Introduction to Game Theory

1. The Rules of the Game

1.1 Basic Definitions

Players are the individuals who make decisions. Each player's goal is to maximize his utility by choice of actions.

Nature is a pseudo-player who takes random actions at specified points in the game with specified probabilities.

An **action** or **move** by player I, denoted a_i, is a choice he can make.

Player i's **action set**, $A_i = \{a_i\}$, is the entire set of actions available to him.

An **action profile** is an ordered set $a = \{a_i\}$, $(i = 1,...,n)$ of one action for each of n players.

One must also specify when the actions are available to a player. The is called the **order of play**.

Player i's **strategy** s_i is a rule that tells him which action to choose at each instant of the game, given his **information set**, his knowledge at a particular time if the values of different variables and of what actions have been previously played.

Player i's **strategy set** or **strategy space** $S_i = \{s_i\}$ is the set of strategies available to him.

A **strategy profile** $s = (s_1,...,s_n)$ is an ordered set consisting of one strategy for each of then players in the game.

> The concept of the strategy is useful because the action a player wishes to pick depends on the past actions of Nature and the other players. However, only rarely can we predict how he will respond to the outside world unconditionally. The player's strategy is a complete set of instructions which tells him what actions to pick in <u>every</u> conceivable situation, even if he does not expect to reach that situation.

By player i's **payoff** $\pi_i(s_1,...,s_n)$, we mean either:

(1) The utility player i receives after all players and nature have picked their strategies and the game has been played out; or

(2) The expected utility player i receives as a function of the strategies chosen by himself and other players.

> As implied by the above definitions, the term “payoff” is used for actual payoff and the expected payoff.

The **outcome** of a game is the set of interesting elements that the modeler picks from the values of the actions, payoffs, and other variables after the game is played out.

Example: The OPEC Model

Players
Saudi arabia, Libya, Venezuela, Kuwait, Nigeria.

Information
All players know the value of the demand, but they choose their actions simultaneously.

Order of Play
(1) The players simultaneously choose supply schedules consisting of the quantity each will supply at each possible market price in 1988 (this model does concern itself with 1989)
(2) Nature picks world demand for oil to be either *Weak* or *Strong*, with equal probability.

Strategies
The strategies are the same as actions, since no information is revealed that might affect the action a player chooses.

Payoffs
for each country the payoff is +100 if its oil revenue remains above a country specific amount necessary to avoid a military coup, -100 if revenue galls below this amount.

Outcomes
The outcome includes the quantities supplied, the state of demand, the resulting revenues and market price, and whether or not each country has a coup.

Equilibrium
To predict the outcome of a game, the modeler focuses on the possible strategy profiles, since it is the interaction of the different players’ strategies that determines what happens.

> An **equilibrium** $s^* = \left(s_1^*, \ldots, s_n^*\right)$ is a strategy profile consisting of the best strategy for each of the n players in the game.

The **equilibrium strategies** in set s* are the strategies players pick in trying to maximize their individual payoffs.

In the OPEC Model the resulting market price would be part of the **equilibrium outcome**, but the equilibrium itself would be the strategy profile that generated the outcome.

To find the equilibrium, the modeler must, in addition to defining the players, strategies, and payoffs, decide what "best strategy" means.

The Two Best-Known Concepts: Dominant Strategy and Nash Equilibrium

1.2 Dominant Strategies

For any vector $y = (y_1,...,y_n)$, denoted by y_{-1} the vector $(y_1,...,y_{i-1},y_{i+1},...,y_n)$, which is the portion of y not associated with player i.

Using this notation y_{-i} is the profile of every player except i.

Player i's **best response** or **best reply** to the strategies s_{-i} chosen by the other players is the strategy s_1^* that yields him the greatest payoff; that is,

$$\pi_i(s_i^*, s_{-1}) \geq \pi_i(s_i', s_{-1}) \forall s_i' \neq s_i^*$$

In other words, the best response is strongly best if no other strategies are equally good, and weakly best otherwise.

From this we get the first important equilibrium concept: the dominant-strategy equilibrium.

