

Information and Equilibria in Economic Games

- I. We now turn to settings in which dominant pure strategies do not exist and/or in which informational problems are explicitly addressed.**
- A.** A good deal of economic analysis and model building implicitly ignores information problems.
- i. That is to say, the people modeled are assumed to know as much as can be known about the problem facing them.
 - ii. For example, the setting analyzed may be assumed to be one in which there is complete certainty, and the decision makers simply optimize in a setting where they know all that can be known about their particular problem: e. g. they know their objective function (profit or utility) and their constraints.
 - iii. A similar case occurs when there are random phenomena that affect expected utility levels or payoffs, but where the functions are well understood. The rent-seeking model in the previous lecture was such a case.
 - iv. If the probability functions are well understood, the persons being modeled simply maximize the "expected" value of their objective functions given their various certain and probabilistic constraints.
 - v. Optimal purchases of lottery tickets, investment in risky processes, rent-seeking, insurance markets etc. are usually modeled in this way.
- B.** Models that take account of uncertainty are implicitly acknowledging information problems that are relevant for the decisions to be made.
- C.** In addition to modeling uncertainty, models that take explicit account of imperfect information may also attempt to model learning or changes in information.
- i. In such settings, one may be able to predict how new information will affect future choices and so model informational strategies with models that resemble Stackelberg games.
 - ii. In other cases, one may simply attempt to model how people adapt to different circumstances as they learn a bit more about a particular probability function, constraint, or their own objectives.
 - iii. What is known, how one comes to know "it" and the manner in which one can learn more about "it" are all matters of interest in models with information problems.
 - For example, consider the model of "rational ignorance" developed in the early Public Choice literature by Downs in the 1950s and Tullock in the 1960s.
 - If voters or consumers are rationally ignorant, the extent of their knowledge can be "manipulated" by manipulating information costs.
- D.** In a game setting, knowledge problems can be modeled in several ways.
- i. For example, most Cournot models implicitly assume that game participants do not know the strategies adopted by their competitors. In contrast, in Stackelberg games the second mover knows what the first has done.
 - ii. Attempts to prevent the other players from knowing what they will do can also be modeled, as when a player randomly chooses among strategies in a repeated game.

E. This lecture provides an overview of some models of strategic choice in which uncertainty and other informational imperfections play a role. The applications are mostly from the economics literature.

II. Mixed Strategy Equilibria: "Random" Strategies

A. It bears noting that many "finite" games that **do not** have a Nash Equilibrium in Pure Strategies.

- For example, many zero sum games lack an equilibrium in pure strategies.
- The rock-scissors-paper game also lacks an equilibrium in pure strategies.
- Nonetheless, if random play is allowed, there is (nearly) always an equilibrium in mixed strategies.

B. A **mixed strategy** is a probability function defined over a player's strategy set.

- i. The players select among probability functions defined over their strategy sets.
 - Note that a pure strategy can be regarded as a probability function where the probability of a particular strategy is 1, and the probability of all the other strategies is zero.
- ii. It turns out that if the probability of playing one or another strategy is "the" strategy, a 2-strategy game (or other finite strategy game) becomes a game with an continuous strategy set.
 - The equilibrium that emerges, however, differs from the Cournot model (in most cases) in that it is driven by consideration of the responses of the opposition.
- iii. To calculate payoffs when using mixed strategies, players are normally assumed to maximize their **expected utility** (the expected payoff) associated with their mixed strategy.
 - For example, if there are four possible payoffs, $U_1, U_2, U_3,$ and U_4 and the probability of realizing U_1 is $P_1,$ probability of realizing U_2 is $P_2,$ the probability of realizing U_3 is P_3 and or realizing U_4 is $P_4.$
 - The expected payoff is $U^E = P_1U_1 + P_2U_2+P_3U_3+P_4U_4$
- iv. The best reply functions that emerge from this process often have a rather peculiar shape as illustrated below, although the functions will cross at some point.

C. Consider the equilibrium of what Buchanan calls the *Samaritan's dilemma* (see Buchanan, 1972, Tullock, 1983, or Rasmussen 1994, p. 68).

