Lectures 10 -11

Repeated Games

14.12 Game Theory

Muhamet Yildiz

Road Map

1. Forward Induction – Examples
2. Finitely Repeated Games with observable actions
    1. Entry-Deterrence/Chain-store paradox
    2. Repeated Prisoners' Dilemma
    3. A general result
    4. When there are multiple equilibria
3. Infinitely repeated games with observable actions
    1. Discounting / Present value
    2. Examples
    3. The Folk Theorem
    4. Repeated Prisoners' Dilemma, revisited –tit for tat
    5. Repeated Cournot oligopoly
4. Infinitely repeated games with unobservable actions

Forward Induction

Strong belief in rationality: At any history of the game, each agent is assumed to be rational if possible. (That is, if there are two strategies s and s’ of a player i that are consistent with a history of play, and if s is strictly dominated but s’ is not, at this history no player j believes that i plays s.)

Burning Money

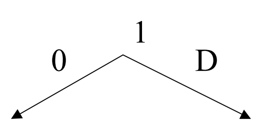

	B	S
B	3,1	.1,.1
S	.1,.1	1,3

	B	S
B	2,1	-.9,.1
S	-.9,.1	0,3

	BB	BS	SB	SS
0B	3,1	3,1	.1,.1	.1,.1
0S	.1,.1	.1,.1	1,3	1,3
DB	2,1	-.9,.1	2,1	-.9,.1
DS	-.9,.1	0,3	-.9,.1	0,3
O	T	H	E	R

Repeated Games

Entry deterrence

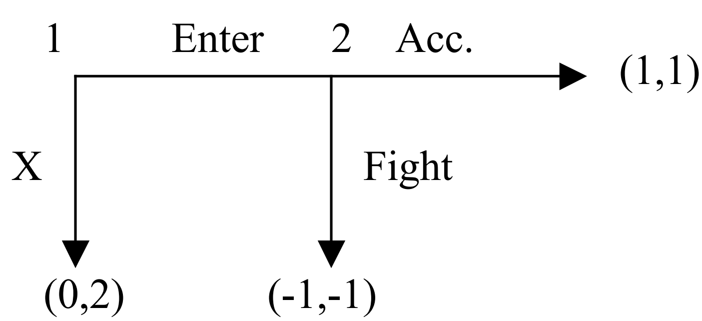

Entry deterrence, repeated twice, many times

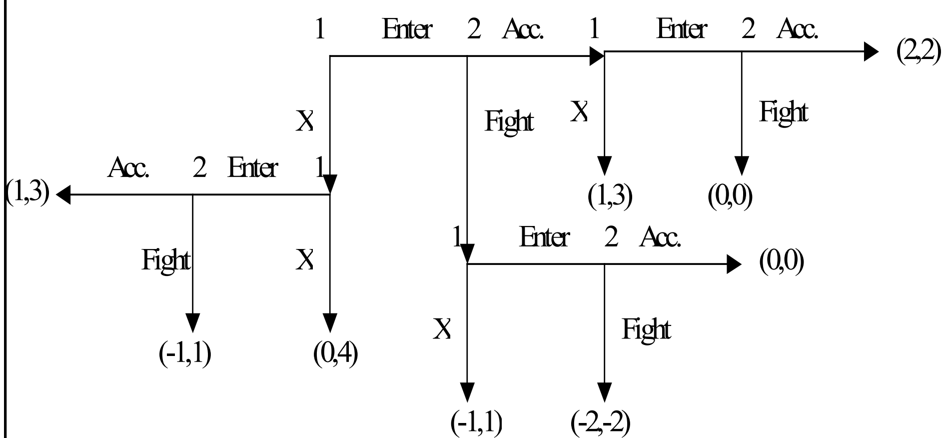

What would happen if repeated n times?

Prisoners’ Dilemma, repeated twice, many times

- Two dates T = {0,1};
- At each date the prisoners’ dilemma is played:

	C	D
C	5,5	0,6
D	6,0	1,1

- At the beginning of 1 players observe the strategies at 0. Payoffs= sum of stage payoffs.