The strategy s_1^* is a **dominant strategy** if it is a player's strictly best response to any strategies the other players might pick, in the sense that whatever strategies they pick, his payoff is highest with s_1^*. Mathematically,

$$\pi_i(s_i^*, s_{-1}) \geq \pi_i(s_i', s_{-1}) \forall s_{-i}, \forall s_i' \neq s_i^*$$

His inferior strategies are **dominated strategies.**

A **dominant-strategy equilibrium** is a strategy profile consisting of each player's dominant strategy.

Dominant strategy is not as important a concept as Nash equilibrium. Most games do not have dominant strategies, and the players must try to figure our each others' actions in order to arrive at their own strategies. However it does crop up in under special circumstances. An important example of this is the Prisoner's Dilemma in the table below. It is a **2-by-2 game**, because each of the two players

-- Row and Column -- has two possible actions in his action set -- *Confess* and *Deny*. The numbers represent the payoffs.

The Prisoner's Dilemma

		Column		
		Deny		*Confess*
	Deny	-1,-1	→	-10,0
Row:		↓		↓
	Confess	0,-10	→	-8,-8

Payoffs to: (Row, Column)

Each player has a dominant strategy. It is *Confess* (prove it by comparing the payoffs the receive when the other player either strategy). he dominant-strategy equilibrium is (*Confess, Confess)*, and the equilibrium payoffs are (-8, -8), which is worse for both players that (-1, -1).

The apparently irrational results of the Prisoner's Dilemma occurs because of the we have modeled it as a cooperative game instead of a non-cooperative game. A **cooperative game** is a game in which the players can make binding commitments, as opposed to a **non-cooperative game**, in which they cannot. Thus Row and Column could have committed themselves to Deny, giving themselves a better payoff. Cooperation does not have to be the produce of outright collusion however. Through **side-payments** in which transfers change the prescribed payoffs, for example, one player can bribe the other payer to commit to a specific action.

1.4 Nash Equilibrium

For the vast majority of games, which lack even iterated dominance equilibria, modelers use the Nash equilibrium, which is the most important and most frequently encountered equilibrium concept.

A strategy profile s* is a **Nash equilibrium** if no player has incentive to deviate from his strategy given that the other players do not deviate. Formally,

$$\forall i \ \pi_i(s_i^*, s_{-1}^*) \geq \pi_i(s_i', s_{-1}^*) \ \forall s_i' \ .$$

The way to approach Nash equilibrium is to propose a strategy profile and test whether each player's strategy is a best response to the others' strategies. The definition of Nash equilibrium lacks the "$\forall s_{-i}$" of a dominant-strategy equilibrium, so a Nash strategy need only be a best response to the other Nash strategies, not all possible strategies. Thus while every dominant-strategy equilibrium is a Nash equilibrium, not every Nash equilibrium is a dominant-strategy equilibrium. Also, it is often based on a belief or prediction about what the other player will do.

In the Prisoner's Dilemma above, equilibrium (*Confess, Confess*) is a Nash equilibrium because for each player to *Confess* is a best strategy given that it is reasonable to expect the other player to *Confess*. However, (*Deny, Deny*) is Pareto superior. If *Deny* is not an irrational possibility then we have the common problem of Nash equilibrium -- the possibility of multiple equilibria. To solve this problem the modeler could add more details to the rules of the game, or he could use an **equilibrium refinement**, adding conditions to the basic equilibrium concept until only one strategy profile satisfies the refined equilibrium concept. To be honest, no approach is completely satisfactory. Your assumption that a Nash equilibrium is played will have to rely on there being some mechanism or process that leads all the players to expect the same equilibrium. You must make the case that your choices are reasonable and be prepared to defend yourself against criticism.

> A Pareto-dominant Equilibria in Subgames (Pareto Perfection) is when we use the refinement of ruling out any strategy profiles that are Pareto dominated by Nash equilibria.

