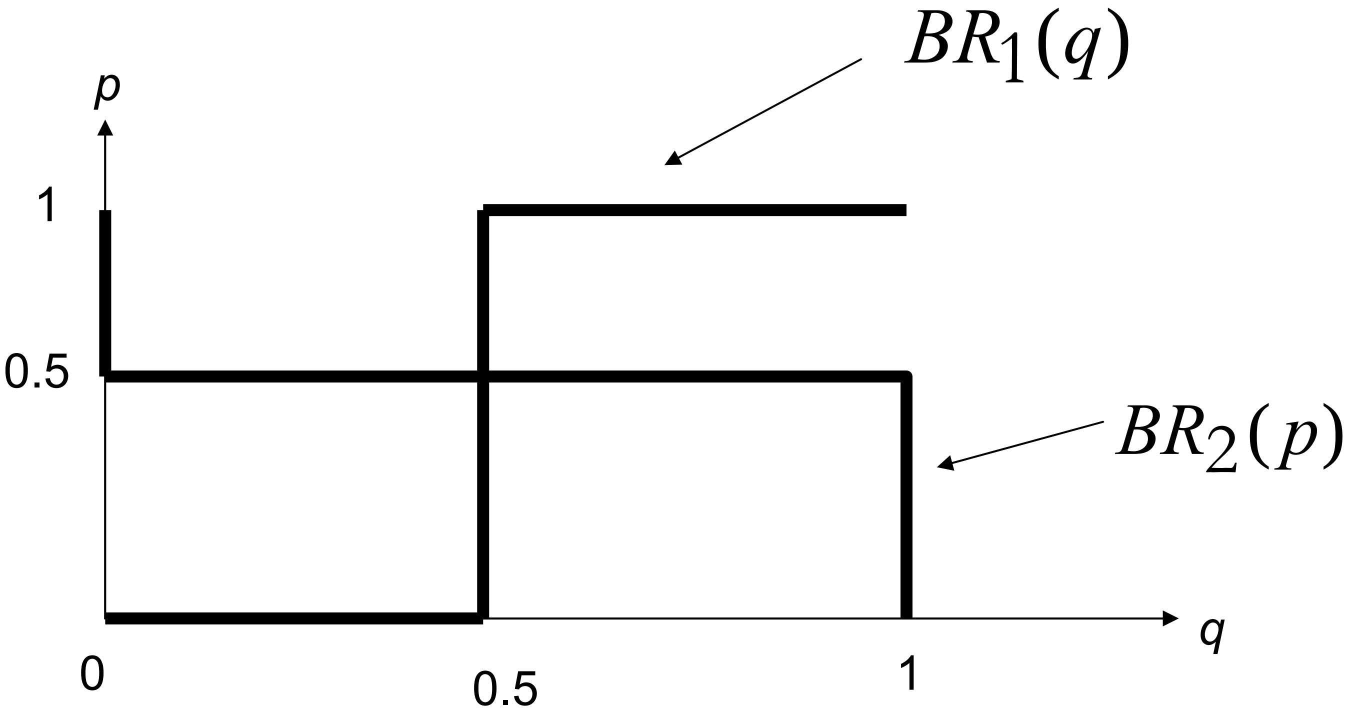

What if it’s difficult to draw the best responses?

- Restating the expected utility of the mixed strategy profile α:

$$U_i(\alpha) = \sum_{a_i \in A_i} \alpha_i(a_i) E_i(a_i, \alpha_{-i}),$$

where A_i is player i's action set,

$E_i(a_i, \alpha_{-i})$ is the expected utility

from playing the pure strategy a_i, given

that your opponents play the mixed strategy α_{-i}.

Characterization of a mixed strategy NE of finite game

A mixed strategy profile α^* is a NE iff for each player i;

- Given α^*_{-i}, the expected payoff $E(a_i,\alpha^*_{-i})$ to every action a_i to which α_i^* assigns positive probability is the same, that is, E*.
- Given α^*_{-i}, the expected payoff $E(a_i,\alpha^*_{-i})$, to every action a_i to which α_i^* assigns zero probability is at most the expected payoff to any action to which α_i^* assigns positive probability, i.e. E*.

Each player's expected payoff in an equilibrium is her expected payoff to any of her actions that she uses with positive probability. That is, in equilibrium, given α^*_{-i} *individual i is indifferent between any of her actions that she plays with a positive probability.*

Pure and mixed equilibria

- A pure NE is simply a special case of a mixed NE where the assigned probability of some action is one.
- If we have, say two NE in a finite strategic game with vNM preferences where we do not allow players to randomize, will these equlibria survive if we now allow players to randomize?

Formation of players' beliefs

- Idealized situation: learn from others in your population.
- What if no experience in the population? – elimination of dominated actions.
- Learning.
  - Best response dynamics.
    - Beliefs about pure strategies.
  - Fictitious play.
    - Beliefs about mixed strategies.

Soccer penalties – the Neeskens effect

German player Berti Vogts brings Johan Cruijff down in the world cup final of 1974.

Johan Neeskens shoots the penalty in the middle of the goal while the goalkeeper, Maier, dives to the right:1 – 0.

Simple penalty game pre 1974

- Until 1974 kickers either shot to the right or left…
- Goal keepers understood this.
- What is the NE?
- What is the probability to score (assuming that neither kicker nor goal keeper will fail in their preferred strategy profiles)?

			Kicker	
			L	*R*
			q	*(1-q)*
Goal keeper	*L*	*p*	1, 0	0, 1
	R	*(1-p)*	0, 1	1, 0

Subtle institutional evolution through behavioral innovation (within the rules)

- Neeskens' innovation changed the institutional setting of the penalty game.
- What is the NE?
- What is the probability to score (assuming that neither kicker nor goal keeper will fail in their preferred strategy profiles)?

			Kicker		
			L	*M*	*R*
			q_1	q_2	q_3
	L	p_1	1, 0	0, 1	0, 1
Goal keeper	*M*	p_2	0, 1	1, 0	0, 1
	R	p_3	0, 1	0, 1	1, 0

New rules, new game: L, M, R.

- Show that there exists no NE in pure strategies and a unique NE in mixed strategies which is:

$$(p_1, p_2, p_3) = (1/3, 1/3, 1/3)$$

$$(q_1, q_2, q_3) = (1/3, 1/3, 1/3)$$

Fundamental Principle of Mixed-Strategy NE

- In a mixed-strategy Nash equilibrium **each player chooses a mixed strategy that makes her opponent indifferent** between playing any[*] of her pure strategies or any mixture of her pure strategies.
- When all players are making the others indifferent, every player is indifferent between playing her current mixture and doing anything else.
- Thus no player has an incentive to deviate (*i.e.* they are all best responding), so this is a Nash equilibrium.

[*] That is, any of the pure strategies that she plays with positive probability in her mixed strategy.

Finding Mixed-Strategy NE

- With small and simple games—such as when there are two players and each has two strategies—the best response functions of each player can be graphed. Any intersection points of the best-response functions are Nash equilibria.
- For larger games—with more than two players or when two players have more than two strategies each—a three-(or *n*-)dimensional graph would be required.
- For either small or large games, the mixed-strategy Nash equilibrium can be found algebraically.

Example: Penalty Kick (1)

- Consider the example of a penalty kick introduced in the last lecture:

Kicker

			q1	q2	q3
			L	M	R
	p1	L	$\underline{1}$, 0	0, $\underline{1}$	0, $\underline{1}$
Goal Keeper	p2	M	0, $\underline{1}$	$\underline{1}$, 0	0, $\underline{1}$
	p3	R	0, $\underline{1}$	0, $\underline{1}$	$\underline{1}$, 0

Underlining best responses, we see there are no pure-strategy NE. Also, we see that no strategies are dominated for either player.

Example: Penalty Kick (2)

- First, since we know $p_1 + p_2 + p_3 = 1$, we can instead rewrite p_3 as $1 - p_1 - p_2$.
- Next, as a function of p_1 and p_2, find the Kicker's expected payoff (or expected utility) for each of his strategies:

$$EU_K(L) = p_1(0) + p_2(1) + (1 - p_1 - p_2)(1) = 1 - p_1$$

$$EU_K(M) = p_1(1) + p_2(0) + (1 - p_1 - p_2)(1) = 1 - p_2$$

$$EU_K(R) = p_1(1) + p_2(1) + (1 - p_1 - p_2)(0) = p_1 + p_2$$

- The Goal Keeper sets p_1 and p_2 such that the Kicker is indifferent between playing any of his strategies.
- For the Kicker to be indifferent, it must be the case that:

$$EU_K(L) = EU_K(M) = EU_K(K) \Rightarrow 1 - p_1 = 1 - p_2 = p_1 + p_2$$

Example: Penalty Kick (3)

- To make the Kicker indifferent, the Goal Keeper chooses p_1 and p_2 such that:

  $1 – p_1 = 1 – p_2 = p_1 + p_2$

- Now use algebra to solve for p_1 and p_2:

  By the equality of the first two terms we have:

  $1 – p_1 = 1 – p_2 \Rightarrow p_1 = p_2$

  Substituting p_1 for p_2 in the third term, by the equality of the first and third terms we have:

  $1 – p_1 = p_1 + p_1 \Rightarrow 1 = 3p_1 \Rightarrow p_1 = 1/3$

  Since $p_1 = p_2$, $p_2 = 1/3$. Also, $(1 – p_1 – p_2) = 1/3$.

