Information and Equilibria in Economic Games

I. A good deal of economic analysis and model building implicitly ignores information problems. That is to say, the people modeled are assumed to know as much as can be known about the problem facing him.

A. 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.

B. This assumption can be weakend. One might assume that there are random phenomena that are well understood, and the persons modeled simply maximize the "expected" value of their objective functions given their various certain and probabilistic constraints. Optimal purchases of lottery tickets, investment in risky processes, rent-seeking, insurance markets etc. are usually modeled in this way.

C. However, in settings of imperfect information one has to explicitly model learning. In this setting 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 to the analyst.

D. This lecture provides an overview of some of the tools and concepts that are most applied in the economics literature on the "B" and "C" literatures.

II. Mixed Strategy Equilibria

A. In the last two lectures, we demonstrated that mixed strategy equlibria exist for games with finite numbers of strategies.

i. The ability to vary the probability of using various pure strategies makes the (expected) payoff function continuous which in turn allows fixed point theorems to be used to prove the existence of equilibrium strategies.

ii. Moreover, continuous expected payoff functions often allow ordinary optimization methods to determine *best probabilistic reply functions* and equilibria.

iii. Here a form of **unpredictability allows an equlibrium** to occur where certainty does not..

B. Consequently, any "finite" game that does not have a Nash Equilibrium in Pure Strategies may nonetheless have a Nash Equilibrium in Mixed strageies

C. Consider the equilibrium of what Tullock calls the *Samaritan's dilemma* (see 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

ii. There is no pure strategy equilibrium to this game. *Note also that it would not be rational for either player to allow their behavior to be entirely predictable.*

iii. 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.

iv. The donor's expected payoff from participating in this game is:

$$\Pi_D = g [3w + (1-w)(-1)] + (1-g) [(-1)w + (1-w)0]$$
$$\Pi_D = 3gw - (1-w)g - (1-g)w = 5gw - g - w$$

v. Differentiating with respect to g and solving allows us to characterize g* as *any g* that occurs when 5w - 1 = 0, e.g. when w = .2 (other wise g* = 0 or 1)

vi. A similar calculation for the Bum yields

$$\Pi_B = w [2g + (1-g)(1)] + (1-w) [3g + 0 (1- g)]$$
$$\Pi_B = 2gw + w - gw + 3g - 3gw = -2gw + w + 3g$$

vii. Differentiating we find that the first order condition will be satisfied for the Bum for *any* probability of working w* whenever -2g + 1 = 0

e. g. whenever g = .5 (otherwise w* = 0 if g>.5, or 1 if g < .5)

D. Note that if the Donor gives with probability .5 and the Bum works with probability .2, both are in equilibrium. Neither has a reason to change their mixed strategy.

i. Neither can do better than this combination. (Show figure of reaction functions.)

ii. Once a Nash equilibrium combination of probabilities is chosen, it is clearly an equilibrium. because no player can achieve a higher expected payoff by changing his or her strategy.

iii. (Moreover, using any other strategy is likely to be worse, if that alternative strategyis discovered by the opponent.)

E. One peculiar property of most mixed strategy equilibria is that at the equilibrium each player is *indifferent* among other (nonequilibiurm) probabilities given the other's behavior, but *none* of those other probabilities are an equilibrium.

i. That is to say, *given* the behavior of other players, one is indifferent among all strategies when all the others are playing their Nash Equilibrium Mixed Strategies.

ii. Unless one understands the game, *and* expects the opponents to change their behavior as soon as they notice that a deviation from equilibrium strategies, there is no reason to adopt the Mixed Strategy equilibrium strategies if everyone else has.

iii. So realizing a Mixed Equilibrium Strategy requires a mode of thinking (a bit more like a Stackelburg model) that differs from that used in non-stochastic Nasj games where, best reply functions are computed taking the other players' strategies to be *given*.

III. A Digression on Von Neumann Morgenstern Utility Functions

A. It is worth noting in passing that *expected utility functions* used in modern economic analysis of decisions under uncertainty do not simply represent "ordinal" relations. One has to be able to perform arithmetic operations on utilities and to use the results. The special "cardinal" utility functions used for these analyses are referred to as Von Neumann Morgenstern utility functions.

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

a. Assume that Xmin and Xmax are the best and worse things that can happen.. Recall, that normally, all the payoffs are known so Xmin and Xmax are often very easy to calculate.)

b. The utility of Xi where Xmin<Xi<Xmax can be calculated as a convex combination of the arbitrary values U(Xmin) and U(Xmax) where U(Xmin) < U(Xmax).

c. Vvary the probability, p, of getting Xmin which implies a probability (1- p) of getting Xmax until the individual states that "he" is indifferent between the gamble and Xi. The utility value of Xi is then defined to be: U(Xi) = p U(Xmin) + (1-p) U(Xmax).

d. (That is to say, assign utility numbers to events (Xi) by varying P until the person of interest is indifferent between the event (Xi) with certainty and the gamble over Xmin and Xmax defined by P.)

B. The Von-Neumann Morgenstern utility index was worked out in their very important game theory text *The Theory of Games and Economic Behavior*, 1944.

IV. Information Sets, Ignorance, and Game Theory

A. In the last lecture, 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*.

V. 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 $F_1(x) = p(x \mid F_o , S)$ where F_1 is the conditional (*posterior*) density function of x given prior probability function Fo, 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 $F_2(x) = p(x \mid F_1, S)$.

a. 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.

b. (On the other hand, occasionally, one learns jsut 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(xi), where i = 1 ..... N , is 1/N for all i.

v. Bayes law is derived from statistical definitions of joint probabilities:

a. The joint probability of E and H being true is

$$P(E, H) = P(H|E)\, p(E) = P(E|H)\, p(H)$$

b. 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)

a. Bayes' law is important because it shows how a rational individual *should* (rationally) update his priors as new information is received.

b. Since Bayes' law follows directly from statistical notions of joint probability, the law applies equally to internally consistent subjective and consistant 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? 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.) If there is no penalty for lying there is no reason to beleive the signals in the first round.

ii. (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. Here all potential employees attempt to signal the employer about their productivity. The challenge is to devise a signalling game that actually produces useful information, or, alternatively, self selection into separate pools of productivitys. (Example: commissions for salesmen)

VI. Problems

A. 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.

B. Is a costly signal necessarily a creditable signal? Analyze and Discuss.

C. 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.)