Twice-repeated PD

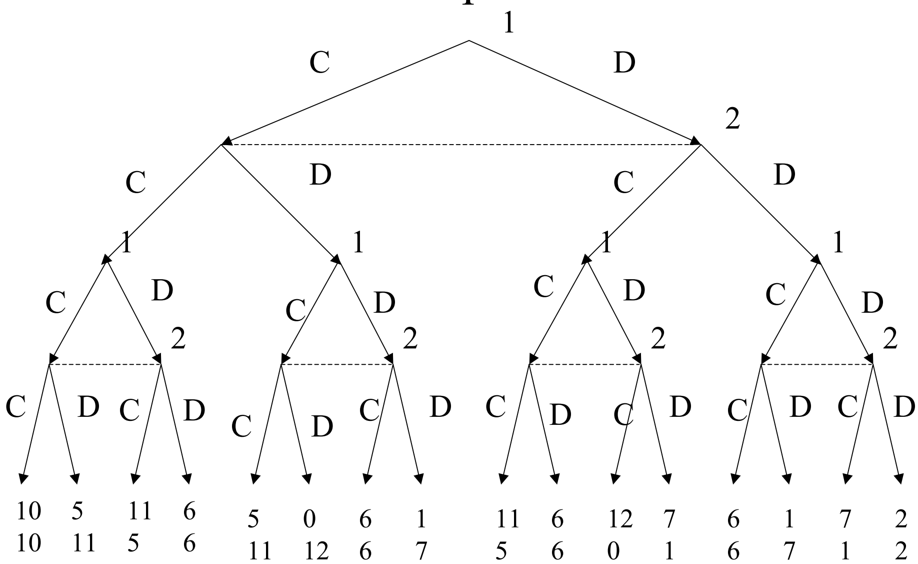

What would happen if T = {0,1,2,…,n}?

A general result

- G = "stage game" = a finite game
- T = {0,1,…,n}
- At each t in T, G is played, and players remember which actions taken before t;
- Payoffs = Sum of payoffs in the stage game.
- Call this game G(T).

Theorem: If G has a unique subgame-perfect equilibrium s*, G(T) has a unique subgame-perfect equilibrium, in which s* is played at each stage.

With multiple equilibria

T = {0,1}

1 \ 2	L	M2	R
T	1,1	5,0	0,0
M1	0,5	4,4	0,0
B	0,0	0,0	3,3

s* =
- At t = 0, each i play Mi;
- At t = 1, play (B,R) if (M1,M2) at t = 0, play (T,L) otherwise.

	L	M2	R
T	2,2	6,1	1,1
M1	1,6	7,7	1,1
B	1,1	1,1	4,4

Infinitely repeated Games with observable actions

- T = {0,1,2,…,t,…}
- G = "stage game" = a finite game
- At each t in T, G is played, and players remember which actions taken before t;
- Payoffs = Discounted sum of payoffs in the stage game.
- Call this game G(T).

Definitions

The *Present Value* of a given payoff stream $\pi = (\pi_0, \pi_1, \ldots, \pi_t, \ldots)$ is

$$PV(\pi;\delta) = \Sigma^{\infty}_{t=1} \delta^t \pi_t = \pi_0 + \delta\pi_1 + \ldots + \delta^t \pi_t + \ldots$$

The *Average Value* of a given payoff stream π is

$$(1-\delta)PV(\pi;\delta) = (1-\delta)\Sigma^{\infty}_{t=1} \delta^t \pi_t$$

The *Present Value* of a given payoff stream π at t is

$$PV_t(\pi;\delta) = \Sigma^{\infty}_{s=t} \delta^{s-t} \pi_s = \pi_t + \delta\pi_{t+1} + \ldots + \delta^s \pi_{t+s} + \ldots$$

Infinite-period entry deterrence

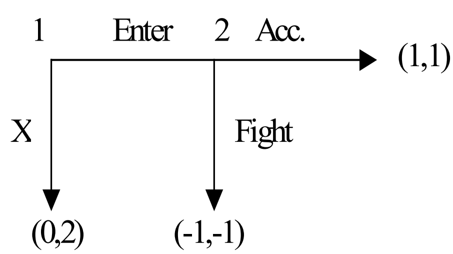

Strategy of Entrant:
Enter iff
Accomodated before.