- In equilibrium the Goal Keeper plays L with probability 1/3, M with probability 1/3, and R with probability 1/3.

Example: Penalty Kick (4)

- Now find the Kicker's equilibrium mixture.
- First, rewrite q_3 as $1 – q_1 – q_2$.
- Next, as a function of q_1 and q_2, find the Goal Keeper's expected payoff (or expected utility) for each of his strategies:

$$EU_G(L) = q_1(1) + q_2(0) + (1 – q_1 – q_2)(0) = q_1$$

$$EU_G(M) = q_1(0) + q_2(1) + (1 – q_1 – q_2)(0) = q_2$$

$$EU_G(R) = q_1(0) + q_2(0) + (1 – q_1 – q_2)(1) = 1 – q_1 – q_2$$

- The Kicker sets q_1 and q_2 such that the Goal Keeper is indifferent between playing any of his strategies.
- For the Goal Keeper to be indifferent, it must be the case that:

$$EU_G(L) = EU_G(M) = EU_G(K) \Rightarrow q_1 = q_2 = 1 – q_1 – q_2$$

Example: Penalty Kick (5)

- To make the Goal Keeper indifferent, the Kicker chooses q_1 and q_2 such that:

  $q_1 = q_2 = 1 – q_1 – q_2$

  Substituting q_1 for q_2 in the third term, by the equality of the first and third terms we have:

  $q_1 = 1 – q_1 – q_1 \Rightarrow 3q_1 = 1 \Rightarrow q_1 = 1/3$

  Since $q_1 = q_2$, $q_2 = 1/3$. Also, $(1 – q_1 – q_2) = 1/3$.

- In equilibrium the Kicker plays L with probability 1/3, M with probability 1/3, and R with probability 1/3.
- The mixed-strategy Nash equilibrium of the Penalty Kick game is:

  (1/3 L + 1/3 M + 1/3 R, 1/3 L + 1/3 M + 1/3 R)

Gambit: a Shortcut to Finding NE

- The open-source software Gambit can quickly and reliably find pure-strategy and mixed-strategy Nash equilibria for even large games.
- You can download Gambit for free at:

  http://gambit.sourceforge.net/download.html

- Remember to use Gambit wisely.
  - Gambit can be a valuable tool in helping you gain an intuition for how mixed-strategy Nash equilibria work.
  - Gambit cannot help you on an exam. If you don't learn how to find all of a game's Nash equilibria by hand, you will be in deep, deep trouble during the midterm or final.

Example: Tennis (1)

- Consider the following game:

			Navratilova	
			q	1 - q
			DL	CC
Evert	p	DL	50, $\underline{50}$	$\underline{80}$, 20
	1 - p	CC	$\underline{90}$, 10	20, $\underline{80}$

Best response analysis shows that there is no equilibrium in pure strategies.

- To arrive at a mixed-strategy Nash equilibrium, each player chooses her mixture to make the other player indifferent.

Example: Tennis (2)

- To find p, Evert considers the expected utility that Navratilova gets from playing each of her strategies as a function of p:

$EU_N(DL) = 50p + 10(1 – p) = 40p + 10$

$EU_N(CC) = 20p + 80(1 – p) = -60p + 80$

- Evert wants Navratilova to be indifferent between playing DL or CC, so she sets $EU_N(DL) = EU_N(CC)$; that is:

$40p + 10 = -60p + 80$

$\Rightarrow \quad 100p = 70$

$\Rightarrow \quad p = 0.7$

			Navratilova	
			q	1 - q
			DL	CC
Evert	p	DL	50, 50	80, 20
	1 - p	CC	90, 10	20, 80

Example: Tennis (3)

- To find q, Navratilova considers the expected utility that Evert gets from playing each of her strategies as a function of q:

  $EU_E(DL) = 50q + 80(1 - q) = -30q + 80$

  $EU_E(CC) = 90q + 20(1 - q) = 70q + 20$

- Navratilova wants Evert to be indifferent between playing DL or CC, so she sets $EU_E(DL) = EU_E(CC)$; that is:

  $-30q + 80 = 70q + 20$

  $\Rightarrow \quad 60 = 100q$

  $\Rightarrow \quad q = 0.6$

			Navratilova	
			q	1 - q
			DL	CC
Evert	p	DL	50, 50	80, 20
	1 - p	CC	90, 10	20, 80

Example: Tennis (4)

- The mixed-strategy Nash equilibrium occurs at p = 0.7, q = 0.6, which could also be written (0.7 DL + 0.3 CC, 0.6 DL + 0.4 CC).
- The best-response graph is as follows:

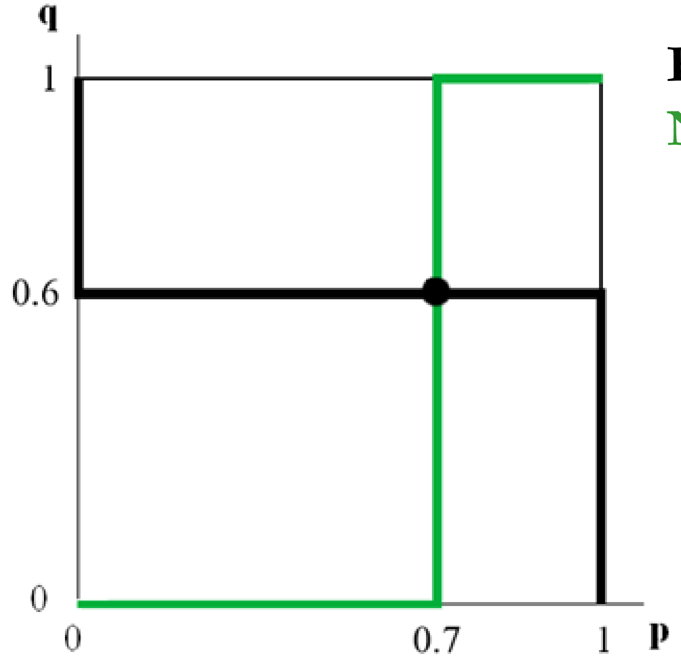

Evert's best response is in black

Navratilova's best response is in green

			Navratilova	
			q	1 - q
			DL	CC
Evert	p	DL	50, 50	80, 20
	1 - p	CC	90, 10	20, 80

Example: Chicken (1)

- Consider the following game:

			Dean	
			q	1 - q
			Swerve	Straight
James	p	Swerve	0, 0	$\underline{-1}, \underline{1}$
	1 - p	Straight	$\underline{1}, \underline{-1}$	-2, -2

Best response analysis shows that there are two pure-strategy Nash equilibria:

NE @ (Straight, Swerve)
with payoffs (1, -1)

NE @ (Swerve, Straight)
with payoffs (-1, 1)

- To arrive at a mixed-strategy Nash equilibrium, each player chooses her mixture to make the other player indifferent.

Example: Chicken (2)

- To find p, James considers the expected utility that Dean gets from playing each of his strategies as a function of p:

  $EU_D(\text{Swerve}) = 0p + (-1)(1 – p) = p – 1$

  $EU_D(\text{Straight}) = 1p + (-2)(1 – p) = 3p – 2$

- James wants Dean to be indifferent between playing Swerve or Straight, so he sets $EU_D(\text{Swerve}) = EU_D(\text{Straight})$, that is:

  $p – 1 = 3p – 2$

  $\Rightarrow \quad 1 = 2p$

  $\Rightarrow \quad p = 1/2$

			Dean	
			q	1 - q
			Swerve	Straight
James	p	Swerve	0, 0	-1, 1
	1 - p	Straight	1, -1	-2, -2

Example: Chicken (3)

- To find q, Dean considers the expected utility that James gets from playing each of his strategies as a function of q:

  $EU_J(\text{Swerve}) = 0q + (-1)(1 – q) = q – 1$

  $EU_J(\text{Straight}) = 1q + (-2)(1 – q) = 3q - 2$

- Dean wants James to be indifferent between playing Swerve or Straight, so he sets $EU_J(\text{Swerve}) = EU_J(\text{Straight})$, that is:

  $q – 1 = 3q - 2$

  $\Rightarrow \quad 1 = 2q$

  $\Rightarrow \quad q = 1/2$

			Dean	
			q	1 - q
			Swerve	Straight
James	p	Swerve	0, 0	-1, 1
	1 - p	Straight	1, -1	-2, -2