Samaritan's Dilemma		Bum Work	Loaf	
		Donor Aid	3 , 2	-1, 3
		No aid	-1, 1	0, 0

- i. Note that there is no pure strategy equilibrium to this game.
 - *Consequently, it would not be rational for either player to allow their behavior to be entirely predictable.*
- ii. To explore the possibility of a mixed strategy equilibrium, suppose that the probability of the Bum working is w and that the probability that the donor will aid the Bum is $g.$

iii. The Samaritan's expected payoff from participating in this game is:

- $U_S^E = g [3w + (1-w)(-1)] + (1-g) [(-1)w + (1-w)0]$
- Recall from statistics that the **joint probability** that **g** and **w** both happen is **gw** as long as g and w are chosen independently of one another.)

iv. Differentiating with respect to g and solving allows us to characterize g^*

- $[3w + (1-w)(-1)] - [(-1)w + (1-w)0] = 0$
- Note that g is not included in this equation!

v. This suggests that g does not have a simple "interior solution."

- Solving for w, using the first order condition $\Rightarrow [3w - 1 + w] + [w] = 0 \Rightarrow 5w = 1$
- Thus, we find that $dU^E/dg = 0$ for all g if $w^* = 0.2$
- Notice that if $w > 0.2$, then utility always increases with g, so g^* should be set at its highest value, 1
- And, if $w < 0.2$ then utility always falls with g, so g^* should be set at its lowest value, 0.
- The best reply function, thus equals 0 for $w < .2$, 1 for $w = .2$, and our player is indifferent among strategy choices when $w = .2$. [Draw this best reply function.]

vi. A similar calculation for the Bum yields

- $U_B^E = w [2g + (1-g)(1)] + (1-w) [3g + 0 (1-g)]$
- or $U_B^E = 2gw + w - gw + 3g - 3gw = -2gw + w + 3g$

vii. Differentiating with respect to g, we find that the first order condition for w^* :

- $-2g + 1 = 0$ which implies that $U_B^E = 0$ only if $g = 0.5$
- e. g. if the Samaritan sets $g = .5$, then any value of w will generate the same utility for the Bum.
- However, if $g > 0.5$ the Bum's utility falls as w increases, so w^* should be set at its lowest value, in which case $w^* = 0$.
- If the Samaritan sets $g < 0.5$, the utility rises as w increases, and w should be set at its highest value, in which case $w^* = 1$.
- Again the best reply function has a strange shape. It takes the value 1 for $g < 0.5$, 0 for $g > 0.5$ and the Bum is indifferent among strategies when $g = .5$.

viii. [Draw these "very strange looking" reaction functions, and note the Nash equilibrium.]

D. Note that if the Donor gives with probability 0.5 and the Bum works with probability 0.2, both are in equilibrium.

- With $g=0.5$ and $w=0.2$, both are on their "best reply function" and neither has a reason to change their mixed strategy.
- (See your diagram of the reaction functions.)
- No player can achieve a higher expected payoff by changing his or her strategy.
- Moreover, using any other strategy is likely to be worse, *if that alternative strategy is discovered by the opponent.*

- E.** One **peculiar property** of most **mixed strategy equilibria** is that at the equilibrium each player is *indifferent* among other (non-equilibrium) probabilities given the other's behavior, but *no other probability* combination generate an equilibrium.
- i. *Given* the behavior of other player(s), each player is indifferent among all strategies because they all generate the same expected payoff (expected utility)
 - (Recall the first order conditions above.)
 - ii. This equilibrium is also a bit demanding in that **it assumes that all players understand the entire game**, not simply their part of it.
 - This requires them to "know" the other player's strategies, payoffs, and method for making decisions in a setting of uncertainty.
 - At a Nash Equilibria in pure strategies, players do not have to understand the behavior of other players to realize that they should not change their strategies.
 - iii. In other words, achieving a "Mixed Equilibrium Strategy" requires a very sophisticated **consciously game theoretic mode of thinking** by the players (as in a Stackelberg model, rather than a Cournot model).
 - All players have to understand the nature of the game and anticipate the rational responses of other players.
 - To do this, they must know a lot about the other players.
 - [Note that if one of the players mistakenly plays a pure strategy, no one is made worse off, but if this pure strategy players is discovered the other players will want to change their strategies. Why? (Hint: think of the paper, rock scissors game.)]
 - iv. Of course, the same mode of reasoning applied to a game with a deterministic equilibrium implies that players will go directly to the Nash equilibrium.
 - Fortunately, in many cases the Nash equilibrium *may emerge* from a series of "moves" even if the players behave "myopicly" rather than fully rationally.
 - (Myopic players optimize only one move at a time, rather than taking account of all possible future moves of opponents.)
 - (See diagram of myopic adjustment in, for example, the lottery game.)
- F.** Our results demonstrate that **equilibrium mixed strategies can be computed by solving for a probability function that makes the other player(s) indifferent among their strategy choices.**
- [Student Puzzle: find the mixed strategy equilibria for a "zero sum" game; and for the Paper, Stone, Scissors game played by children. How do your choices of payoffs affect the "mixing probabilities?]
 - [Student puzzle: It can be said that a mixed strategy equilibrium depends on "random play" being "perfectly predictable" How is this possible?]
 -
 -
 -
 -

III. A Quick Review of Rational Choices in a Stochastic Environment

A. The mixed strategy equilibrium is only one of many settings in which choices must be made in an environment in which particular outcomes are not known at the moment choices have to be made.

