ECON/MGT 395 Behavioral Game Theory
Lecture title: "Introduction to (behavioral) game theory"

Dr. Angelino C. G. Viceisza
Spelman College
Economics

August 26, 2013

1 Acknowledgements

This note is created by the author and is in part based on the following sources:

- Chapters 1 and 2 of Dixit, A., S. Skeath, and D. Reiley. 2009. *Games of strategy.* Third Edition. New York, NY: W. W. Norton & Company.
- Gibbons, R. 1997. An Introduction to applicable game theory. *Journal of Economic Perspectives* 11 (4): 127-149.
- Camerer, Colin F. 1997. Progress in behavioral game theory. *Journal of Economic Perspectives* 11 (4): 167-188.

2 Motivation and day-to-day examples

- We play games of strategy (also called games) all the time. Apart from day-to-day (parlor) games such as tic-tac-toe, monopoly, words with friends, or sports, we also play strategic games with parents, siblings, friends, roommates, and so on.
- As such, having a proper understanding of what constitutes a game and how one can go about "strategizing" for it, is a useful skill to have. It is anticipated that most of us will play games for the rest of our life through different forms of social interaction such as with employers, employees, partners, children, and even children.
- From a disciplinary standpoint, games appear in many literatures including but not limited to: (1) economics (e.g. cooperation, coordination, competition), (2) business (e.g. management strategy and negotiation), (3) psychology (e.g. areas such as psychological game theory and discussions of people's beliefs), (4) political science (e.g. political competition, international negotiation and diplomacy), and (5) biology (e.g. evolutionary dynamics).
- Some concrete day-to-day examples of games:
  - *Tennis–Navratilova versus Evert*: Mixing one's plays as an example of mixed strategy equilibrium. Evert must introduce random variation in order to prevent perfect guessing by Navratilova. (chapter 7-8)

- *The GPA Rat Race*: Enrolled in a course that is graded on a curve (for lack of a better example, ECON/MGT 395). What is the best strategy? Suppose you all collude and agree to do equally badly; is that wise? Why yes; why not? This is the same dilemma that the prisoners face! (chapter 11)
- *Flat Tire*: Two friends were chemistry students at Duke and were doing well up to the final exam. The Saturday before the final, they went to a party at the University of Virginia. They overslept Sunday and got back too late to study for the exam on Monday. They agreed to tell the professor a sad story: 'We had a flat tire on the way back; could not get help; and now are too tired–could they have a makeup final?' The professor agreed to do so Tuesday. They studied all of Monday and showed up for the exam the next day. The professor placed them in separate rooms and handed a test to each of them. The first question on the first page, worth 10 points, was very easy. Each of them wrote a good answer, and greatly relieved, turned the page. It had just one question, worth 90 points. It was: "Which tire?"
  What is the issue? What would you do? Moral of the story?
  The point is not whether the choice of tire is obvious or logical, but whether it is obvious to the other that it is obvious to you that it is obvious to the other......This is a so-called *infinite regress*/convergence of expectations. Does a focal point for coordination exist? (chapter 4)
- *Others (mean professors; roommates and families; and dating)*: See pp. 10-14.

3 Some concepts and definitions

- Strategic thinking is essentially about interactions with others: someone else is also doing similar thinking at the same time and about the same situation.

- Behaving rationally is when you think carefully before you act–when you are aware of your objectives or preferences and of any limitations or constraints on your actions and you choose your actions in a calculated way to do the best according to your own criteria. In the context of ECON 242 or ECON 315, this typically meant that you maximized utility subject to constraints. Recall the so-called $MU = MC$ condition.

- Game theory adds another dimension to rational behavior–namely, interaction with equally rational decisionmakers. So, it is the 'science' of rational behavior in interactive situations. It was developed by John von Neumann and Oskar Morgenstern in their book called "The theory of games and economic behavior".

- Game theory has limitations. It does not give you an exact prescription of how to play all games (most games are too complex) and in particular, how to win. However, it is useful as a guide. Recall the purpose of models more broadly–they are intended to be an abstract representation of reality–they do not contain all elements.

- Behavioral game theory is about how real humans interact relative to what the pure game theory would predict. This is described by Camerer in his book called "Behavioral game theory" (the one that will be on course reserve) and also in the article "Progress in behavioral game theory" (see Moodle). See also the section 'Observation and experiment' in Games of Strategy (pp. 35-36).

- Behavioral economics is about incorporating psychological principles into economics models more generally. Behavioral game theory is a subset thereof; specifically, building psychology into game theoretic models.

- *Strategic games* or *games* describe interactions between mutually aware players versus *decisions* or *(individual) choices*, which describe action situations where each person can choose without concern for reaction or response from others.

  - *Simple rule: Unless there are two or more decisionmakers, each of whom responds to what others do (or what each thinks the others might do–beliefs), it is not a game.*

- A game G consists of the following components:

  1. A set of N consisting of n players, i.e. $N = \{1, 2, 3, ..., n-1, n\}$ (Think of this set N as {Briana, Charleen, Kibet, ..., Tasha, Thandiwe}).
  2. Each player has a set A_i consisting of m actions, i.e. $A_i = \{a_1, a_2, ..., a_{m-1}, a_m\}$ where i indicates the player (This depends on the game but for example in the tic-tac-toe example discussed in class, each player would have the set $A_i =${Circle, Cross}).
  3. Each player has a set S_i consisting of strategies that can be formulated based on her and others' actions. The strategy space depends on the type of game in question. But, loosely speaking, a strategy is a *complete plan of action*. I.e. for most games, strategies are *not* the same as actions. They are more complicated. We will elaborate further as we go along.
  4. Each possible outcome of the game has an associated payoff for each player. So, there is a set of payoffs P_i which depends on what players do in the game. Two important aspects need to be understood about payoffs:
     - The payoffs for one player capture everything in the outcomes of the game that she cares about (we will come back to this when we read the article "Testing game theory" by Jorgen Weibull).
     - They may be calculated based on an "expected" manner if there is risk. For example, if in one player's ranking, outcome A has payoff 0 and outcome B has payoff 100, then the prospect of a 75% probability of A and 25% probability of B should have the payoff $0.75 \times 0 + 0.25 \times 100 = 25$.
  5. Information (I) and common knowledge of the above rules of the game. To understand what common knowledge means, suppose the game is being played between two players: 1 and 2. Then, common knowledge implies that 1 and 2 know the rules of the game, but also that 1 knows that 2 knows and that 2 knows that 1 knows. Taking it further it also means that 1 knows that 2 knows that 1 knows and that 2 knows that 1 knows that 2 knows. And so on...again, it is an infinite regress of common knowledge.
  6. So, in a sense a game can be described as $G = \{S, P, A, I, N\}$.
  7. Furthermore, it is often assumed that all players are rational and that there is common knowledge of rationality.
  8. Finally, when all these rational players interact, the typical solution to the game is one of equilibrium. This means that each player is using the strategy that is the best response to the strategies of the other players. There are also non-equilibrium solutions to games (e.g. level-k thinking in guessing games and backward induction). There are also evolutionary approaches/solutions to games.

- Uses of game theory: (1) explanation, (2) prediction, and (3) advice or prescription.

4 Classifying games

- This section is based on Games of Strategy pp. 20-27 and Gibbons (1997).
- Games of strategy can be categorized according to different dimensions. Some of these dimensions are:
  1. The timing of moves. If moves are taking place *simultaneously*, each player must anticipate what the other players will do and recognize that they are doing the same. If moves are taking place *sequentially*, each player must think of the future consequences of moves.
  2. Limitations of information. (1) A player may not know all the information that is pertinent for the choice that she has to make at every point in the game; this could be due to external or strategic uncertainty. If a game has neither form of uncertainty, we say the game is one of perfect information; otherwise, the game has imperfect information. (2) When one player knows more than another does; this is called incomplete or asymmetric information. Otherwise, it is complete or symmetric information. This leads to situations of signaling and screening.
  3. Whether players are in total conflict or there is some commonality. This gives rise to zero/constant-sum games or not.
  4. Whether rules are fixed or manipulable. This gives rise to so-called pregames as opposed to standard games.
  5. Whether agreements are enforceable. If they are, we have so-called cooperative games; otherwise, we have noncooperative games. Most of what we will deal with in this course will be focused on noncooperative game theory.
  6. Repetition and fixed or flexible partnership. This can give rise to reputational aspects.
- Gibbons (1997) in particular uses 1 and 2 above to distinguish between (1) static games with complete information (with Nash equilibrium as the solution concept); (2) dynamic games with complete information (with backward induction as the solution concept); (3) static games with incomplete information (with Bayesian Nash equilibrium as the solution concept); and (4) dynamic games with incomplete information (with perfect Bayesian equilibrium as the solution concept).
- Food for thought questions:
  1. What examples does Gibbons discuss of these different games?
  2. What is the punchline of Camerer's article?
  3. How do these two relate?
  4. What does Camerer mean by "games as social allocations"?

ECON/MGT 395 Behavioral Game Theory
Lecture title: "Simultaneous-move Games"

Dr. Angelino C. G. Viceisza
Spelman College
Economics

September 11, 2013

1 Acknowledgements

This note is created by the author and is in part based on the following sources:

- The course textbook–Games of strategy by Dixit et al; chapters 4 and 5.
- Parts of Behavioral game theory by Camerer such as appendix A1.1 on basic game theory and chapter 7 on coordination (this is related to the student presentations)
- The article by Gibbons–in particular, the section on static games of complete information.

2 Simultaneous-move games with discrete action sets (GoS, chapter 4 and other references throughout)

- Recall from the previous note and several example discussed thus far, that games are said to have simultaneous moves if players must move without knowledge of what their rivals have chosen to do. It is clearly so if players choose their actions at the same time.
- A game is also simultaneous when players choose their actions in isolation, with no information about what other players have done or will do, even if the choices are made at different hours of the clock. For this reason, simultaneous-move games have *imperfect* information in the sense defined previously (recall lecture note and chapter 2 of GoS; a game has imperfect information if it has either external or strategic uncertainty or both).
- Many familiar strategic situations can be described as simultaneous-move games: producers designing product features without knowing what competitors are doing; individual voters casting ballots without knowing what other voters are doing; and so on.
- Recall that a strategy is a complete plan of action.
- In a purely simultaneous-move game, each player can have at most one opportunity to act, although that action can have many component parts; if a player had multiple opportunities to act, there would be an element of sequentiality. Therefore, there is no real distinction between strategy and action in simultaneous-move games, and the terms are often used as synonyms in this context.

- There is one complication however. A strategy can be a probabilistic choice from the basic actions initially specified. Recall the tennis example discussed previously–Navratilova might have changed aims to confuse her opponent; otherwise, she would be too predictable. Such probabilistic strategies are called *mixed* strategies to distinguish from *pure* strategies (really the pure actions in simultaneous-move games). Chapters 4 and 5 look at pure strategies; chapters 7 and 8 look at mixed strategies.

- To analyze simultaneous games, we need to ask/consider how players choose their actions/ strategies. In the previous note, we saw two such approaches: (1) iterated elimination of dominated strategies (IEDS) and (2) Nash equilibrium (NE) solution.