Example: Chicken (4)

- The mixed-strategy Nash equilibrium occurs at p = 1/2, q = 1/2, or: (1/2 Swerve + 1/2 Straight, 1/2 Swerve + 1/2 Straight).
- The best-response graph is as follows:

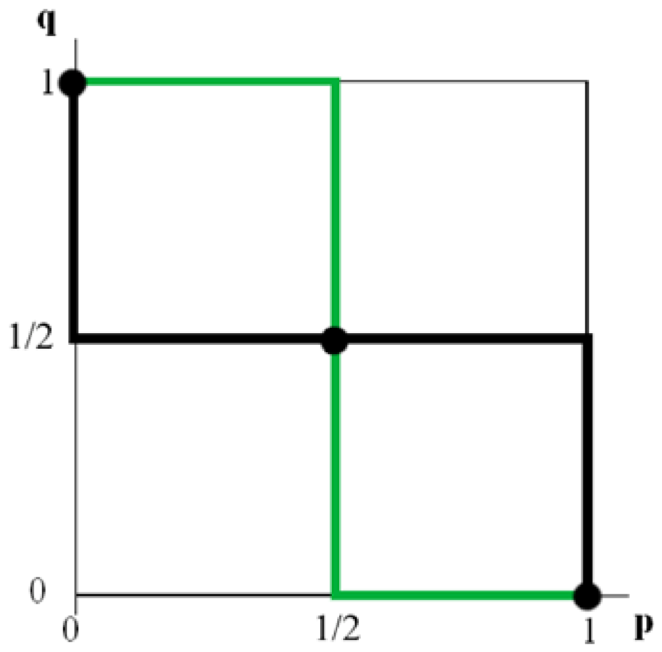

James's best response is in black

Dean's best response is in green

			Dean	
			q	1 - q
			Swerve	Straight
James	p	Swerve	0, 0	-1, 1
	1 - p	Straight	1, -1	-2, -2

Example: Battle of the Sexes (1)

- Consider the following game:

			Sally	
			q	1 - q
			Hockey	Ballet
Harry	p	Hockey	2, 1	0, 0
	1 - p	Ballet	0, 0	1, 2

Best response analysis shows that there are two pure-strategy Nash equilibria:

NE @ (Hockey, Hockey)
with payoffs (2, 1)

NE @ (Ballet, Ballet)
with payoffs (1, 2)

- To arrive at a mixed-strategy Nash equilibrium, each player chooses her mixture to make the other player indifferent.

Example: Battle of the Sexes (2)

- To find p, Harry considers the expected utility that Sally gets from playing each of her strategies as a function of p:

  $EU_S(\text{Hockey}) = 1p + (0)(1 – p) = p$

  $EU_S(\text{Ballet}) = 0p + 2(1 – p) = 2 – 2p$

- Harry wants Sally to be indifferent between playing Hockey or Ballet, so he sets $EU_S(\text{Hockey}) = EU_S(\text{Ballet})$, that is:

  $p = 2 – 2p$

  $\Rightarrow \quad 3p = 2$

  $\Rightarrow \quad p = 2/3$

			Sally	
			q	1 - q
			Hockey	Ballet
Harry	p	Hockey	2, 1	0, 0
	1 - p	Ballet	0, 0	1, 2

Example: Battle of the Sexes (3)

- To find q, Sally considers the expected utility that Harry gets from playing each of his strategies as a function of q:

  $EU_H(\text{Hockey}) = 2q + (0)(1 – q) = 2q$

  $EU_H(\text{Ballet}) = 0q + (1)(1 – q) = 1 – q$

- Sally wants Harry to be indifferent between playing Hockey or Ballet, so she sets $EU_H(\text{Hockey}) = EU_H(\text{Ballet})$, that is:

  $2q = 1 – q$

  $\Rightarrow \quad 3q = 1$

  $\Rightarrow \quad q = 1/3$

			Sally	
			q	1 - q
			Hockey	Ballet
Harry	p	Hockey	2, 1	0, 0
	1 - p	Ballet	0, 0	1, 2

Example: Battle of the Sexes (4)

- The mixed-strategy Nash equilibrium occurs at p = 2/3, q = 1/3, or: (2/3 Hockey + 1/3 Ballet, 1/3 Hockey + 2/3 Ballet).
- The best-response graph is as follows:

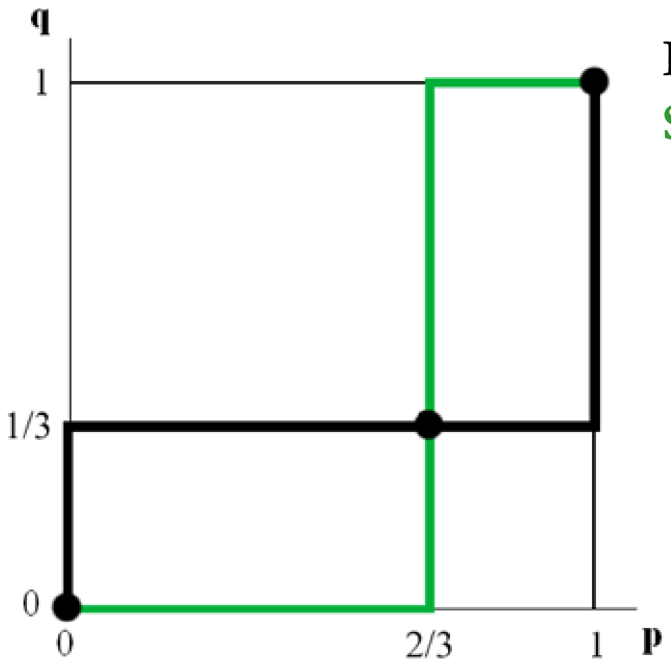

Harry's best response is in black

Sally's best response is in green

			Sally	
			q	1 - q
			Hockey	Ballet
Harry	p	Hockey	2, 1	0, 0
	1 - p	Ballet	0, 0	1, 2

General Formula for p and q for 2x2 Games

- Any two-player simultaneous game where each player has two strategies can be written in the following form:

			Column	
			q	1 - q
			L	R
Row	p	U	*a*, *A*	*b*, *B*
	1 - p	D	*c*, *C*	*d*, *D*

The equilibrium values of p and q are as follows:

$$\mathrm{p} = \frac{D-C}{A-B+D-C}$$

$$\mathrm{q} = \frac{d-b}{a-c+d-b}$$

Remember that p and q cannot be less than zero or greater than one.

Counterintuitive Yet Important Observation

- After all dominated strategies are eliminated, **a player's equilibrium mixture depends only on her opponent's payoffs**, not her own. (Weird, right?)
- Consider this pair of games:

Column

			q	1 - q
			L	R
	p	U	2, -1	4, -2
Row	1 - p	D	3, 0	1, 1

$p = 0.5, \ q = 0.75$

Column

			q	1 - q
			L	R
	p	U	2, -1	4, -2
Row	1 - p	D	999, 0	1, 1

$p = 0.5, \ q = 0.003$

Pure-Strategy NE *Are* Mixed-Strategy NE

- Pure-strategy Nash equilibria are simply mixed-strategy Nash equilibria where each player plays a single strategy with probability one and all other strategies with probability zero.
- Restated:
  - The concept of pure-strategy Nash equilibrium is a special case of the concept of mixed-strategy Nash equilibrium.
  - The concept of mixed-strategy Nash equilibrium is a generalization of the concept of pure-strategy Nash equilibrium.

Applying the General Formula to PD

- What happens when we apply the general formulas for p and q to a game with one *pure*-strategy Nash equilibrium? In short, the formulas usually "break".
- Consider a prisoners' dilemma (PD) game:

			Column	
			q	1 - q
			Deny	Confess
Row	p	Deny	-1, -1	-9, 0
	1 - p	Confess	0, -9	-6, -6

The formulas for p and q give the following:

$$p = \frac{-6-(-9)}{-1-0+(-6)-(-9)} = \frac{3}{2}$$

$$q = \frac{-6-(-9)}{-1-0+(-6)-(-9)} = \frac{3}{2}$$

Finding Mixed-Strategy NE of Larger Games

- To find mixed-strategy Nash equilibria in larger games the principle is still the same: we use algebra to find the mixture(s) that will make the other player(s) indifferent.
- The difference is that the algebra required to find mixed-strategy Nash equilibria becomes more complicated and tedious.
- Remember to eliminate dominated strategies first!
- Knowing how to find the mixed-strategy Nash equilibria of larger games using algebra is an important skill that requires practice.

How Many Nash Equilibria Will There Be?