Strategy of Incumbent:
Accommodate iff
accomodated before.

Incumbent:

- $V(Acc.) = V_A = 1/(1-\delta)$;
- $V(Fight) = V_F = 2/(1-\delta)$;
- Case 1: Accommodated before.
  - Fight => $-1 + \delta V_A$
  - Acc. => $1 + \delta V_A$.
- Case 2: Not Accommodated
  - Fight => $-1 + \delta V_F$
  - Acc. => $1 + \delta V_A$
  - Fight ⇔ $-1 + \delta V_F \geq 1 + \delta V_A$

  ⇔ $V_F - V_A = 1/(1-\delta) \geq 2/\delta$

  ⇔ $\delta \geq 2/3$.

Entrant:

- Accommodated
  - Enter => $1+V_{AE}$
  - X => $0 +V_{AE}$
- Not Acc.
  - Enter =>$-1+V_{FE}$
  - X => $0 +V_{FE}$

Infinitely-repeated PD

	C	D
C	5,5	0,6
D	6,0	1,1

A Grimm Strategy: Defect iff someone defected before.

- $V_D = 1/(1-\delta)$;
- $V_C = 5/(1-\delta) = 5V_D$;
- Defected before (easy)
- Not defected
  - D =>
  - C =>
  - C ⇔

Tit for Tat

- Start with C; thereafter, play what the other player played in the previous round.
- Is (Tit-for-tat,Tit-for-tat) a SPE?
- **Modified:** There are two modes:
  1. Cooperation, when play C, and
  2. Punishment, when play D.

  Start in Cooperation; if any player plays D in Cooperation mode, then switch to Punishment mode for one period and switch back to the Cooperation period next.

Folk Theorem

Definition: A payoff vector $v = (v_1,v_2,\ldots,v_n)$ is feasible iff v is a convex combination of some pure-strategy payoff-vectors, i.e.,

$$v = p_1 u(a^1) + p_2 u(a^2) + \ldots + p_k u(a^k),$$

where $p_1 + p_2 + \ldots + p_k = 1$, and $u(a^j)$ is the payoff vector at strategy profile a^j of the stage game.

Theorem: Let $x = (x_1,x_2,\ldots,x_n)$ be s feasible payoff vector, and $e = (e_1,e_2,\ldots,e_n)$ be a payoff vector at some equilibrium of the stage game such that $x_i > e_i$ for each i. Then, there exist $\underline{\delta} < 1$ and a strategy profile s such that s yields x as the expected average-payoff vector and is a SPE whenever $\delta > \underline{\delta}$.

Folk Theorem in PD

	C	D
C	5,5	0,6
D	6,0	1,1

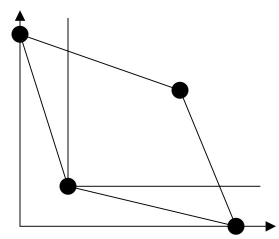

- A SPE with PV (1.1,1.1)?
  - With PV (1.1,5)?
  - With PV (6,0)?
  - With PV (5.9,0.1)?