The table below demonstrates much of what has been said about Nash equilibrium. The Battle of the Sexes represent the conflict between a man who wants to go to a prize fight and a woman who wants to go to a ballet. While selfish, they are deeply in love, and would, if necessary, sacrifice their preferences to be with each other. Less romantically, their payoffs are given by the following table:

Battle of the Sexes

		Woman		
		Prize Fight		*Ballet*
	Prize Fight	**2,1**	←	0,0
Man		↑		↓
	Ballet	0, 0	→	**1,2**

Payoffs to: (Man, Woman)

Battle of the Sexes does not have an iterated dominance equilibrium. It has two Nash equilibria, one of which is the strategy profile (*Prize Fight, Prize Fight)*. Given that the man chooses *Prize Fight*, so does the woman; given that the woman chooses *Prize Fight*, so does the man. The strategy profile (*Ballet, Ballet*) is another Nash equilibrium, by the same line of reasoning. If they do not talk beforehand, the man might go to the ballet and the woman to the fight, each mistaken about the other's belief's. But even if the players do not communicate, Nash equilibrium is sometimes justified by repetition of the game. If the couple do not talk, but repeat the game night after night, one may suppose that eventually they'll settle on one or another of the Nash equilibria. Who moves first can also be an important refinement. If the man could buy the fight ticket in advance, his commitment would induce the woman to go to the fight. The player who moves first has a **first-mover advantage**.

(An economic application of the Battle of the Sexes game is the choice of an industry-wide standard when two firms have different preferences but both want a common standard so as to encourage consumers to buy the product. These are part of a large class of games known as **coordination games** which share the common feature that the players need to coordinate on one of multiple Nash equilibria. In other words, games in which the payoff is high if one follows the leader.)

1.5 Focal Points and Boundaries

Focal points: Nash equilibria which for psychological reasons are particularly compelling. Formalizing what makes a strategy profile a focal point is not an easy task and depends on the context.

The **Boundary** is a particular kind of focal point. There is an arbitrary discontinuity in behavior at the boundary. In other words, where the boundary is located is arbitrarily determined but once the boundary is established, it takes on additional significance, because behavior with respect to the boundary conveys information.

2. Information

2.1 The Extensive Form and the Game Tree (Sequential decision making)

In addition to the matrix forms, there are two other ways to describe a game are the extensive form and the game tree. First we need to define their building blocks.

A **node** is a point in the game at which some player or Nature takes an action, or the game ends.

A **successor** to node X is a node that may occur later in the game if X has been reached.

A **predecessor** to node X is a node that must be reached before X can be reached.

A **starting node** is a node with no predecessors.

An **end node** or **end point** is a node with no successor.

A **branch** is one action in a player's action set at a particular node.

A **path** is a sequence of nodes and branches leading from the starting node to an end node.

These concepts can be used to define the extensive form and the game tree.

The **extensive form** is a description of a game consisting of:

(1) A configuration of nodes and branches running without any closed loops from a starting node to its end nodes.

(2) An indication of which node belongs to which player.

(3) The probabilities that Nature uses to choose different branches at its nodes.

(4) The information sets into which each player's nodes are divided.

(5) The payoffs for each player at each end node.

The **game tree** is the same as the extensive form except (5) is replaced with:

(5′) The outcomes at each node.

If the outcome is defined as the payoff profile, the extensive form is the same as the game tree. The Follow-the-Leader game in the extensive form would look like this:

Follow-the-Leader in Extensive Form

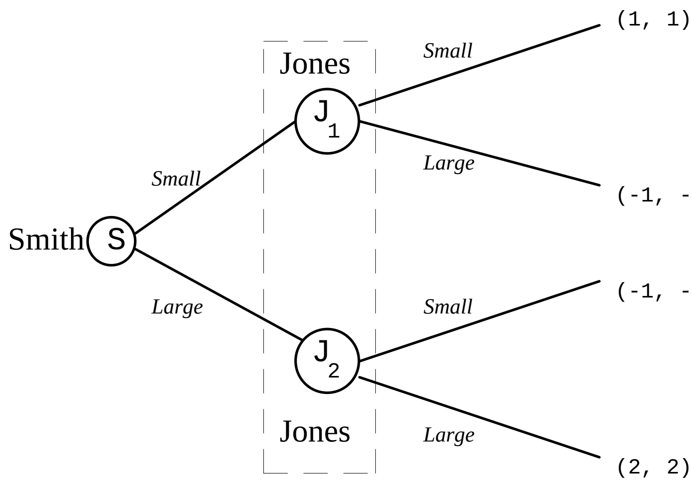

Payoffs to: (Smith, Jones)

2.2 Perfect, Certain, Symmetric, and Complete Information

Information Categories

Information category	**Meaning**
perfect	Each information set is a singleton
certain	Nature does not move after any player moves
symmetric	No player has information different from other players when he moves, or at the end node
complete	Nature does not move first, or her initial move is observed by every player

In a game of **perfect information** each information set is a singleton. Otherwise the game is one of **imperfect information**.