- All finite games (*i.e.* games with a finite number of players and a finite number of strategies) have at least one Nash equilibrium in mixed strategies (where here the term "mixed strategies" includes pure strategies).
- A finite game may have an infinite number of mixed-strategy Nash equilibria.
- *Almost all* games with a finite number of Nash equilibira have an odd number of them.
  - If there are two pure-strategy Nash equilibria, it is almost certain that there is another equilibrium in mixed strategies.

Example: Game With Infinitely Many NE

- Consider the following game:

			Column	
			q	1 - q
			L	R
Row	p	U	3, 1	2, 2
	1 - p	D	0, 2	2, 3

The best-response graph is as follows:

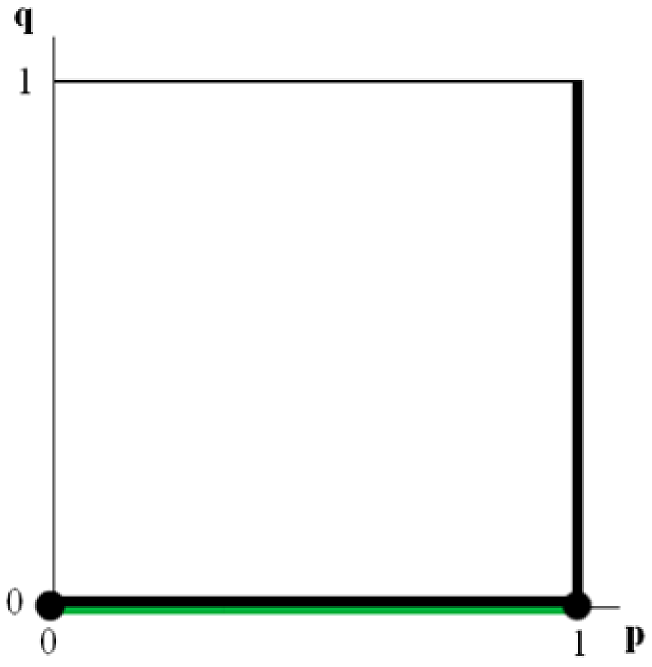

Row's best response is in black

Column's best response is in green

Example: Game With Exactly Two NE

- Consider the following game:

Column

		q	1 - q
		L	R
Row p	U	1, 1	0, 1
Row 1 - p	D	1, 0	1, 1

The best-response graph is as follows:

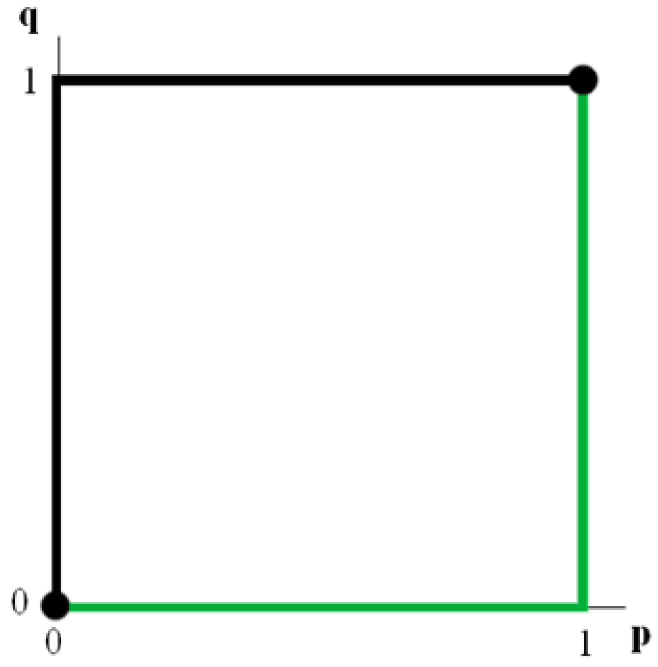

Row's best response is in black

Column's best response is in green

Domination by Mixed Strategies (1)

- A pure strategy that isn't dominated by another pure strategy *might* be dominated by a mixture of other pure strategies if it is never a best response.
- Consider the following game:

		Professor Plum		
		Revolver	Knife	Wrench
Mrs. Peacock	Conservatory	1, 3	2, -2	0, 6
	Ballroom	3, 1	1, 4	5, 0

Revolver is neither dominated by Knife nor Wrench, but it is also never a best response for Column given Row's pure strategies.

Domination by Mixed Strategies (2)

- To see if a pure strategy is dominated by a mixture of other pure strategies, graph the expected payoff of all pure strategies as a function of the opponent's mixture:

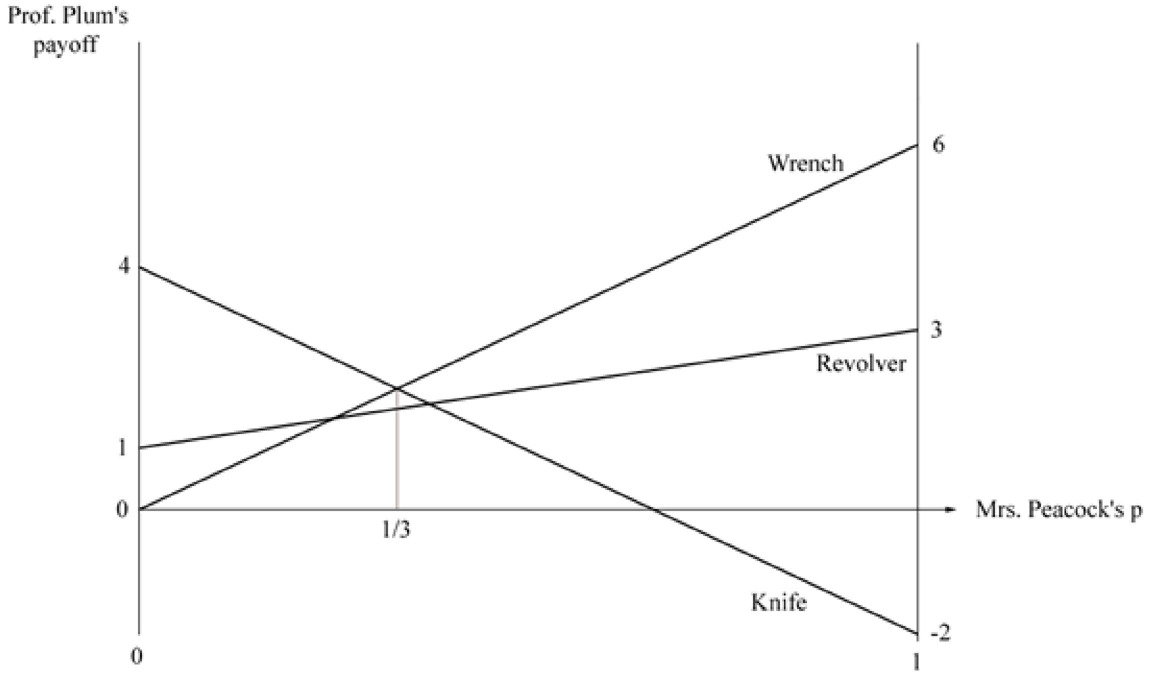

No matter what mixture Mrs. Peacock plays, the expected payoff from Revolver is always less than the expected payoff from either Knife or Wrench.

Revolver is thus strictly dominated by some mixture of Knife and Wrench (many, actually).

Domination by Mixed Strategies (3)

- Consider, however, a slight modification of the game:

Professor Plum

		Revolver	Knife	Wrench
Mrs. Peacock	Conservatory	1, 3	2, -2	0, 6
	Ballroom	3, **2**	1, 4	5, 0

Again, Revolver is neither dominated by Knife nor Wrench, but it is also never a best response for Column given Row’s pure strategies.