Infinitely-repeated Cournot oligopoly

- N firms, MC = 0; P = max{1-Q,0};
- Strategy: Each is to produce q = 1/(2n); if any firm defects produce q = 1/(1+n) forever.
- V_C =
- V_D =
- V(D|C) =
- Equilibrium ⇔

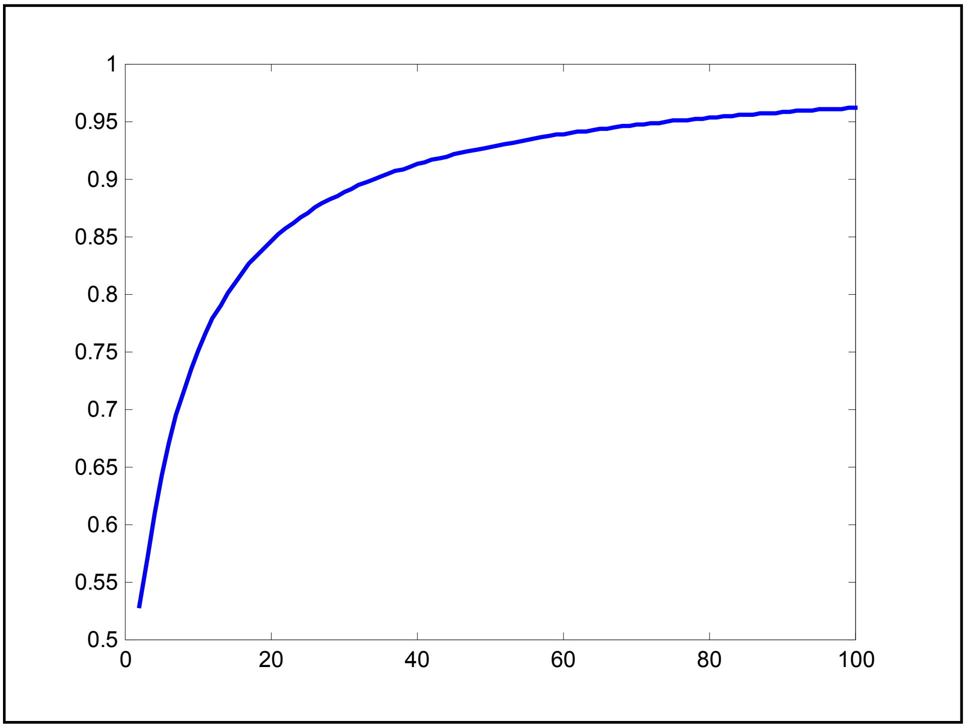

1
0.95
0.9
0.85
0.8
0.75
0.7
0.65
0.6
0.55
0.5
0
20
40
60
80
100

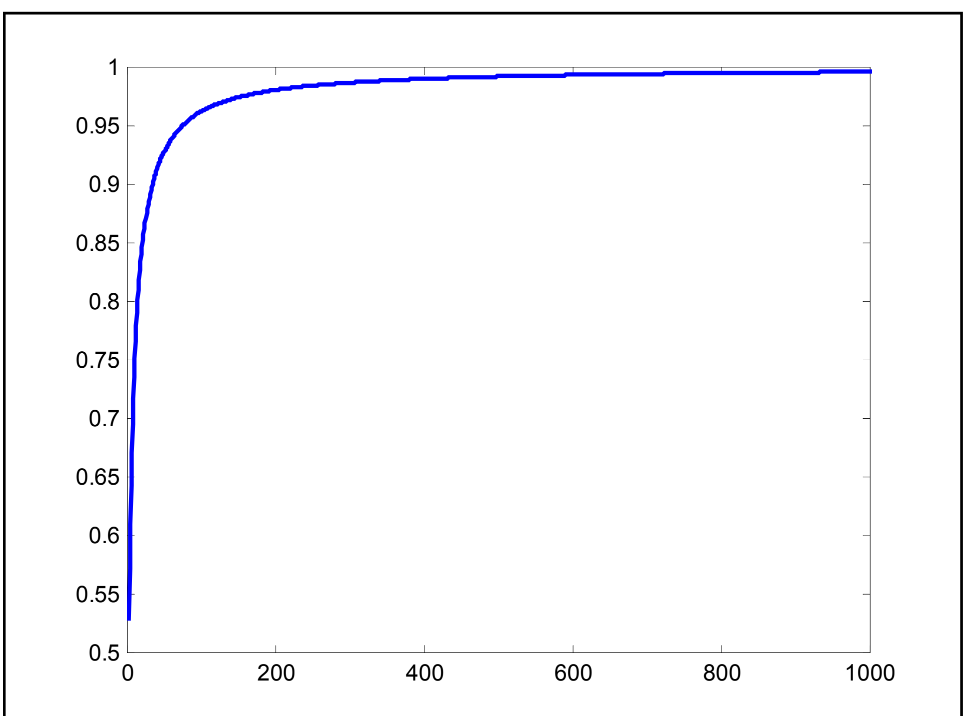

1
0.95
0.9
0.85
0.8
0.75
0.7
0.65
0.6
0.55
0.5
0
200
400
600
800
1000

IRCD (n=2)

- Strategy: Each firm is to produce q*; if any one deviates, each produce 1/(n+1) thereafter.
- $V_C = q^*(1-2q^*)/(1-\delta)$;
- $V_D = 1/(9(1-\delta))$;
- $V_{D|C} = \max q(1-q^*-q) + \delta V_D = (1-q^*)^2/4 + \dfrac{\delta}{9(1-\delta)}$
- Equilibrium iff

$$q^*(1-2q^*) \geq (1-\delta)(1-q^*)^2/4 + \delta/9$$

- $\Leftrightarrow$

$$q^* \geq \frac{9-5\delta}{3(9-\delta)}$$

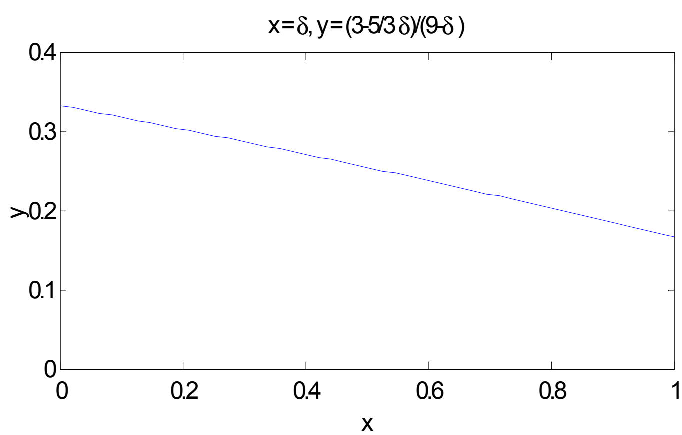

Carrot and Stick

Produce ¼ at the beginning; at ant $t > 0$, produce ¼ if both produced ¼ or both produced x at t-1; otherwise, produce x.