- The simplest cases are coin tosses, rolling dice, and lotteries.
- The rent-seeking contests of the previous lecture are also examples.
- More complex cases arise whenever a long term effect has to be estimated in some way, as with future interest rates, weather, or political contests.

B. A Review of Probabilities and Probability Functions. One way to characterize a stochastic environment is with a **probability function** mapped over the "event space."

- i. A probability function has several important properties.
- ii. Every **possible outcome** is assigned a number between 0 and 1.
 - If the possible outcomes are $X_1, X_2, X_3 \dots X_N$ then a **probability function** defined over those possibilities looks like $P_i = f(X_i)$, for $i = 1, 2, \dots N$.
 - Probabilities cannot, by definition, be a negative number.
 - *Every impossible outcome is assigned the value 0.*
- iii. The sum of the probabilities, by definition, has to add up to 1.
 - $(\sum_{i=1}^N P_i) = 1$

C. The "**expected value**" of a random event such as X above is a technical statistical term that means "average" outcome.

- It is widely used in rational choice models and in applied work by statisticians.
- The expected value of X can be calculated as: $X^e = \sum_{i=1}^N P_i X_i$
- For example in rolling a dice, $P_i = 1/6$ for $i = 1, 2 \dots 6$ and $X_1 = 1, X_2 = 2, \dots X_6 = 6$
- $X^e = \sum_{i=1}^N P_i X_i = (1/6)(1) + (1/6)(2) + (1/6)(3) + (1/6)(4) + (1/6)(5) + (1/6) 6 = 3.5$
- (As this illustration illustrates, sometimes the expected value is actually impossible. However, if one calculated the average of a large number of dice throws (say 25 or more) the result would be number quite close to 3.5.)

D. Expected Utility. For many purposes, in rational choice theory only the "ordinal" properties of utility functions matter.

- That is to say, the fact that $U(A) = 2 * U(B)$ only means that A is preferred to B, not that A is twice as desirable as B.
- i. However, **to calculate expected utility** in the conventional ways, requires utility functions for which such algebra has meaning.
 - Only a subset of possible utility functions can be used to calculate "expected utilities."

- These are called Von Neuman Morgenstern utility functions.
- ii. **Von-Neuman Morgenstern** utility functions are bounded and continuous.
- Von-Neuman Morgenstern utility functions are "unique" up to a linear transformation.
 - Consequently, they are regarded by many economists and decision theorists to represent *quasi-cardinal* utility (in addition to ordinal utility).
 - V-N Morgenstern Utility functions have the property that if $U^e(A) > U^e(B)$, then A is preferred to B, where $U^e(A)$ is the "expected utility of stochastic setting A" and $U^e(B)$ is the expected utility of stochastic setting B.
- iii. The *expected utility* associated with such a stochastic setting is calculated as the expected value of utility:
- $$E(U(\mathbf{X})) = U^e = \sum_{i=1}^N P_i U(X_i)$$
- where the *value of the possible outcomes is now measured in utility terms*.
 - $E(U(V))$ is more **commonly written as $U^e(\mathbf{X})$**
- iv. It is important to remember that $U^e(\mathbf{X})$ is **not** generally equal to $U(\mathbf{X}^e)$.
- $U^e(\mathbf{X})$ is **the same as $U(\mathbf{X}^e)$ only if U is a linear** function of V.
 - In such cases, the individual is said to be **risk neutral**.
- v. DEF: An individual is said to be **risk averse** if the expected utility of some gamble or risk is less than the utility generated at the expected value (mean) of the variable being evaluated.
- A *risk neutral* individual is a person for whom the expected utility of a gamble (risky situation) and utility of the expected (mean) outcome are the same.
 - A risk preferring individual is one for whom the expected utility of a gamble is greater than the utility of the expected (mean) outcome.
- vi. The more risk averse a person is, the greater is the amount that he or she is willing to pay to avoid a risk.