Domination by Mixed Strategies (4)

- Again, graph the expected payoffs of each pure strategy as a function of the opponent's mixture:

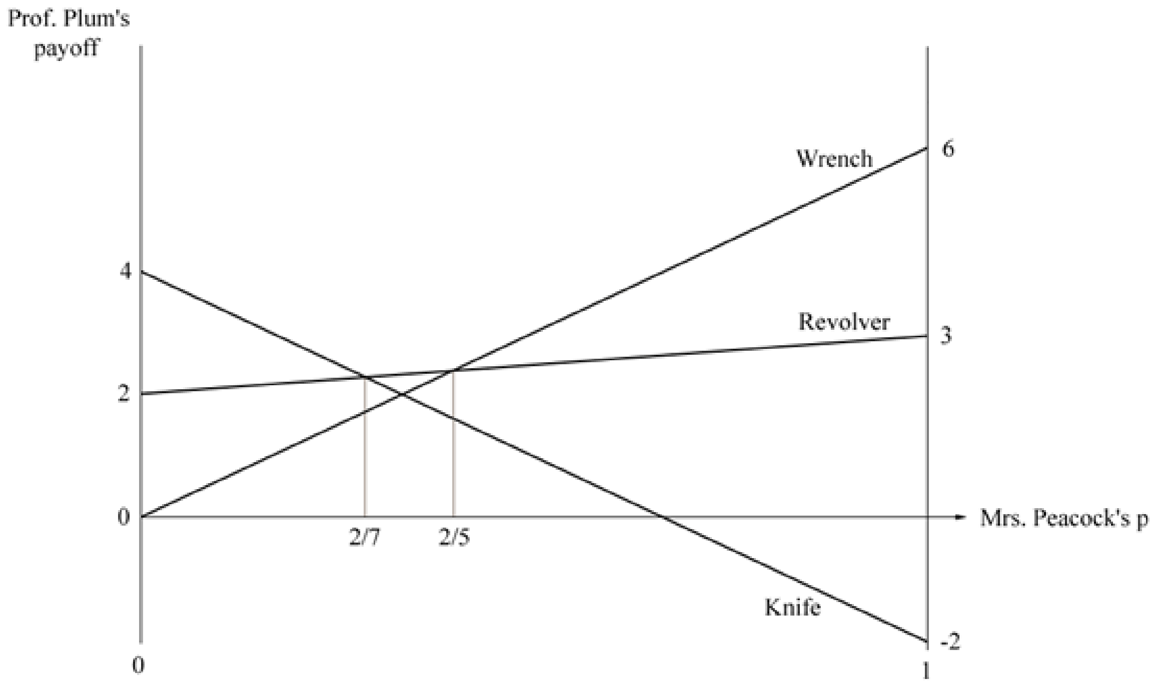

While Revolver is not a best response to either of Mrs. Peacock's pure strategies, Revolver *is* the best response when Mrs. Peacock's mixture is between p = 2/7 and p = 2/5.

In this case Revolver is not dominated by any mixture of Knife and Wrench.

Domination by Mixed Strategies (5)

- Lesson #1:
  - The fact that a pure strategy may be dominated by a mixture of other pure strategies can greatly simplify an *m* x 2 or 2 x *n* game (*i.e.*, a two-player game in which one player has two strategies and the other has more than two strategies).
- Lesson #2:
  - Be very careful when contemplating the elimination of a pure strategy that "looks like" it's dominated by a mixture of other pure strategies.

A given pure strategy s_i may be dominated by a mixed strategy, even though no other pure strategy is able to dominate s_i.

Consider the following matrix:

$$\begin{array}{ccc} & L & R \\ T & (3, \_) & (0, \_) \\ M & (0, \_) & (3, \_) \\ B & (1, \_) & (1, \_) \end{array}$$

Player has the following mixed strategy: $(p, q, 1 - p - q)$.

Let us verify if a mixed strategy $(p, q, 1-p-q)$ dominates a pure strategy T, or M, or B for player 1, given that $0 \leq p \leq 1$ and $0 \leq q \leq 1$, $p + q < 1$.

To check if the mixed strategy $(p, q, 1 - p - q)$ dominates the pure strategy T, we need to solve:

$$\begin{cases} 3p + 0 \times q + 1 \times (1 - p - q) > 3 \\ 0 \times p + 3 \times q + 1 \times (1 - p - q) > 0 \end{cases}$$

$$\Leftrightarrow$$

$$\begin{cases} p > 1 + q/2 \\ p < 2q + 1 \end{cases}$$

The first inequality is impossible since p cannot be greater than 1.

To check if the mixed strategy $(p, q, 1 - p - q)$ dominates the pure strategy M, we need to solve:

$$\begin{cases} 3p + 0 \times q + 1 \times (1 - p - q) > 0 \\ 0 \times p + 3 \times q + 1 \times (1 - p - q) > 3 \end{cases}$$

$$\Leftrightarrow$$

$$\begin{cases} q < 2p + 1 \\ q > 1 + p/2 \end{cases}$$

The second inequality is impossible since q cannot be greater than 1.

To check if the mixed strategy $(p, q, 1-p-q)$ dominates the pure strategy B, we need to solve:

$$\begin{cases} 3p + 0 \times q + 1 \times (1-p-q) > 1 \\ 0 \times p + 3 \times q + 1 \times (1-p-q) > 1 \end{cases}$$

$$\Leftrightarrow$$

$$\begin{cases} \frac{p}{q} > \frac{1}{2} \\ \frac{p}{q} < 2 \end{cases}$$

$$\Leftrightarrow \frac{1}{2} < \frac{p}{q} < 2$$

This is possible. For example satisfied for $(p, q, 1-p-q) = (1/2, 1/2, 0)$ since $\frac{p}{q} = 1$.

As a result, the mixed strategy $(1/2, 1/2, 0)$ strictly dominates the pure strategy B.

Rationalizability and Iterated Dominance

- Consider the following game:

1/2	X	Y	Z
A	3,3	0,5	0,4
B	0,0	3,1	1,2

Rationalizability and Iterated Dominance

- Player 1 has strategies that are dominated.
- Thus player 1 can play A or B depending on her beliefs.
- Denote by p proba that 2 play X (belief) and q proba that 2 play Y, and thus 1-p-q proba that 2 play Z.
- This makes sense if $p \geq 0$, $q \geq 0$, $p+q \leq 1$.

Rationalizability and Iterated Dominance

- Given these beliefs,

EU of playing A: $3p$

EU of playing B: $1 - p + 2q$

- Thus, play A if: $4p > 1+2q$
- Play B if: $4p < 1+2q$

1/2	X	Y	Z
A	3,3	0,5	0,4
B	0,0	3,1	1,2

Rationalizability and Iterated Dominance

- Are these beliefs correct?
- Suppose it is *common knowledge* for the 2 players that both players are *rational* and understand exactly the game.
- Can player 1 rationalize playing strategy A against a rational opponent?
- Answer: NO

Rationalizability and Iterated Dominance

- Think of player 2.
- One can see that she will never plays X (st. dom. by Y and Z!).
- There is no belief that she could have about player 1's strategy that would cause her to play X.
- Knowing this, player 1 should assign proba zero to strategy X, i.e. Player 1's belief should be $p = 0$.

1/2	X	Y	Z
A	3,3	0,5	0,4
B	0,0	3,1	1,2

Rationalizability and Iterated Dominance

1/2	~~X~~	Y	Z
A	~~3,3~~	0,5	0,4
B	~~0,0~~	3,1	1,2

Rationalizability and Iterated Dominance

- In this *new game*, player 1 has two strategies A and B and player 2 has two strategies Y and Z.
- Knowing that player 2 will never play X, player 1's rational strategy is to play B.
- In terms of beliefs, if $p = 0$, then it is impossible that:

$4p \geq 1+2q$

1/2	~~X~~	Y	Z
A	~~3,3~~	0,5	0,4
B	~~0,0~~	3,1	1,2

Rationalizability and Iterated Dominance

- Now, player 2 knows that that player 1 knows that player 2 will not play X.
- Thus, recognizing that player 1 is rational, she deduces that player 1 will not play A.
- She puts proba 1 that player 1 plays A.
- Thus her best response is to play Z since 2 > 1.

1/2	~~X~~	Y	Z
~~A~~	~~3,3~~	~~0,5~~	~~0,4~~
B	~~0,0~~	3,1	1,2

Rationalizability and Iterated Dominance

- Thus, with rational players and common knowledge of rationality, the only outcome that can arise is (B,Z).
- This is called *iterative removal of strictly dominated strategies* or **iterated dominance** for short.

1/2	~~X~~	~~Y~~	Z
~~A~~	~~3,3~~	~~0,5~~	~~0,4~~
B	~~0,0~~	~~3,1~~	1,2

Rationalizability and Iterated Dominance

- General idea: Each player should play a best response to her belief.
- The belief *rationalizes* playing the strategy.
- Each player should assign positive probability only to strategies of other players that can be rationalized.
- Thus, the set of strategies that survive iterated dominance is called the **rationalizable strategies**.
- Rationalizability requires that the players’ beliefs and behavior be consistent with common knowledge of rationality.

Exercise

- Consider the following matrix:

1 / 2	L	C	R
U	5,1	0,4	1,0
M	3,1	0,0	3,5
D	3,3	4,4	2,5

- Determine the set of **rationalizable strategies** of this game (you should look at strategies that are strictly dominated by both pure and mixed strategies).

Answer

- If we look at the matrix, once can see that player 2’s strategy L is strictly dominated by the mixed strategy (0,1/2,1/2), i.e. The strategy that puts proba ½ on both C and R.