- The principle way to solve simultaneous-move games is by finding NE.[1] As a reminder, a NE in a game is a list of strategies, one for each player, such that no player can get a better payoff by switching to some other strategy that is available to her while all the other players adhere to the strategies specified for them in the list. It is found by asking whether a player's strategy is a *best response* given any of the other players' strategies. In games that are small enough, it can be found by means of cell-by-cell inspection, or enumeration. Other approaches such as IEDS sometimes give the same solution as NE. Recall the example discussed in the previous note: (Up,Middle) was a NE but it was also the solution that emerged from iterated elimination of dominated strategies.

- Let us consider another classic example for which IEDS leads to the same solution as NE: the prisoner's dilemma. For a detailed description of this game, see GoS, chapter 4, pp. 97-98; Gibbons (1997), p. 131; BGT by Camerer, appendix A1.1 and chapter 2.

- Example. The Prisoner's dilemma (payoffs are years to be served).

		Player 2	
		Confess (Defect)	Deny (Cooperate)
Player 1	Confess (Defect)	(10,10)	(1,25)
	Deny (Cooperate)	(25,1)	(3,3)

- How do we solve the game?
  NE: What is player 1's best response? If player 2 confesses, player 1's best response is to confess also. If player 2 denies, player 1's best response is to still to confess. The same logic applies to player 2's best response. So, (Confess,Confess) is a NE.
  IEDS: Well, note that Deny is a *dominated* strategy for both players while Confess is a *dominant* strategy. So, each player should eliminate Deny as a plausible/rational strategy. As such, (Confess,Confess) emerges from the process of IEDS.

- Any game with the same general payoff pattern as that illustrated in the above game is typically given the generic label "prisoner's dilemma". It has three essential features: (1) each player has two strategies: to cooperate with one's rival or defect from cooperation; (2) each player also has a dominant strategy (defect/confess); and (3) its dominance solution equilibrium is worse for both players than the nonequilibrium situation in which each plays the dominated strategy (cooperate/deny). Games of this type have a wide range of applicability. They are used to model "free-riding" behavior since each player has an incentive to

[1] This solution concept is named after John Nash. He also won the Nobel Prize in economics in 1994 for his work on games. The movie "A beautiful mind" featuring Russell Crowe is based on a biography of Nash.

cheat/confess/defect. They are also used to model how people contribute/behave with regard to public goods. So, sometimes people will use the term public goods games to indicate prisoner's dilemma games. While they are not the same, they are related.

- IEDS can be limited. Many simultaneous-move games have no dominant strategies and no dominated strategies. Others may have one or several dominated strategies, but IEDS will not yield a unique outcome/solution to the game. In such cases, we need a next step in the process of finding a solution to the game.

- Best-response analysis, as applied in the example in the previous note and in the prisoner's dilemma, is a systematic method for finding NE. It is a comprehensive way of locating all possible NE of a game. Let us consider another game from GoS, p. 104.

- Example. GoS, figure 4.1 and 4.7.

		Column		
		Left	Middle	Right
Row	Top	(3,1)	(2,3)	(10,2)
	High	(4,5)	(3,0)	(6,4)
	Low	(2,2)	(5,4)	(12,3)
	Bottom	(5,6)	(4,5)	(9,7)

- How do we solve the game using best-response analysis?

- Well, suppose the column player plays Left, what is the row player's best response? To play Bottom. If she plays Middle? To play Low. If she plays Right? Low also. How about the column player's best response? If row plays Top, the column player should play Middle. High, Left. Low, Middle. Bottom, Right. Low, Middle is a best response for both players! So, this is the NE.

- The games discussed thus far have had one unique NE. This need not always be the case. A game can have multiple equilibria. We can illustrate this using a class of games called "coordination games". See also BGT, chapter 7, in particular sections 7.1 and 7.2 (recall student presentations) and Gibbons, pp. 132-133 (the dating game and matching pennies).

- The flat tire example was a coordination game–the students could have matched on four possible tires. What is the normal (or strategic) form representation of the game?

- Example. Flat tire example in normal/strategic form (payoff is full exam score assuming both students got 10 points for question 1)

		Student B			
		LF	RF	LB	RB
Student A	LF	(100,100)	(10,10)	(10,10)	(10,10)
	RF	(10,10)	(100,100)	(10,10)	(10,10)
	LB	(10,10)	(10,10)	(100,100)	(10,10)
	RB	(10,10)	(10,10)	(10,10)	(100,100)

This game is also said to be one of *pure coordination* because it has multiple Nash equilibria for which it does not matter which is achieved. All that matters is that the two students manage to coordinate on one of the tires! This is the same as the "Will Harry Meet Sally" example on pp. 111-113 of GoS. Is there a focal point for coordination?

- Not all coordination games are "symmetric" in payoffs. Assurance or stag-hunt games also require coordination, but if players manage to coordinate on one particular strategy/action, they get higher payoffs than another strategy/action. See pp. 113-114. Again, there are multiple equilibria but one is superior in that it leads to higher payoffs. Will this be achieved? Not necessarily. It depends on common knowledge and convergence of expectations. Much of the lack of coordination can be achieved through communication.

- *Activities/Discussions: (1) classroom stag-hunt experiment and (2) relation to Senegal research (with cheap talk).*

		other player	
		yellow	green
you	yellow	(1,1)	(1,0)
	green	(0,1)	(3,3)

- For this particular game, also see the discussion in BGT, section 7.4.

3 Simultaneous-move games with continuous action sets (GoS, chapter 5)

- Many simultaneous-move games differ from those considered so far; they entail players choosing strategies from a wide range of possibilities. Examples include games in which firms choose prices or quantities for their products; philanthropists choose donation amounts; and contractors choose project bid levels. Technically, prices, quantities and dollar amounts do have a minimum unit, such as a cent, but having to account for those is handled more easily by regarding such choices as continuously variable real numbers. When games have such large action/strategy spaces/sets, game tables/payoff matrices become useless as analytical tools. So, we need a different way to approach such games.