A game of **certainty** has no moves by Nature after any player has moved. Otherwise the game is one of **uncertainty**.

When players face uncertainty, we need to specify how they evaluate their uncertain future payoffs. The obvious way to model their behavior is to say that the players maximize the expected values of their utilities. Players who behave in this way are said to have **von Neumann-Morgenstern utility functions**.

In a game of **symmetric information**, a player's information set at

(1) any node where he chooses an action,
or
(2) an end node contains at least the same elements as the information sets of every other player. Otherwise the game is one of **asymmetric information**.

The essence of asymmetric information is that some player has useful **private information**: an information partition that is different and not worse than any other player's.

In a game of **incomplete information**, Nature moves first and is unobserved by at least one of the players. Otherwise the game is one of **complete information**.

2.2 Harsanyi Transformation and Bayesian Games

The Harsanyi Transformation

Games of incomplete information cannot be analyzed using standard game theory techniques because the equilibrium concept requires that each player knows the others' payoffs (motives). Harsanyi (1967 and 1968) showed that an incomplete information game can be modeled as a complete but imperfect information game. This is accomplished by including Nature as another player, wherein Nature initially chooses the type of the other players (we assume each player knows his type). Now player A's incomplete information about player B's type becomes imperfect information about nature's move. Players share a common "belief" as to how Nature makes its probabilistic choice. The new game is called a *Bayesian game* and is analyzed using standard game theory techniques.

Bayesian Equilibrium

There are two useful concepts we will need in this section:

A player's **type** is the strategy set, information partition, and payoff function which Nature chooses for him at the start of a game of incomplete information.

and

A **state of the world** is a move by Nature.

The term Bayesian equilibrium is used to refer to a Nash equilibrium in which players update their **prior beliefs** according to Bayes' Rule. The two-step procedure of checking a Nash equilibrium now becomes a three-step procedure.

(1) Propose a strategy profile.
(2) See what beliefs the strategy profile generates when players update their beliefs in response to each others' moves.
(3) Check that, given those beliefs together with the strategies of the other players, each player is choosing a best response for himself.

The following terminology may be helpful in understanding Bayes' Rule:

Bayesian Terminology

Name	**Meaning**
likelihood	$Prob(data\|event)$
marginal likelihood	$Prob(data)$
conditional likelihood	$Prob(data\ \mathrm{X}\|data\ \mathrm{Y},\ event)$
prior	$Prob(event)$
posterior	$Prob(event\|data)$

(where "|" denotes "given that")

Bayes' Rule shows how to revise a prior belief in the light of new information such as a given move by the other player. It uses two pieces of information, the likelihood of seeing the other player make the specific move (marginal likelihood) and the likelihood of seeing the player make the move given that Nature did not choose some state *x*. In a general form, for Nature's move *x* and the observed data, Bayes' Rule is

$$Prob(x|data) = \frac{Prob(data|x)\,Prob(x)}{Prob(data)},$$

which for our example would be

$$(Posterior\ for\ Nature's\ Move) = \frac{(Likelihood\ of\ Player's\ Move)\cdot(Prior\ for\ Nature's\ Move)}{(Marginal\ likelihood\ of\ Player's\ Move)}.$$

As always in Nash equilibrium, the modeler assumes that the players know which equilibrium strategies are being played out, even though they do not know which particular actions are being chosen.

Much of this can be kept simple if one follows the **Harsanyi doctrine:** a modeler should assume that beliefs begin the same but diverge because of private information. This can be seen in terms of a plaintiff and defendant who are bargaining for a pretrial settlement. Their perception of their relative chance of winning at trial is based on private information, not because the modeler assigns them different beliefs at the start of the game.

3. Mixed and Continuous Strategies

3.1 Mixed Strategies

We invoked the concept of Nash equilibrium to provide predictions of outcomes without dominant strategies, but some games lack even a Nash equilibrium. It is often useful and realistic to expand the strategy space to include random strategies, in which case a Nash equilibrium almost always exists. These random strategies are called "mixed strategies."