1 / 2	~~L~~	C	R
U	~~5,1~~	0,4	1,0
M	~~3,1~~	0,0	3,5
D	~~3,3~~	4,4	2,5

- In this new game,

1 / 2	~~L~~	C	R
~~U~~	~~5,1~~	~~0,4~~	~~1,0~~
M	~~3,1~~	0,0	3,5
D	~~3,3~~	4,4	2,5

- In this new game,

1 / 2	~~L~~	~~C~~	R
~~U~~	~~5,1~~	~~0,4~~	~~1,0~~
M	~~3,1~~	~~0,0~~	3,5
D	~~3,3~~	~~4,4~~	2,5

- In this new game, the set of **rationalizable strategies** is {(M,R)}

1 / 2	~~L~~	~~C~~	R
~~U~~	~~5,1~~	~~0,4~~	~~1,0~~
M	~~3,1~~	~~0,0~~	3,5
~~D~~	~~3,3~~	~~4,4~~	~~2,5~~

Exercise

- Consider a two-player game with the following strategy spaces: $S_1 = [0,5]$ and $S_2 = [0,5]$. Suppose the players' best-response functions are as pictured here:

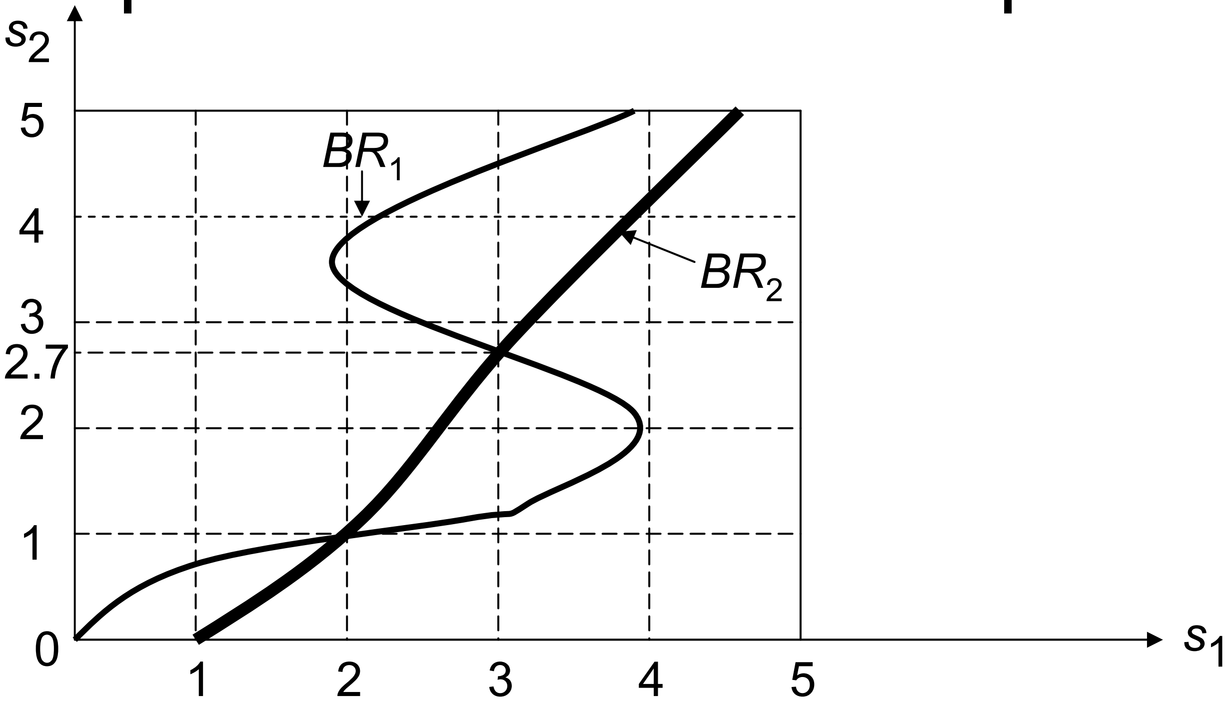

Questions

- (a) Does this game have any Nash equilibria? If so, what are they?
- (b) What is the set of rationalizable strategies for this game?

Answers

- (a) Nash equilibria: (2, 1) and (3, 2.7).
- (b) Set of rationalizable strategies: $[2,3]\times[1,2.7]$

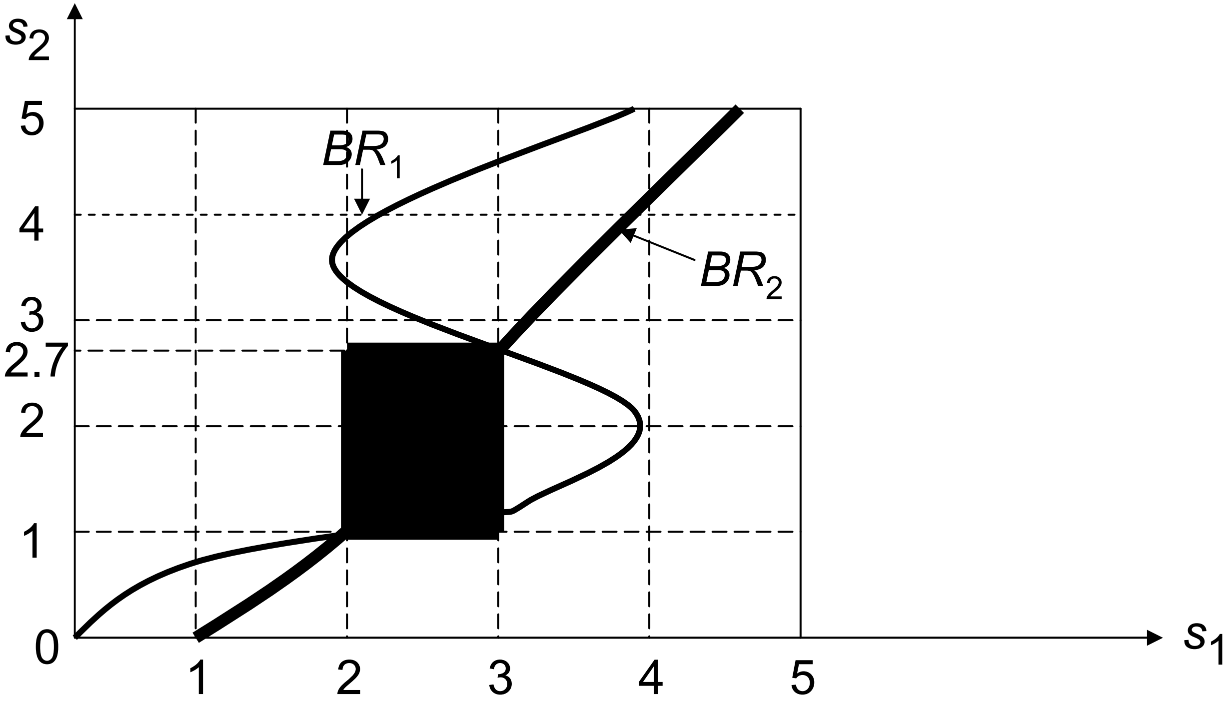

Example: Airport Security

- When you go to a foreign country and go through customs, you can either declare what you’re bringing in or you can not declare.

- A Tourist can Declare or Not Declare. An Inspector can Inspect or Not Inspect.

Example: Airport Security

- I’m the Tourist, and I have $1,000 that I’m trying to bring in.
- If I declare, there is a $270 charge, and I keep $730. If I don’t declare, and customs catches me, I lose all my money and end up with $0. If I don’t declare, and customs does not catch me, I keep all my money $1,000.

- It costs Customs $100 to inspect me. If I declare, then Customs makes $270 because they didn’t inspect me, but I declared anyway and paid the $270 fee.
- If I declare and Customs inspects me, they get my $270 fee, but it cost $100 to inspect me, so Customs has a net gain of $170.
- If I don’t declare, and Customs inspects me, then they make $900. If I don’t declare, and Customs does not inspect me, then Customs gets nothing $0.

Example: Airport Security

		Inspector	
		Inspect	Not
Tourist	Declare	730, 170	730, 270
	Not	0, 900	1000, 0

Example: Airport Security

- **Pure Strategy-** Players either play one strategy or another. You take one action with probability 100%. Here no pure-strategy NE.
- **Mixed Strategy-** Players play strategies with varying probabilities.

Example: Airport Security (Mixed Strategy NE)

- The Tourist will play Declare with probability *d* and Not Declare with probability 1-*d*.
- The Inspector will play Inspect with probability *i* and Not Inspect with probability 1-*i*.

Example: Airport Security (Mixed Strategy NE)

			Inspector	
			i	$1-i$
			Inspect	**Not**
Tourist	d	**Declare**	730, 170	730, 270
	$1-d$	**Not**	0, 900	1000, 0

Example: Airport Security (Mixed Strategy NE)

- Tourist’s Perspective:
- Assume Tourist’s risk preference is the cubed root of the monetary payoff:
- $$U(m) = m^{1/3}$$
- This means that she is **Risk Averse** because the graph of this function curves down.