- Basically, due to the continuous action set, we can now use algebraic formulations (functional forms) to express the relationship between strategies/actions and payoffs. We will also be in a position to take derivatives. Nonetheless, we still seek to characterize best responses.

- Let us consider an example regarding price competition in oligopoly markets (aks Bertrand model; recall ECON 242 and/or 315). Suppose two firms–X and Y–need to set their respective menus and get them printed without knowing the other's price. In setting its price, each firm has to calculate the consequences for its profit. For simplicity, suppose the two restaurants are symmetric such that the cost of serving each customer is \$8. Suppose further that experience or market surveys have shown that price and quantity are related as follows:

$$Q_x = 44 - 2P_x + P_y \tag{1}$$
$$Q_y = 44 - 2P_y + P_x \tag{2}$$

X's profit per week can be calculated as:

$$\Pi_x = (P_x - 8)Q_x = (P_x - 8)(44 - 2P_x + P_y) = -8(44 + P_y) + (60 + P_y)P_x - 2(P_x)^2. \tag{3}$$

Since X sets its price P_x to maximize the payoff in equation 3 for any given level of Y's price, we can treat P_y as a constant. Recognizing this equation 3 can be seen as a polynomial of order two. Recall the so-called "abc formula", which allows you to solve for the roots of a polynomial (quadratic). For this purpose, recognize that equation 3 takes the following form:

$$\Pi_x = a + bP_x + c(P_x)^2 \tag{4}$$

where $a = -8(44 + P_y)$, $b = (60 + P_y)$ and $c = -2$.
Alternatively, we can use calculus as we saw in math basics. X will seek to maximize the function in equation 3 with respect to P_x. To do so, it will take the derivative with respect to P_x and set it equal to zero. It will also verify second-order conditions to make sure it is at a maximum profit and not a minimum. Doing so, we get:

$$\frac{d\Pi_x}{dP_x} = (60 + P_y) - 4P_x = 0, \tag{5}$$

which gives

$$P_x = 15 + 0.25P_y. \tag{6}$$

Equation 6 gives X's best response function. Similarly, due to symmetry, Y's best response function is $P_y = 15 + 0.25P_x$. The intersection of these two best response functions is the NE, see pp. 137 of GoS.

4 Related issues/sections you should read/check on your own...

- NE as a system of beliefs and choices, GoS, pp. 95-96.
- Minimax method for zero/constant-sum games, GoS, pp. 106-107.
- Extending to three or more players, GoS, pp. 108-111.
- See other examples of coordination games: (a) Battle of the sexes (GoS, pp. 115-116 AND BGT, section 7.2), (b) Game of chicken (GoS, pp. 116-118), (c) Matching pennies (Gibbons, pp. 133 and BGT, chapter 7).
- No equilibrium in pure strategies, GoS, pp. 118-120.
- Other types of best-response functions in continuous contexts, GoS, pp. 138-142.
- Empirical evidence concerning NE. Make sure to relate this to the chapters from BGT, chapter 7, 2, 4. These readings talk about whether real people actually behave according to equilibrium. We will see more of this in the second and third parts of the course on behavioral and experiments. As a result of failures of NE to predict how people *actually* behave, (behavioral) game theoretists have criticized the solution concept and started seeking for alternatives to rationalize people's behavior in strategic contexts.
- Rationalizability has been proposed as another solution concept that can carry us all the way to a unique NE in a game, see GoS, pp. 157-162. *Rationalizability* identifies strategies that survive elimination after having ruled out those strategies that are *never a best response*. It is a stronger refinement than elimination of dominated strategies. Time permitting, we may discuss an example or two further down the road.
- Reminder: Read section 7.2 and skim other sections of the chapter, in particular section 7.4.

ECON/MGT 395 Behavioral Game Theory
Lecture title: "Sequential-move Games"

Dr. Angelino C. G. Viceisza
Spelman College
Economics

September 11, 2013

1 Acknowledgements

This note is created by the author and is in part based on the following sources:

- The course textbook–Games of strategy by Dixit et al; chapters 3 and 6.
- Parts of Behavioral game theory by Camerer such as chapter 2 on dictator, trust, and ultimatum games and chapter 4 on bargaining games (this is related to the student presentations).
- The article by Gibbons–in particular, the section on dynamic games of complete information.

2 Sequential-move games (GoS, chapter 3)

- Sequential-move games entail strategic situations in which there is a strict order of play. Players take turns making their moves, and they know what players who have gone before them have done.
- In the terminology of Gibbons (p. 133), the timing of such a game is as follows:
  1. Player 1 chooses an action a_1 from a set of feasible actions A_1.
  2. Player 2 observes 1's choice and then chooses an action a_2 from a set of feasible actions A_2.
  3. After the players choose their actions, they receive payoffs $u_1(a_1, a_2)$ and $u_2(a_1, a_2)$ respectively.
- To play well in such a game, players must use a particular type of interactive thinking. Whenever actions are taken, players need to think about how their current actions will influence future actions, both for their rivals and for themselves. Players thus decide their current moves on the basis of calculations of future consequences.
- Most actual games combine aspects of both sequential- and simultaneous-move situations, but the concepts and methods of analysis are more easily understood if they are first developed separately for the two pure cases. Chapter 6 and parts of chapters 7 and 8 discuss combined games.

- A graphical technique for displaying and analyzing sequential-move games is based on game trees. This tree refers to the extensive form representation of a game.

- An example is the game tree in Figure 3.1 on page 49 of GoS. The rules of the game gives the firm move to Ann; this is shown at the leftmost point, or *node*, which is called the initial node or root of the game tree. At this action or decision node, Ann has two actions available to her: "Stop" or "Go". If Ann chooses Stop, then Bob has the move; if she chooses Go, then Chris has the move, and so on. Once all players have made their decisions and have collectively followed a certain "path of play", the terminal node indicates the payoffs for each player as components of a payoff vector.