A pure strategy constitutes a rule that tells the player what action to choose, while a mixed strategy constitutes a rule that tells him what dice to throw in order to choose an action. If a player pursues a mixed strategy, he might choose any of several different actions in a given situation. In other words, a player's mixed strategy in a given situation is to take action *x* 20 percent of the time, action *y* 35 percent of the time and action *z* 45 percent of the time. There is probably a reason why a player picks one strategy over another, but all a model with mixed strategies requires to be a good description of the world is that the actions appear random to observers. Another way to reinterpret the mixed strategy game is to imagine instead of a single player there are many, with identical tastes and payoff functions, all of whom must be treated alike by the other player. Then we can say that 20 percent of the group will choose action *x*, 35 percent will choose *y*, and 45 percent will choose *z*. Another interpretation is that a single player is drawn from a random population with the above distribution.

Existence of Equilibrium

Sometimes the modeler cannot find a particular equilibrium easily, but is able to show what characteristics an equilibrium must have if it exists, and would like to prove that it does indeed exist. One of the strong points of **Nash equilibria** is that they exist, in mixed strategies if not in pure, in practically every game one is likely to encounter.

3.2 Mixed Strategies with General Parameters and N Players

A General 2-by-2 Game

Column

		Left (θ)		*Right*(1 - θ)
	Up(γ)	*a, w*	→	*b, x*
Row:		↑		↓
	Down(1 - γ)	*c, y*	←	*d, z*

Payoffs to: (*Row*, *Column*)

To find the game's equilibrium, equate the payoffs from the pure strategies. For Row, this yields

$$\pi_{\text{Row}}\,(Up) = \theta a + (1- \theta)b \tag{1}$$

and

$$\pi_{\text{Row}}\,(Down) = \theta c + (1- \theta)d \tag{2}$$

Equation (1) and (2) gives

$$\theta(a + b - c) + b - d = 0, \tag{3}$$

Which yields

$$\theta^* = \frac{d-b}{(d-b)+(a-c)}. \tag{4}$$

Similarly, equating the payoffs for Column gives

$$\pi_{\text{column}}\,(Left) = \gamma w + (1- \gamma)y = \pi_{\text{column}}\,(Right) = \gamma x + (1- \gamma)z \tag{5}$$

which yields

$$\gamma^* = \frac{z-y}{(z-y)+(w-x)}. \tag{6}$$

4 Dynamic Games with Symmetric Information

(By "dynamic game" we mean one with a sequence of moves before reaching an equilibrium)

4.1 Backward Induction and Subgame Perfectness

A **subgame** is a game consisting of a node which is a singleton in every player's information partition, that node's successors, and the payoffs at the associated end nodes.

A strategy profile is a **subgame perfect Nash equilibrium** if (a) it is a Nash equilibrium for the entire game; and (b) its relevant action rules are a Nash equilibrium for every subgame. Basically this is a refinement of the model to find one Nash equilibrium. It is a carry over of the concept of backward induction to games of imperfect information.

Backward Induction applies to finite games of perfect information. The algorithm begins by determining the optimal choices in the final stage K for each history h^K -- that is, the action for the player on move, given history h^K, that maximizes that player's payoff conditional on h^K being reached. Then we work back to stage K - 1, and determine the optimal action for the player on move there, given that the player on mover at stage K with history h^K will play the action we determined previously. Backward induction is the first and simplest form of an equilibrium refinement. It relies on the presumption that an equilibrium (U,R) is not credible because it relies on an "empty threat" by player 2 to play R. The Cournot model is basically an example of this, as we pick one firm's strategy based on our assumptions of what the other firm previously would have chosen.

Subgame perfectness, on the other hand, is best understood with an example. In section 2.1 a flaw of Nash equilibrium was revealed in the game Follow the Leader, which has three pure strategy Nash equilibria of which only one is reasonable. The players are Smith and Jones, who choose disk sizes. Both their payoffs are greater if they choose the same size and greatest if they coordinate on *Large*. Smith moves first, so his strategy set is (*Small*, *Large*). Jones' strategy is more complicated, because it must specify an action for each information set, and Jones' information set depends on what Smith chose. A typical element of Jones' strategy set is (*Large*, *Small*), which specifies that he chooses *Large* if Smith chose *Large*, and *Small* if Smith chose *Small*. From the strategic form we found the following three Nash equilibria.