Example: Airport Security (Mixed Strategy NE)

- So now let's translate the Tourist's monetary payoffs into VNM payoffs.

$$U(m) = 1000^{1/3} = 10$$

$$U(m) = 730^{1/3} = 9(approx)$$

- We insert the Utility values into our chart:

Example: Airport Security (Mixed Strategy NE)

			Inspector	
			i	$1-i$
			Inspect	**Not**
Tourist	d	**Declare**	730, 170 9	730, 270 9
	$1-d$	**Not**	0, 900 0	1000, 0 10

Example: Airport Security (Mixed Strategy NE)

- **U(Declare given i=.9):**
- =U(0)*(0.9) + U(100)*(0.1)
- =9*(0.9) + 9*(0.1)=9
- **U(Not Declare given i=.9):**
- =U(0) *(0.9) + U(100)*(0.1)
- =0*(0.9) + 10*.1

Example: Airport Security (Mixed Strategy NE)

- Because the Utility of Declaring and Not Declaring is the same, the payoff is always going to be 9 whatever the probability of inspection *i*.
- In this case, everything is a best reply *and* everything is a worst reply.

Example: Airport Security (Mixed Strategy NE)

Prob(i)	**U(Declare)**	**U(Not)**	**brT(i)**
.9	9	1	d=1
.7	9	3	d=1
.4	9	6	d=1
.1	**9**	**9**	**$0 \leq d \leq 1$**
.05	9	9.5	d=0

Example: Airport Security (Mixed Strategy NE)

- Best-reply function of the turist:

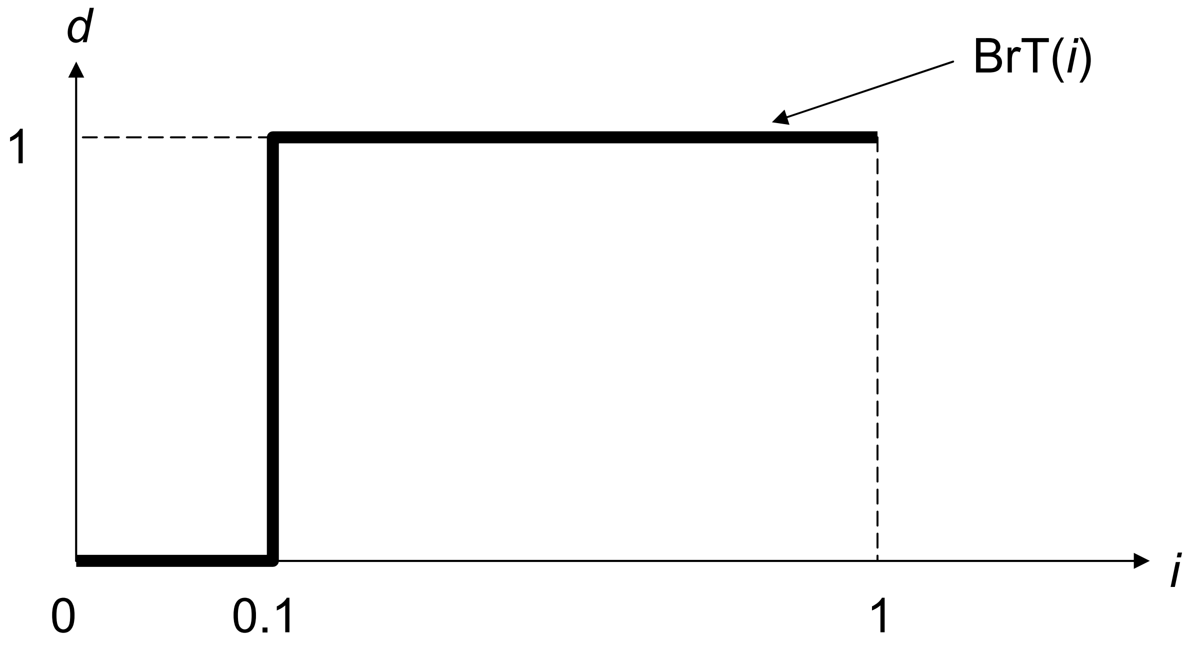

Example: Airport Security (Mixed Strategy NE)

- Same analysis from the Inspector’s Point of View. **Risk Neutral**: U(m)=m
- **U(Inspect, given Declare):**
- =170*d + 900*(1-d)
- **U(Not Inspect, given Not Declare):**
- =270*d + 0*(1-d)

Example: Airport Security (Mixed Strategy NE)

- When is Inspecting going to be a best reply?

- When the payoff for Inspecting is greater than the payoff for Not Inspecting.

Example: Airport Security (Mixed Strategy NE)

- Inspector’s Best Response:
- $170*d + 900*(1-d) > 270*d + 0*(1-d)$
- So now let’s solve for *d*.
- $170*d+900-900*d > 270*\mathrm{d}$
- $900 > 1000*d$
- $9/10 > d$
- So when the *d* is less than 9/10, you want to Inspect for sure (with probability = 1).

Example: Airport Security (Mixed Strategy NE)

- Best-reply function of the Inspector:

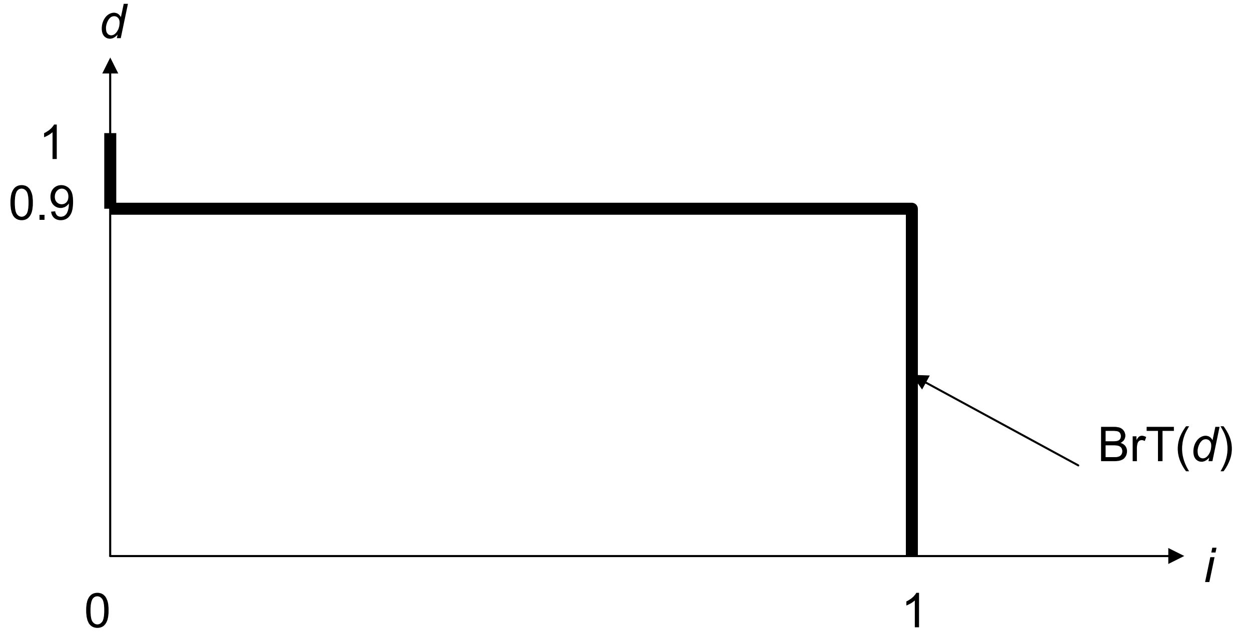

Example: Airport Security (Mixed Strategy NE)

- Mixed-strategy NE: (0.1;0.9)