- Notice that this game tree also contains *external uncertainty*, which is handled by introducing an outside player called "Nature". This is thus a so-called move of nature.

- Similarly to before, a strategy is a complete plan of action. Unlike simultaneous-move games in which strategies and actions were one and the same, in sequential-move games a player should plan ahead for each possible move that another player can make.

- When solving games by using trees, we often times use the so-called concept of *rollback* or *backward induction*. This is typically used when a sequential-move game either has a finite set of moves or is combined with a simultaneous-move game. To illustrate, consider the example discussed in GoS of a teenager named Carmen who is deciding whether to smoke. How can this be a game? Well, suppose that there is Today's Carmen and a Future Carmen. Basically, Future Carmen different from Today's Carmen in that the former is a grown-up Carmen whose tastes and preferences are different. Thus, Today's Carmen is playing a sequential game with Future Carmen. Today's Carmen wants to try smoking for a bit, but then give it up. Future Carmen is Today's Carmen after it has smoked some and is likely to become addicted. So, the following game tree can be constructed:

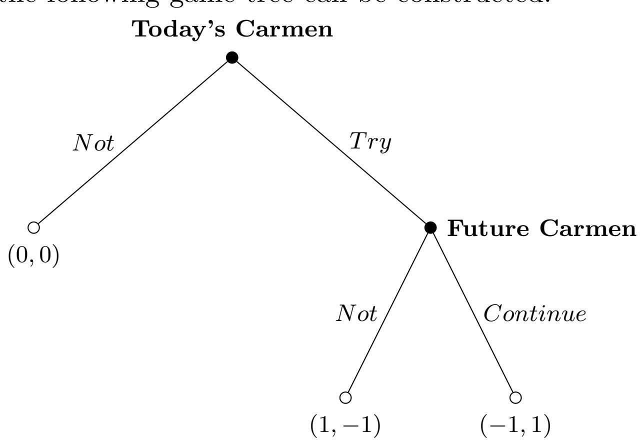

- What would rollback or backward induction have to say about how Today's Carmen and Future Carmen should play this game? Well, according to backward induction, we ask: At the end of the day, so to speak, what would the final player do? In this game, what would Future Carmen do? Well, she would prefer to continue smoking than not. Given this, the game that Today's Carmen is really playing is the following:

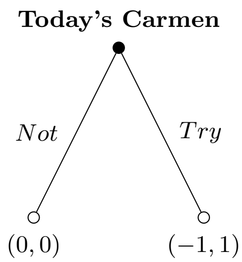

- Given this, what should Today's Carmen do? Choose Not. That is, she should never start smoking. This process is referred to as rollback or backward induction.

- When all players choose their optimal strategies found by doing rollback analysis/backward induction, we call this set of strategies the *rollback equilibrium outcome*. Game theory predicts this as the equilibrium of a sequential game when all players are rational calculators in pursuit of their respective best payoffs.

- All finite sequential-move games discussed in GoS have at least one rollback equilibrium–in fact, most have exactly one. In the aforementioned game, the rollback/backward induction equilibrium is where Today's Carmen chooses Not and Future Carmen chooses the strategy Continue. When Today's Carmen takes her optimal strategy, the addicted Future Carmen does not come into being at all and therefore gets no actual opportunity to move. But Future Carmen's shadowy presence and the strategy that she would choose if Today's Carmen chose Try are important parts of the game. If that so-called threat from Future Carmen did not exist, Today's Carmen would play the game differently!

3 Combined games (GoS, chapter 6)

- Three key take-aways:

  1. Simultaneous-move games and sequential-move games can be combined. This is the main focus of chapter 6.
  2. For purposes of this more general class of games, we use the so-called subgame perfect Nash equilibrium (SPNE). We consider a so-called smaller game, known as "subgame". We find the solution of such a game and use backward induction to solve the remainder of the game. In the aforementioned smoking example, the smaller game that Future Carmen plays could be seen as a simple subgame. But, the subgame can also be more involved, for example, it can be a simultaneous-move game in itself.
  3. An important aspect that is discussed by Gibbons (p. 134) and further deepened by GoS (pp. 185-199) is that simultaneous-move games need NOT be represented in normal form and sequential-move games need NOT be represented in extensive form. In other words, sequential games can be represented in normal form and simultaneous games can be represented by extensive forms. Using the appropriate solution concepts, we should still get the same equilibrium strategies. See for example Gibbons on the trust game.

4 Related issues/sections you should read/check on your own...

- GoS, Adding more players, pp. 57-62.

- GoS, First-mover advantage does not always exist, p. 62-63.
- GoS, Adding more moves with examples of day-to-day sequential games (tic-tac-toe, chess, checkers).
- GoS, Evidence concerning rollback/backward induction, pp. 71-74.
  - Link this to the article by Weibull on "Testing Game Theory" that will be discussed around 9/19-9/24.
  - Link this to BGT/Camerer's discussions of how real people play sequential-move games– link it to the student presentations!
- Some common sequential-move games that are used in behavioral/experimental economics: (1) trust game, (2) dictator game, (3) ultimatum game, (4) bargaining game which can be related to (3) and the so-called centipede game. Gibbons, GoS, BGT (in particular the student presentations) will all highlight these different games. Make sure to make the links!