Equilibrium	**Strategies**	**Outcomes**
E_1	{*Large*, (*Large*, *Large*)}	Both pick *Large*.
E_2	{*Large*, (*Large*, *Small*)}	Both pick *Large*.
E_3	{*Small*, (*Small*, *Small*)}	Both pick *Small*.

Only Equilibrium E_2 is reasonable, because the order of the moves should matter to the decisions which the players make. The problem with the strategic form, and, thus, with simple Nash equilibrium, is that it ignores who moves first. Smith moves first, and it seems reasonable that Jones should be allowed -- should be required -- to rethink his strategy after Smith moves. Consider Jones' strategy of (*Small*, *Small*) in equilibrium E_3. If Smith deviated from equilibrium by choosing *Large*, if would be unreasonable for Jones to stick to the response *Small*. Instead, he should also choose *Large*. But if Smith expected a response of *Large*, he would have chosen *Large* in the first place, and E_3 would not be an equilibrium. A similar argument shows that it would be irrational for Jones to Choose (*Large*, *Large*), and we are left with E_2 as a unique equilibrium.

We say that E_1 and E_2 are Nash equilibria but not "perfect" Nash equilibria. A strategy profile is a perfect equilibrium if it remains an equilibrium on all possible paths, includes not only the equilibrium path (the path through the game tree that is followed in equilibrium) but all the other paths, which branch off into different subgames.

In the context of subgame perfect Nash equilibrium, the term **sequential rationality** is often used to denote the idea that a player should maximize his payoffs at each point in the game, re-optimizing his decisions at each point and taking into account the fact that he will re-optimize in the future. This is a blend of the economic ideas of ignoring sunk costs and rational expectations.

One reason why perfectness is a good equilibrium concept is because it represents the idea of sequential rationality. A second reason is that a weak Nash equilibrium is not robust to small changes in the game. Remember, though, it is limited to games of symmetric information. Other equilibrium concepts are needed to extend the concept of perfectness to games of asymmetric information.

4.2 Discounting

A model in which the action takes place in real time must specify whether payments and receipts are valued less if they are made later, i.e., whether they are **discounted**. Discounting is measured by the discount rate or the discount factor.

> The **discount rate**, r, is the extra fraction of a payoff unit needed to compensate for delaying receipts by one period.
>
> The **discount factor**, δ, is the value in present payoff units of one payoff unit to be received one period from the present. [$\delta = 1/(1 + r)$]

Whether to put discounting into a model involves two question. The first is whether the added complexity will be accompanied by a change in the results or by a surprising demonstration of no change in the results. A second, more specific question is whether the events of the model occur in real time, so that discounting is appropriate.

Discounting has two important sources: time preference and a probability that the game might end, represented by the rate of time preference, ρ, and the probability each period that the game ends, θ. It is usually assumed that ρ and θ are constant. If they both take the value zero, the player does not care whether his payments are scheduled now or ten years from now. Otherwise, a player is indifferent between $x/(1 + \rho)$ now and x guaranteed to be paid one period later. With probability $(1 - \theta)$ the game continues and the later payment is actually made, so the player is indifferent between $(1 - \theta)\, x/(1 + \rho)$ now and the promise of x to be paid one period later contingent upon the game still continuing. The discount factor is therefore

$$\delta = \frac{1}{1+r} = \frac{(1-\theta)}{(1+\rho)}.$$

Discounting is especially important in bargaining theory.