Two Phase: Cartel & Punishment

$V_C = 1/8(1-\delta)$. $V_x = x(1-2x) + \delta V_C$.

$V_{D|C} = \max q(1-1/4-q) + \delta V_X = (3/8)^2 + \delta V_X$

$V_{D|x} = \max q(1-x-q) + \delta V_X = (1-x)^2/4 + \delta V_X$

$V_C \geq V_{D|C} \Leftrightarrow V_C \geq (3/8)^2 + \delta^2 V_C + \delta\, x(1-2x)$

$\Leftrightarrow (1-\delta^2)\, V_C - (3/8)^2 \geq \delta\, x(1-2x) \Leftrightarrow (1+\delta)/8 - (3/8)^2 \geq \delta\, x(1-2x)$

$V_X \geq V_{D|C} \Leftrightarrow (1-\delta)V_x \geq (1-x)^2/4 \Leftrightarrow (1-\delta)(x(1-2x) + \delta/8(1-\delta)) \geq (1-x)^2/4$
$\Leftrightarrow (1-\delta)x(1-2x) + \delta/8 \geq (1-x)^2/4$

$$2x^2 - x + 1/8 - 9/64\delta \geq 0$$

$$(9/4-2\delta)x^2 - (3-2\delta)x + \delta/8(1-\delta) \leq 0$$

Incomplete information

Incomplete information

We have incomplete (or asymmetric) information if one player knows something (relevant) that some other player does not know.