ECON/MGT 395 – Guessing Game, Strategic Thinking, and Nonequilibrium

Dr. Angelino C. G. Viceisza

Spelman College

September 18, 2013

- Instructions
  - Everyone has to simultaneously pick a number from 0 to 100 (i.e. 0, 1, 2, 3, ..., 98, 99, 100).
  - Once you pick your number, write it on a piece of paper and submit it to me.
  - I will collect all the numbers and calculate the average.
  - The winner(s) of the "contest" is (are) the person(s) whose number is (are) closest to 2/3 times the average of all numbers submitted.
  - Is this clear? (no talking, no interaction, this is an individual game)

- And the winner is...
- Why?
  - Suppose you reason that everyone else chooses randomly from [0,100]. The average will be around 50. So, you should pick 2/3*50=33.33.
  - However, if you reason that way, then others might too. So, you should pick 2/3*33.33=22.22.
  - However, if you reason that way, then others might too. So, you should pick 2/3*22.22=14.81.
  - And so on...
  - So, the Nash equilibrium prediction (or what results from iterated elimination of dominated strategies) is zero...

- This game has been called the "beauty contest" after Keynes who proposed the following scenario:
  - Suppose a fictional newspaper contest is conducted in which people are asked to pick the six most beautiful faces from a series of photos.
  - The winner is s/he who chooses the most beautiful faces according to everyone (her/himself included).
  - One strategy is to choose the faces that you think are most beautiful, but a more "rational" or perhaps "strategic" view is to pick the faces you think others will find most beautiful.
  - In principle, there can be an infinite regress of reasoning.
  - Recall the candidate for President matching game (similar idea).
- The game has been used extensively to study strategic reasoning (e.g., Crawford, Irriberi, Costa Gomez, Nagel).
- Typically, it is found that people correspond to different levels of strategic thinking. This has led to level-k thinking models.
  - Level-0: Those who choose randomly from [0,100].
  - Level-1: Those who reason that all others are level-0. Thus, the average would be around 50 and therefore, they choose 2/3*50=33.33.
  - Level-2: Those who reason that all others are level-1. Thus, they choose 2/3*33.33=22.22.
  - And so on...

ECON/MGT 395 Behavioral Game Theory
Lecture title: "Time and Discounting"

Dr. Angelino C. G. Viceisza
Spelman College
Economics

October 8, 2013

1 Acknowledgements

This note is created by the author and is in part based on the following sources:

- The articles by Laibson (1997) and Frederick et al. (2002) [see syllabus and Moodle].
- See also the articles for the student presentations by Andersen et al. (Econometrica, 2008) and Ashraf et al. (QJE, 2006).

2 Time

- Recall ECON 242 (Principles of Micro), utility functions are introduced over one's own goods. However, as this and other courses show, utility functions can be defined over other arguments. We already saw some applications hereof in this course: (1) social preferences (utility functions not just defined over one's own goods/payoffs, but also over those of others) and (2) risk/uncertainty (utility functions or value functions not just defined over goods, but also over money/terminal wealth; furthermore, also defined over losses in addition to gains–recall prospect theory).
- This part of the course introduces yet another dimension over which we can define utility: TIME. How so? Well, many decisions are made over time or with the future in mind. Recall the game played between Present and Future Carmen. We were silent about whether or not Carmen's future utility needs to be discounted to today. Why may that be the case? Because preferences/conditions may change over time; also, there is the so-called related concept of time value of money (related to interest rates in finance/macro). So, 10 units of utility in the future may not be the same as 10 units of utility today to Carmen!
- As Andersen et al. (2008) indicate in their opening sentence: *"Utility functions are characterized in three dimensions, reflecting preferences over goods, time, and uncertainty."* ECON 242 takes care of part of the goods component. Since goods can broadly be defined to include other people's goods/utility, this course (ECON/MGT 395) takes care of the other part. Furthermore, it also takes care of uncertainty (risk; see previous note) and time (this note).
- The basic model for time goes back to Samuelson's discounted utility (DU) model proposed in 1937. Notice that Frederick et al. distinguish between time discounting and time preferences.

Time discounting broadly refers to any reason for caring less about a future consequence, including factors that diminish the expected utility generated by a future consequence, such as uncertainty or changing tastes. Meanwhile, *time preference* more specifically refers to the preference for immediate utility over delayed utility.

- The DU model specifies a decisionmaker's intertemporal preferences over consumption profiles $(c_t, ..., c_T)$. Under the usual assumptions, such preferences can be represented by an intertemporal utility function $U^t(c_t, ..., c_T)$. The model goes further by assuming that a person's intertemporal utility function can be described by the following special functional form:

$$U^t(c_t, ..., c_T) = \Sigma_{k=0}^{T-t} D(k) u(c_{t+k}) \tag{1}$$

where $D(k) = \left(\frac{1}{1+\rho}\right)^k$.

- Some aspects to note:

  1. In this formulation,
     (a) $u(c_{t+k})$ is often interpreted as the person's cardinal instantaneous utility function–her wellbeing in period $t + k$
     (b) $D(k)$ is often interpreted as the person's *discount* function–the relative weight she attaches, in period t, to her wellbeing in period $t + k$
     (c) ρ represents the individual's pure rate of time preference (her discount rate), which is meant to reflect the collective effects of certain "psychological" motives (see Section 2 of Frederick et al. for more detail).
  2. Section 3.1 through 3.7 of Frederick et al. discuss several important features of the DU model as it is typically used by economists. One such aspect, which has been challenged by the behavioral literature, is so-called *constant discounting* and *time consistency*. A more general discount function could be represented as $D(k) = \Pi_{n=0}^{k-1}\left(\frac{1}{1+\rho_n}\right)$, where ρ_n represents the per-period discount rate for period n. So, by assuming–as we did above–that $D(k) = \left(\frac{1}{1+\rho}\right)^k$, we are assuming a constant per-period discount rate. This means that delaying or accelerating two dates outcomes by a common amount should not change the preferences between outcomes. This also means that a person's intertemporal preferences are *time-consistent*, which means that later preferences "confirm" earlier preferences.