5 Repeated Games with Symmetric Information

5.1 Infinitely Repeated Games and the folk Theorem.

What if the Prisoner's Dilemma Game was not a **one-shot game** (unrepeated), instead went on for 20 periods? If one player chooses *Deny* one may assume that the other would also choose *Deny* in the next period and so on and so on. Thus they can achieve the Pareto optimal solution. However, the game does end in period 20, and the player who chooses at the moment is better off if he chooses *Confess*. With this expectation it is rational for the other player to choose *Confess* in period 19, but this implies that it is rational for the players to have chosen *Confess* from the beginning. (This is also known as "the chain store paradox)

This solution may not seem realistic, since we know that people do commit themselves to specific behavior when interaction with another is repeated. Where information is symmetric we can incorporate such behavior by assuming the game is infinitely repeated; that is, no end period. In addition, we use a strategy such as the Grim Strategy or Tit-for-Tat. While the assumption that past play does not influence the feasible actions or payoff functions in current period seems unrealistic, it does provide some good approximations of some some-run relationship -- particularly those where "trust" or "social pressure" are important.

Grim Strategy

(*1*) *Start by choosing Deny.*

(*2*) *Continue to choose Deny unless some player has chosen Confess, in which case choose Confess forever.*

Tit-for-Tat

(*1*) *Start by choosing Deny.*

(*2*) *Thereafter, in period n choose the action that the other player chose in period (n - 1).*

The Folk Theorem

In an infinitely repeated n-person game with finite action sets at each repetition, any combination of actions observed in any finite number of repetitions is the unique outcome of some subgame perfect equilibrium.

What the Folk Theorem tells us is that claiming that particular behavior arises in a perfect equilibrium is meaningless in an infinitely repeated game. If an infinite amount of time always remains in the game, a way can always be found to make one player willing to punish some other player for the sake of a better future, even if the punishment currently hurts the punisher as well as the punished. Any finite interval of time is insignificant compared to eternity, so the threat of future reprisal makes the players willing to carry out the punishments needed.

5.2 Reputation

Reputation makes threat to punish credible. The product-quality game, and other *one-sided prisoner's dilemma games* are examples of the use of reputation. (see Rasmussen, pp. 129-134.) Note: a player is most likely to be willing to incur short-run costs to build up his reputation when he is patient and his planning horizon is long.

6 Dynamic Games with Asymmetric Information

6.1 Perfect Bayesian Equilibrium

In games of asymmetric information, we will still require that an equilibrium be subgame perfect, but the mere forking of the game tree might not be relevant to a player’s decision, because with asymmetric information he does not know which fork the game has taken. You cannot eliminate subgames that are not subgame perfect if their beginning nodes are not singletons in the information partition. Thus, the implausible Nash equilibria escape elimination by a technicality. The equilibrium concept needs to be refined, under conditions of asymmetric information. Two general approaches can be taken: “trembling hand-perfect” equilibrium and “perfect Bayesian” or “sequential” equilibrium.

Trembling-Hand Perfectness

Trembling-hand perfectness is an equilibrium concept in which a strategy that is to be part of an equilibrium must continue to be optimal for player even if there is a small chance that the other player will pick an out-of-equilibrium action (i.e., that the other player’s hand will “tremble”). it is defined for games with finite action sets as follows:

Unfortunately, it is often hard to tell whether a strategy profile is trembling-hand perfect, and the concept is undefined for games with continuous strategy spaces because it is hard to work with mixtures of a continuum. Moreover, deciding why one tremble should be more common than another may be difficult.

Perfect Bayesian Equilibrium and Sequential Equilibrium

The second approach to asymmetric information is to start with the prior beliefs, common to all players, that specify the probabilities with which nature chooses the types of the players at the beginning of the game. Some of the players observe Nature's move and update their beliefs, while other players can update their beliefs only by deductions they make from observing the actions of the informed players. When players update their beliefs, they assume that the other players are following the equilibrium strategies, but since the strategies themselves depend on the beliefs, equilibrium can no longer be defined based on strategies alone. It is now a strategy profile and a set of beliefs such that the strategies are best responses. The combination of beliefs and strategies is called an **assessment**.

Notice that perfect Bayesian equilibrium is not defined structurally, like subgame perfectness, but rather in terms of optimal responses. To find a Nash equilibrium, the modeler thinks about his game, picks a plausible strategy profile, and tests whether the strategies are best responses to each other. To make it a perfect bayesian equilibrium, he notes which actions are never taken in equilibrium and specifies the beliefs that players use to interpret those actions. He then tests whether each player's strategies are best responses given his beliefs at each node, checking in particular whether any player would like to take and out-of-equilibrium action in order to set in motion the other players' out-of-equilibrium beliefs and strategies. The concept of perfect Bayesian equilibrium leaves the modeler free to specify how the players form beliefs from the additional information conveyed by other players' action, so long as the beliefs do not violate Bayes' Rule (this does not imply that one can make ridicules assumptions).