E. A Von Neumann Morgenstern Utility function can be constructed as follows.

- i. Assume that X_{min} and X_{max} are the best and worse payoffs of the game.
- Recall, in a finite game, all the payoffs are finite, so X_{min} and X_{max} exist, and are often very easy to calculate.
- ii. The utility of X_i where $X_{min} < X_i < X_{max}$ can be calculated as a "convex combination" of the arbitrary values $U(X_{min})$ and $U(X_{max})$ where $U(X_{min}) < U(X_{max})$.
- Recall that a convex combination of X and Y can be written as $Z = (1-a) X + (a) Y$ where $0 < a < 1$.
 - Note that "a" can thus be interpreted as the probability of X, and Z as the expected value of getting X with probability "a" and Y with probability "1-a."
- iii. Note also that this "probability" can be varied until the individual of interest is indifferent between the gamble and X_i .
- The utility value of X_i is then defined to be:
 - $U(X_i) = p U(X_{min}) + (1-p) U(X_{max})$.

- (That is to say, assign utility numbers to events (X_i) by varying P until the person of interest is indifferent between the event (X_i) with certainty and the gamble over X_{\min} and X_{\max} defined by P .)
- iv. The Von-Neumann Morgenstern utility index was worked out in their very important game theory text *The Theory of Games and Economic Behavior*, 1944.
 - (Student puzzle: Why does the difference between ordinal and cardinal utility matter?)
 - (Student relief: for most of the rest of this course, we require only "ordinal" utility functions.)

IV. Practice Problems

- A.** Create a zero sum game 2x2 game using the payoffs of -1 and 1 in various orders.
- i. Does your game have a Nash equilibrium in deterministic strategies?
 - ii. Does your game have any Pareto optimal outcomes?
 - iii. If there is a deterministic equilibrium, change your payoffs to eliminate this equilibrium.
 - iv. Find the mixed strategy equilibrium for your game, and explain why it is an equilibrium.
- B.** Repeat "A" with a two person three strategy game.
- C.** Repeat "A" again using zero sum payoffs from -3 to +3.
- D.** Create a two stage game in which in stage one, players may send signals (credible commitments) to each other about whether they will cooperate or not in the PD game that follows.
- i. Assume that credible commitments are possible, but that creating them is costly.
 - ii. Find the equilibrium strategies in stages 1 and 2 of this game.
 - iii. Can communication in this type of setting solve the PD game?
 - iv. How costly does a credible commitment have to be to make a commitment to cooperate a poor strategy?
 - v. How might a game be designed to make "truth telling" an optimal strategy? (Discuss a couple of examples from ordinary markets regarding quality or price.)

V. Information Sets, Ignorance, and Game Theory

- A.** In the last section, we also mentioned that games in extended form can be used to represent simultaneous play, by making assumptions about what players know as the game unfolds.
- i. Specifically, if player 2 goes second, but does not know what the first player has done, it is as if player 2 is in a simultaneous game.
 - ii. (Ignorance of the other player's strategy choice is the most important consequence of simultaneous play in most parlor games.)
- B.** Formally, what we have done is to characterize the *information sets* of participants in the game.

- i. Rasmusen, p. 40, defines an *information set* as follows: Player i 's information set ω_i at any particular point of the game is the set of different nodes in the game tree that "he" knows might be the actual node, but between which he cannot distinguish by direct observation.
- ii. We represented this last time with dotted lines surrounding the nodes of the second player after the first had chosen to defect or cooperate in the PD game.

C. A related concept is a player's *information partition*.

- i. Again Rasmusen, p. 42, provides a useful definition: a player's information partition is a collection of his information sets such that: (1) each path is represented by one node in a single information set in the partition and (2) the predecessors of all nodes in a single information set are in one information set.
- ii. The information partition represents the different positions that the player knows he will be able to distinguish from each other at different stages of the game, thereby dividing up the set of all possible nodes into subsets called information sets.
- iii. If a player knows exactly what has happened at every stage in the game to the point of interest, the information partition is simply the set of all nodes that can be reached. E. g. each information set is a "singleton" consisting of the actual node itself.
- iv. The finer the informational partition, (smaller the information sets) the better is the player's information. (The player will better be able to distinguish between nodes in the game if the information sets are smaller.)
- v. Information is said to be *Common Knowledge* if it is known to all players, if each player knows that all the players know it, if each player knows that all the players know that all the players know it, and so forth *ad infinitum*.