3 Miscellaneous aspects

- *Measurement*: The DU model assumes that a person's time preference can be captured by a single discount rate, ρ. Over the past four decades, there have been many attempts to measure this rate. Some of these studies are included in Frederick et al. Table 1 (pp. 378–379). Some recent studies are: Andersen et al. (2008; see reading list, Moodle, and one of the student presentations) and Andreoni and Sprenger (two papers in *The American Economic Review* in 2012).

- These empirical studies have suggested several aspects/features (sometimes also referred to as *anomalies*) of time discounting and preferences:

1. *Hyperbolic discounting*: This term is used to indicate that a person has a declining rate of time preference, in our notation ρ_n is declining in n, contrary to what constant discounting/time consistency would claim. See the paper by Laibson (1997) which is on the reading list.
2. *Gains are discounted more than losses (sign effect)*: This means that people prefer to incur a loss sooner than delay it relative to a gain. *Does this strike you as plausible? Does it remind you of anything we saw previously?*
3. *Small amounts are discounted more than large amounts (magnitude effect)*: Stakes matter for discounting. People prefer to incur larger amounts sooner rather than later relative to smaller amounts.
4. *Preference for improving sequences*: This means that in choices over sequences of outcomes, improving sequences are preferred to declining sequences. Notice that the assumption of positive time preference would dictate the opposite.
5. *Violations of independence*: Independent (irrelevant) alternatives impact people's time preference.

- One question is whether these anomalies can be considered mistakes. The answer depends on who you ask, but it is safe to conclude that not necessarily.

- These empirical findings have suggested several alternative models and ways for eliciting time preferences:

  1. Time preferences should remind you of our discussion of risk preferences! Indeed, they are theoretically related. Andersen et al. (2008) in particular argues for *joint elicitation of risk and time preferences.*
  2. Alternative models (see Section 5 of Frederick et al. for more detail):
     (a) Models of hyperbolic discounting: The so-called (β, δ) model or formulation allows for exponential (constant) discounting and hyperbolic discounting. These models also allows us to explain key issues such as procrastination and lack of self-control/commitment (see also Laibson 1997 on reading list). A subset of these models are models of self-awareness in which a person with time-inconsistent preferences may or may not be aware of this.
     (b) Models that enrich the instantaneous utility function: Among these are habit formation models; reference-point/reference-dependent models that incorporate ideas from prospect theory (departures from status quo); models of anticipated utility; and models of visceral influences such as hunger, sexual desire, and so on. The latter class of models also give room for neuroeconomics models.
     (c) Models that represent more radical departures from the DU model: Models of projection bias (people mispredict how their tastes will change over time–they are boundedly rational); models of mental accounting (not all money is treated as fungible and we may separate funds into different "mental accounts"); multiple (or dual) self models (recall Present and Future Carmen; recall also split-personality question on the exam); and models of temptation (disutility is experienced from not choosing the most enjoyable option now).

- Bottom line: These alternative set of models can explain a wide range of impulsive choices and other self-control phenomena. One such discussion is Laibson (1997) and Ashraf et al.

(2006; also see Moodle and one of the student presentations). These concepts have many applications to how we make decisions; e.g. health (incentives to exercise or signing up for health insurance), finance (retirement planning or investment decisions), crime (current versus future opportunities, and so on.

ECON/MGT 395 Behavioral Game Theory
Lecture title: "Uncertainty and Information"

Dr. Angelino C. G. Viceisza
Spelman College
Economics

September 25, 2013

1 Acknowledgements

This note is created by the author and is in part based on the following sources:

- The course textbook–Games of strategy by Dixit et al; chapters 6, 7, and 9.
- The article by Gibbons–in particular, the sections on incomplete information.
- The articles on 'prospect theory' by Kahneman and Tversky (1979) and Barberis (2013) [see syllabus and Moodle].

2 Information (GoS, chapters 6 and 9)

- At the beginning of the course/in chapter 2, we discussed different ways in which uncertainty can arise in a game–external and strategic–and ways in which players can have limited information about aspects of the game–imperfect or incomplete, symmetric or asymmetric. If you recall, we said that if a game has either strategic uncertainty, external uncertainty, or both, it is said to have *imperfect* information. Furthermore, we said that if one player has more information in a game than other players, this game is said to have asymmetric/incomplete information.
- This chapter discusses "uncertainty and information" in greater detail. With this being said, we have already encountered/analyzed some of these cases:

    1. In simultaneous-move games (chapters 4 and 5; see also the lecture notes), each player faced strategic uncertainty and therefore, these were games of imperfect information. Namely, each player took his/her action without knowing the action(s) that other players were taking. Notice that while the game has imperfect information, it is still a game of *complete* information (as discussed by Gibbons). This imperfect information can be further illustrated by representing the simultaneous-move game in extensive form. This goes back to an aspect we alluded to in the note on sequential-move games and that is also discussed in chapter 6, pp. 193-195. Consider a two-player simultaneous matching pennies game for example. Its normal form representation is given by:

		Player 2	
		Head	Tail
Player 1	Head	(1,1)	(0,0)
	Tail	(0,0)	(1,1)

Now, suppose we represent this in extensive form by using a game tree. Then, we get the following tree:

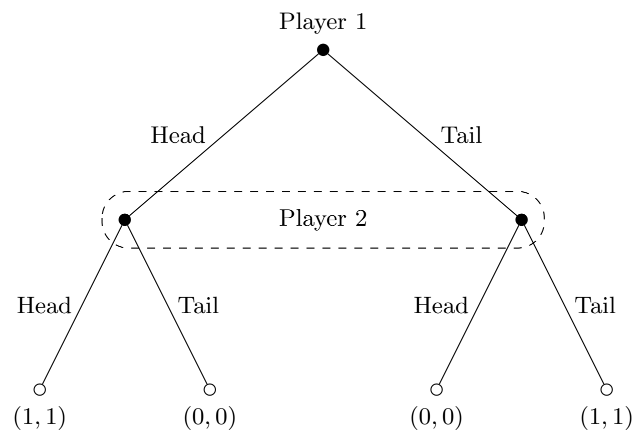

Notice that the actions of Player 2 are lumped into the same information set, as indicated by the broken oval, since Player 2 does not know what action Player 1 took when s/he needs to take an action. So, indeed strategic uncertainty gives rise to imperfect information.