Example: PhD Admissions Game

Suppose that a university knows that 90 percent of the population hate economics and would be unhappy in its PhD program, and that 10 percent love and would do well. In addition, it cannot observe the applicant's type. If the university rejects an application, its payoff is 0 and the applicant's is -1 because of the trouble needed to apply. If the university accepts the application of someone who hates economics, the payoffs of both university and student are -10, but if the applicant loves economics the payoffs are +20 for each player. The extensive form of the game appears below. The population are represented by a node at which nature chooses the student to be a *Lover* or *Hater* of economics.

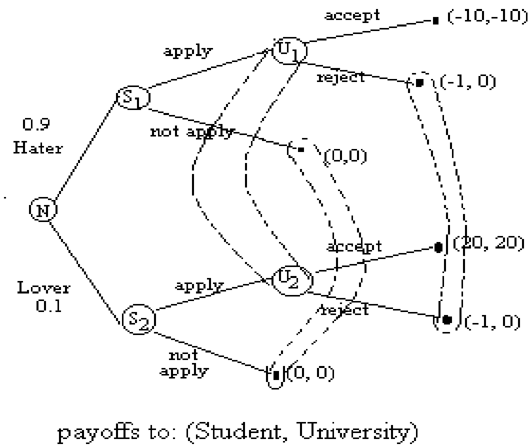

PhD Admissions is a signaling game. It has various perfect Bayesian equilibria that differ in their out-of-equilibrium beliefs, but the equilibria can be divided into two distinct categories, depending on the outcome: the **separating equilibrium,** in which the lovers of economics apply and the haters do not, and the **pooling equilibrium,** in which neither type of student applies.

A Separating Equilibrium for PhD Admissions
Student: *Apply|Lover, Do not Apply|Hater*
University: *Admit*

The Separating equilibrium does not need to specify out-of-equilibrium beliefs, because Bayes' Rule can always be applied whenever both of the two possible actions *Apply* and *Do Not Apply* can occur in equilibrium.

A Pooling Equilibrium for PhD Admissions
Student: *do Not Apply|Lover, Do not Apply|Hater*
University: *Reject*
Beliefs: *Prob*(*Hater|Apply*) = 0.9 (passive conjectures)

The pooling equilibrium is supported by **passive conjectures**; that is, out-of-equilibrium behavior leaves beliefs unchanged from the prior. Here the university is willing to reject any student who foolishly applied, believing that is is a Hater with 90 percent probability; and both types of students refrain from applying because they believe correctly that they would be rejected and receive a payoff of -1.

There are a variety of other refinements of the equilibrium concept.

The Intuitive Criterion.

$Prob(Hater|Apply) = 0$

If there is a type of informed player who could not benefit from the out-of-equilibrium action no matter what beliefs were held by the uninformed player, the uninformed player's belief must put zero probability on that type.

Complete Robustness.

$Prob(Hater|Apply) = m, 0 \leq m \leq 1.$

Under this approach, the equilibrium strategy profile must consist of responses that are best, given any and all out-of-equilibrium beliefs. Complete robustness rules out a pooling equilibrium in PhD Admissions, because a belief like $m = 0$ makes accepting applicants a best response, in which case only the *Lover* will apply. A useful first step in analyzing conjectured pooling equilibria is to test whether they can be supported by extreme beliefs such as $m = 0$ and $m = 1$.

An ad hoc Specification.

$Prob(Hater|Apply) = 1.$

Sometimes the modeler can justify beliefs by the circumstances of the particular game. Here, one could argue that anyone so foolish as to apply knowing that the university would reject them could not possible have the good taste to love economics. This supports the pooling equilibrium also.

An alternative approach to the problem of out-of-equilibrium beliefs is to remove its origin by building a model in which every outcome is possible in equilibrium because different types of players take different equilibrium actions. One then uses Bayes' Rule. This approach is especially attractive if the modeler takes the possibility of trembles literally.