An Example

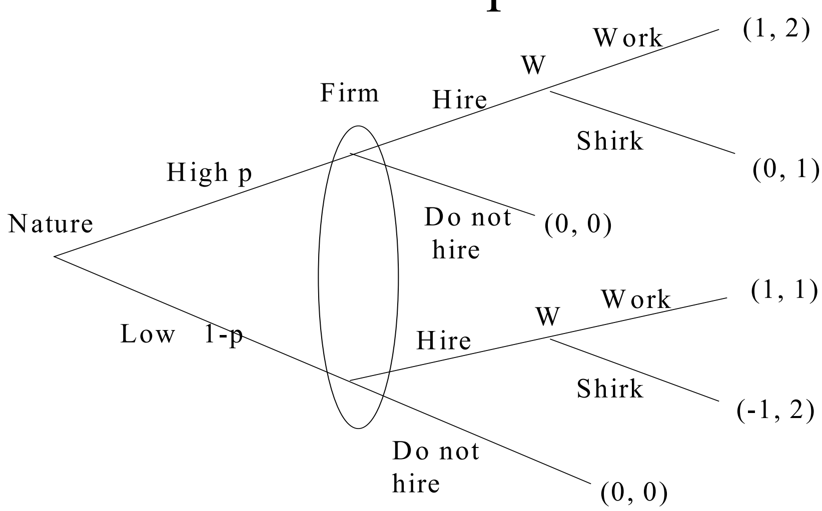

The same example

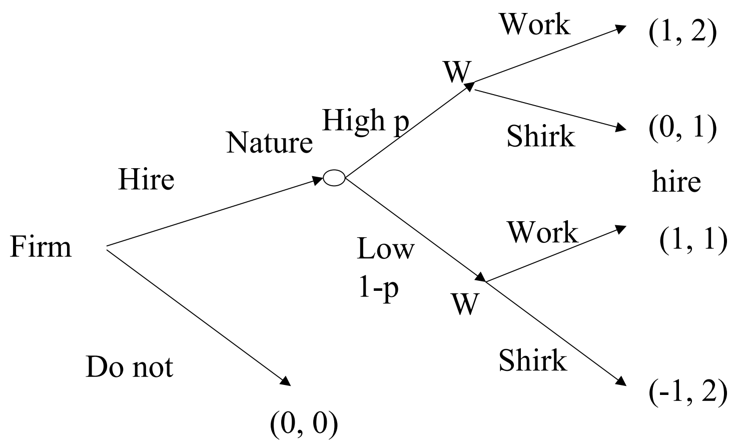

Another Example

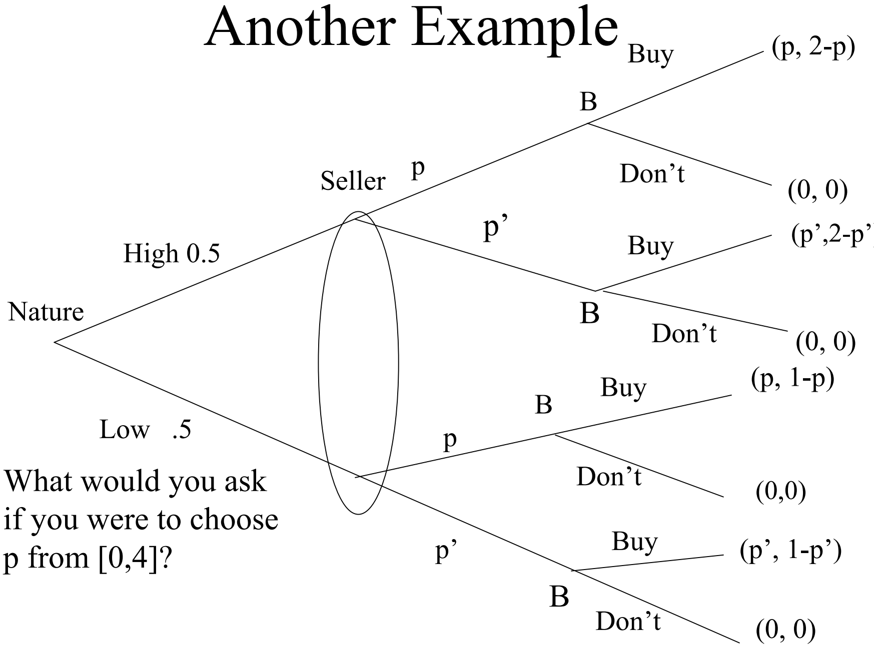

What would you ask
if you were to choose
p from [0,4]?

Same "Another Example"

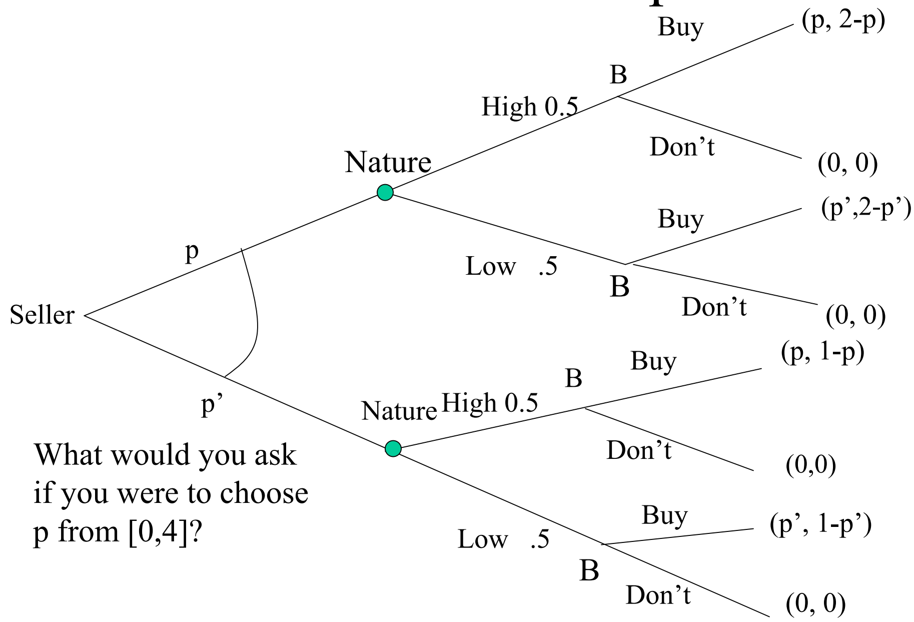

What would you ask if you were to choose p from [0,4]?

Bayes' Rule

- $\text{Prob}(A|B) = \dfrac{\text{Prob}(A \text{ and } B)}{\text{Prob}(B)}$
- $\text{Prob}(A \text{ and } B) = \text{Prob}(A|B)\text{Prob}(B) = \text{Prob}(B|A)\text{Prob}(A)$
- $\text{Prob}(A|B) = \dfrac{\text{Prob}(B|A)\text{Prob}(A)}{\text{Prob}(B)}$

Example

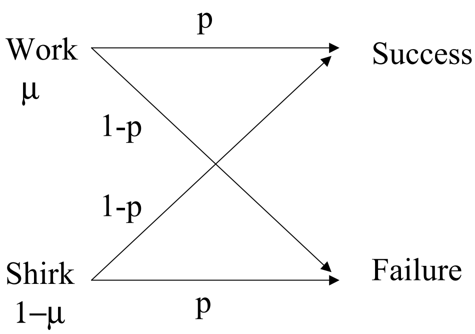

- Prob(Work|Success) = $\mu p/[\mu p + (1-\mu)(1\text{-}p)]$
- Prob(Work|Failure) = $(1\text{-}\mu)p/[\mu(1-p) + (1-\mu)p]$