VI. Bayes' Law, Signaling, and Learning Functions

- A. Learning in most games and in most economic and statistical models is represented as changes in a probability function that characterized *expected* outcomes or possibilities.
- i. For example, a consumer might have an general idea about the range of prices that he or she will have to pay for gasoline in Alaska (this would be called a *prior* probability distribution).
 - ii. Once you have actually purchased gasoline there, or done a bit of research, the probability function that you assign to Alaskan gasoline prices will change to reflect the new information that you have. You have learned something new about that distribution. The new distribution is called a *posterior* probability function.
 - iii. One can model this learning process in a general way as $F_1(x) = p(x | F_0, S)$ where F_1 is the conditional (*posterior*) density function of x given prior probability function F_0 , and S is some signal, observation, message, or event that causes the individual to change (or update) "his" priors. $F_1(x)$ becomes the new prior. So the next updated posterior would be $F_2(x) = p(x | F_1, S)$.
 - Generally as one becomes better informed, the probability function describing beliefs becomes "tighter" and converges toward the real value of x . That is to say, the variance of F generally falls with learning.
 - (On the other hand, occasionally, one learns just how ignorant one really is, in which case the variance of this probability function may increase.)
 - iv. *Diffuse priors* are often the initial starting point of this kind of analysis. With diffuse priors, one regards every possibility to be equally likely. That is to say $p(x_i)$, where $i = 1 \dots N$, is $1/N$ for all i .
 - v. Bayes law is derived from statistical definitions of joint probabilities:

- The joint probability of E and H being true is

$$P(E, H) = P(H|E) p(E) = P(E|H) p(H)$$
 - Using the last equality we can solve for $P(H|E)$ as:

$$P(H|E) = P(E|H) p(H) / p(E)$$
- vi. This is *Bayes' law* that relates the posterior probability $P(H|E)$ to the prior probability that H is true, $P(H)$, the conditional probability that E will be observed if H is true, $P(E|H)$, and the probability that E will be observed in any case, $P(E)$. (Varian, p. 192)
- Bayes' law is important because it shows how a rational individual *should* (rationally) update his priors as new information is received.
 - Since Bayes' law follows directly from statistical notions of joint probability, the law applies equally to internally consistent subjective and consistent objective probabilities.
- B.** In a signaling game individuals send messages to each other with the intent of changing the other player's expectations and thereby their behavior. In most models of signalling behavior, the signals / messages are the only control variable, or one of two control variables. (Illustration as a Nash Game: see class notes)
- i. The problem of *cheap talk* in such games addresses why should anyone believe a signal?
 - ii. Suppose there is a two stage game. In the first stage a person can signal his "intention," in the second which is a traditional PD game. (Either player may signal I "will " or "will not" cooperate in the PD game that follows.)
 - iii. If there is no penalty for lying there is no reason to believe the signals in the first round.
 - (Puzzle: Why then do experiments where people can talk yield more cooperation than those where players do not?)
- C.** A related game concerns **screening** people into different categories, such as productivity, in a setting where the "firm owner" knows less about the productivity of his or her potential employees than they do themselves.
- i. Here all potential employees attempt to signal the employer about their productivity, but in many cases they will have an incentive to exaggerate their productivity (as in team production settings).
 - ii. The challenge for a firm owner or manager is to devise a signaling game that actually produces useful information, or, alternatively, self selection into separate pools of productivities.
 - iii. For example, commissions for sales can be paid, rather than salaries. This induces persons who expect to be good salesmen to accept the commission and induces others to seek other positions in the firm or to shift to other firms.

VII. Homework Problems (Optional)

- A.** Create a zero sum game 2x2 game using the payoffs of -1 and 1 in various orders.
- i. Does your game have a Nash equilibrium in deterministic strategies?
 - ii. Does your game have any Pareto optimal outcomes?
 - iii. If there is a deterministic equilibrium, change your payoffs to eliminate this equilibrium.
 - iv. Find the mixed strategy equilibrium for your game, and explain why it is an equilibrium.

- B.** Create a Stackelberg signaling model in a PD game where one prisoner can attempt to persuade the other that he will cooperate. Use an abstract learning function and assume signals are costly, and that both players maximize expected utility.
- C.** Is a costly signal necessarily a creditable signal? Analyze and Discuss.
- D.** How might a game be designed to make "truth telling" an optimal strategy? (Discuss a couple of examples from ordinary markets regarding quality or price.) Analyze and Discuss.