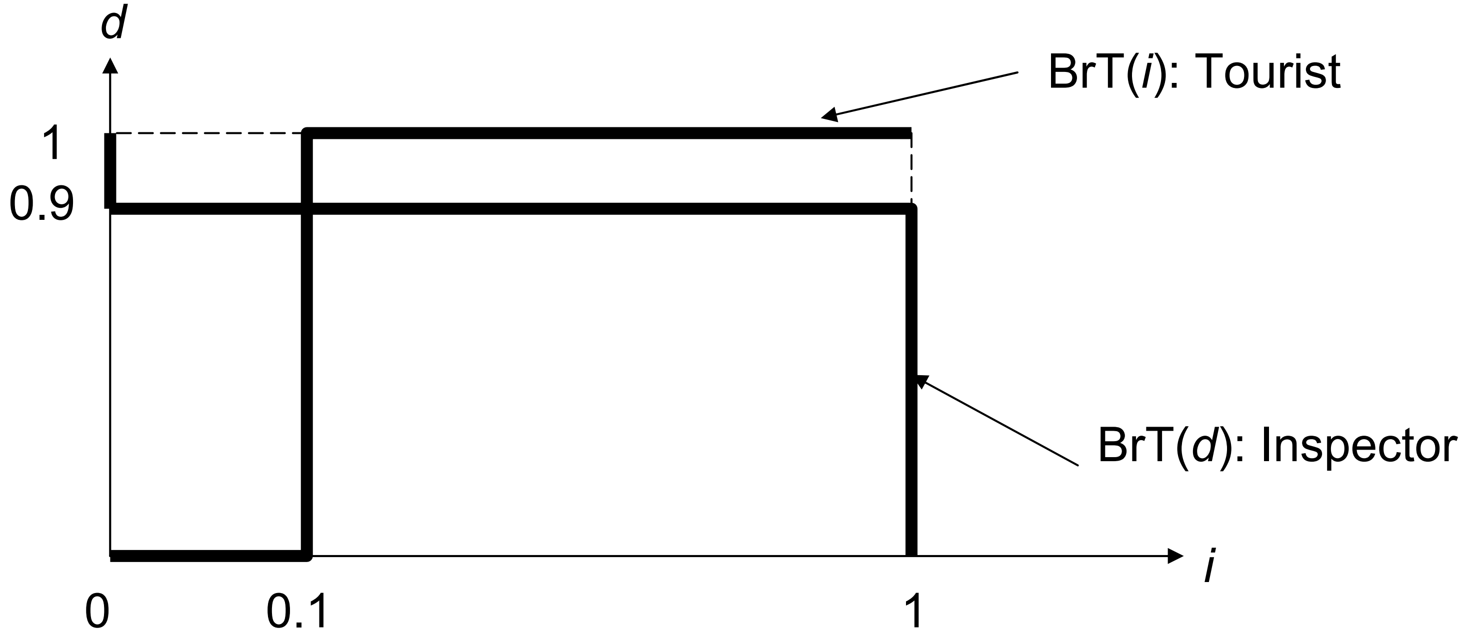

Example: Airport Security (Mixed Strategy NE)

- Where the two lines intersect is the best reply for both players— at (0.1, 0.9), when the Inspector inspects with probability 0.1 and the Tourist declares with probability 0.9.
- This is the Nash Equilibrium for the game.

Exercise: Coordination games (Osborne)

A crime is observed by a group of n people. Each person would like the police to be informed, but prefers that someone else make the phone call.

Specifically, suppose that each person attaches the value v to the police being informed and bears the cost c if she makes the phone call, where $v > c > 0$.

Each person gets a payoff 0 if no one calls (including herself), a payoff $v - c$ if she calls (independently of what the others are doing), and a payoff v if at least one person calls but she does not.

Question (a): Model this situation as a strategic game with vNM preferences.

Answer Question (a)

Players: The n people

Actions: Each player's set of actions is $\{$Call, Don't call$\}$.

Preferences: Each player's preferences are represented by the expected value of a payoff function that assigns 0 to the profile in which no one calls, $v - c$ to any profile in which she calls, and v to any profile in which at least one person calls, but she does not.

Question (b): Show that this game has no symmetric pure Nash equilibrium.

Answer Question (b)

The game has no symmetric pure Nash equilibrium.

If everyone calls, then any person is better off switching to not calling since $v > v - c$.

If no one calls, then any person is better off switching to calling $v - c > 0$.

Question (c): Show that this game has a symmetric mixed strategy equilibrium in which each person calls with positive probability less than one.

Answer Question (c)

In a symmetric mixed strategy equilibrium in which each person calls with positive probability less than one, each person's expected payoff to calling is equal to her expected payoff to not calling.

Each person's payoff to calling is $v - c$;

Each person's payoff to not calling is 0 if no one else calls and v if at least one other person calls

The equilibrium condition (each person's expected payoff to calling is equal to her expected payoff to not calling) is:

$$\begin{aligned} v - c &= 0 \times \Pr\{\text{no one else calls}\} \\ &\quad +v \times \Pr\{\text{at least one other person calls}\} \\ &= v \times \Pr\{\text{at least one other person calls}\} \\ &= v \times (1 - \Pr\{\text{no one else calls}\}) \end{aligned}$$

This is equivalent to:

$$\frac{c}{v} = \Pr\{\text{no one else calls}\}$$

Denote by p the probability with which each person calls.

The probability that *no one else calls* is the probability that every one of the other $n-1$ people does not call, namely $(1-p)^{n-1}$.

Thus the equilibrium condition

$$\frac{c}{v} = \Pr\{\text{no one else calls}\}$$

is now given by:

$$\frac{c}{v} = (1-p)^{n-1}$$

which is equivalent to:

$$p = 1 - \left(\frac{c}{v}\right)^{1/(n-1)}$$

Observe that

$$p > 0 \Leftrightarrow v > c$$

and

$$p < 1 \Leftrightarrow \left(\frac{c}{v}\right)^{1/(n-1)} > 0$$

We have shown that p is between 0 and 1, so we conclude that the game has a unique symmetric mixed strategy equilibrium, in which each person calls with probability $1 - \left(\frac{c}{v}\right)^{1/(n-1)}$.

That is, there is a steady state in which whenever a person is in a group of n people facing the situation modeled by the game, she calls with probability $1 - \left(\frac{c}{v}\right)^{1/(n-1)}$.

Question (d): How does this equilibrium change as the size of the group increases?

Answer Question (d)

We see that as n increases, the probability p that any given person calls decreases.

Indeed,

$$\begin{aligned} a^{x(n)} &= \exp\left[\ln\left(a^{x(n)}\right)\right] \\ &= \exp\left[x(n)\ln a\right] \end{aligned}$$

where $a \equiv c/v$ and $x(n) \equiv 1/(n-1)$. Thus

$$\begin{aligned}
\frac{\partial p}{\partial n} &= \frac{\partial\left[-a^{x(n)}\right]}{\partial n} \\
&= -\frac{\partial a^{x(n)}}{\partial n} \\
&= -\frac{\partial\left\{\exp\left[x(n)\ln a\right]\right\}}{\partial n} \\
&= -\exp\left[x(n)\ln a\right]\times x'(n)\ln a \\
&= \exp\left[\frac{1}{(n-1)}\ln\left(\frac{c}{v}\right)\right]\times\frac{1}{(n-1)^2}\ln\left(\frac{c}{v}\right)
\end{aligned}$$

which is *negative* since ln $\left(\frac{c}{v}\right) < 0$ since $c < v$.

We have shown that p, the probability with which each person calls, is decreasing in the size of the group n.

What about the probability that at least one person calls?

Fix any player i. Then the event “no one calls” is the same as the event “i does not call and no one *other than* i calls”.

Thus

$$\Pr\{\text{no one calls}\} = \Pr\{i \text{ does not call}\} \times \Pr\{\text{no one else calls}\}$$

Now, the probability p that any given person calls decreases as n increases, or equivalently the probability $1-p$ that she does not call increases as n increases.

Further, from the equilibrium condition

$$\frac{c}{v} = \mathsf{Pr}\left\{\text{no one else calls}\right\}$$

we see that $\mathsf{Pr}\left\{\text{no one else calls}\right\}$ is independent of n.

We conclude that the probability that no one calls *increases* as n increases.

That is, the larger the group, the less likely the police are informed of the crime!

REPORTING A CRIME: SOCIAL PSYCHOLOGY AND GAME THEORY

Thirty-eight people witnessed the brutal murder of Catherine (“Kitty”) Genovese over a period of half an hour in New York City in March 1964.

During this period, none of them significantly responded to her screams for help; none even called the police.

Journalists, psychiatrists, sociologists, and others subsequently struggled to understand the witnesses’ inaction.

Some ascribed it to apathy engendered by life in a large city: “Indifference to one’s neighbor and his troubles is a conditioned reflex of life in New York as it is in other big cities” (Rosenthal 1964, 81–82).

Experiments quickly suggested that, contrary to the popular theory, people—even those living in large cities—are not in general apathetic to others' plights.

An experimental subject who is the lone witness of a person in distress is very likely to try to help.

But as the size of the group of witnesses increases, there is a decline not only in the probability that any given one of them offers assistance, but also in the probability that at least one of them offers assistance.

Social psychologists hypothesize that three factors explain these experimental findings.

First, *diffusion of responsibility*: the larger the group, the lower the psychological cost of not helping.

Second, *audience inhibition*: the larger the group, the greater the embarrassment suffered by a helper in case the event turns out to be one in which help is inappropriate.

Third, *social influence*: a person infers the appropriateness of helping from others' behavior, so that in a large group everyone else's lack of intervention leads any given person to think intervention is less likely to be appropriate.

The critical element missing from the socio-psychological analysis is the notion of an *equilibrium*.

Whether any given person intervenes depends on the probability she assigns to some other person's intervening.

In an equilibrium each person must be indifferent between intervening and not intervening, and as we have seen this condition leads inexorably to the conclusion that an increase in group size reduces the probability that at least one person intervenes.