- Chapter 9 goes into further detail on aspects related to imperfect and incomplete information. In particular, the chapter deals with some classic cases in which players have incomplete information that is also asymmetric. As was mentioned in chapter 2, in the presence of incomplete information, the player with superior information may want to *signal* his/her type. Meanwhile, the player with inferior information would like to *screen* players.

- Some typical situations of incomplete information:

  1. **Principal-agent relationships** Typically, the principal is the player with inferior information and the agent is the player with superior information about his/her type. Here are some examples:

     (a) **Employer-employee relationships** Employee knows his type (high or ability); employer does not. The employer would like to screen and the employee would like to signal. Typical employee signal: *education* (Spence's job market signaling model in the *Quarterly Journal of Economics* in 1973; see also pp. 326-332). Typical screening process: *interviewing/job market description/references.*

     (b) **Voter-politician relationships** Politician knows his type (corrupt or not); voter does not. The voter would like to screen and the politician would like to signal. Typical politician signal: *taxes/campaigns/public goods* (Besley and Smart's political agency/signaling model in the *Journal of Public Economics* in 2007). Typical screening process: *elections/debates/townhall.*

     (c) **Insured-insurer relationships** Insured knows his type (risky or not); insurer does not. The insurer would like to screen and the insured would like to signal. Typical insured signal: *good conduct.* Typical screening process: *credit/history/driving record.*

(d) **Stockholder (Board)-management (CEO) relationships**.

(e) **Landlord-Tenant relationships**.

- Related terms: moral hazard and adverse selection (see pp. 308 and 324).
- Typically, determination of the agent's type in signaling games is introduced as a *move of nature*. This is owing to Harsanyi (*International Journal of Game Theory*, 1973), who suggested introducing nature as a separate player (external uncertainty) in order to transform a game of incomplete information into one of imperfect information.
- *Disclaimer: Signaling games are one classic example of games of incomplete information. They are typically sequential-move games. We will not deal with them in detail at this stage. Time permitting, we will return to them in the behavioral/experimental sections of the course. Some of you may also have to deal with them in your research proposals. Beware that we have just barely scraped the surface!*

3 Risk and Uncertainty (GoS, chapter 7 and articles on prospect theory)

- As previously discussed, uncertainty and information are joined at the hip. As such, it makes sense to discuss some details on how we deal with risk and uncertainty. This also serves as a good bridge to cross over into specific behavioral topics.
- To start, let us make a distinction between risk and uncertainty (sometimes also referred to as ambiguity). Consider a decision D that can have three different outcomes: a, b, c. We say that D is *risky* if the probability with which each outcome occurs, that is p_a, p_b, p_c, is known. If any of these probabilities is not known, we say that D is *uncertain* (or ambiguous).
- So, for example, we say that a player's action/strategy in a game gives rise to strategic *uncertainty* for the other players because these players *do not know* the probability with which this person is going to choose any particular action/strategy.
- With this being said, as we saw previously, the other players may have a (subjective) belief (recall the Bayesian notion of probability) about a given player's action/strategy.
- The seminal question in this context is *'how do we define utility/payoff functions in the presence of risk/uncertainty?'* There are two principle models for dealing with this:[1]

    1. *Expected utility theory (EUT; aka the workhorse model)*, which was basically developed by von Neumann and Morgenstern in their 1944 classic title "The theory of games and economic behavior". As such, most of game theory is based on this approach. Simply put, in the above example the expected utility of D is, $EU(D) = p_a \cdot U(a) + p_b \cdot U(b) + p_c \cdot U(c)$. In this context, p refers to probabilities or subjective beliefs, for example in the case of strategic uncertainty in games.
    2. *(Cumulative) prospect theory (CPT; aka the main contender model)*, which was basically developed by Kahneman and Tversky (two psychologists) in their 1979 classic title "Prospect Theory: An Analysis of Decision under Risk". This model is the

[1] There are also other models such as rank dependent utility theory and dual theory, but these will not be discussed here.

main contender to EUT and is based on psychological foundations and as such, it is more 'behavioral'. Simply put, in the above example the prospect D is evaluated as $\pi_a \cdot v(a) + \pi_b \cdot v(b) + \pi_c \cdot v(c)$. In this context, π refers to decision weights, which are transformations of the ps.

- What are the key differences?

  1. Notice that so far we have been silent about $U(.)$, $v(.)$, and their arguments. This is one main difference between EUT and CPT. In EUT, U is assumed to be an increasing ($U' > 0$) and concave ($U'' < 0$) utility function that is defined over terminal wealth or money–an *absolute* value of wealth. In CPT, v is assumed to be a value function that is increasing ($v' > 0$) and exhibits diminishing sensitivity (more below) over *gains* and *losses* measured *relative* to some *reference point (or status quo)*, which is usually calibrated/normalized to zero–in particular, $v(0) = 0$.
  2. EUT allows for the decisionmaker/player to be risk averse, risk neutral, or risk seeking, depending on one's assumption on $U(.)$. To define risk aversion, suppose there is a decision E which gives the following outcome *with certainty*: $p_a \cdot a + p_b \cdot b + p_c \cdot c$. Notice that this outcome is equivalent to the expected outcome for decision D. A player is said to be *risk averse* is s/he prefers E to D. S/he is *risk-neutral* if indifferent and *risk-seeking* if she prefers D to E. Decisionmakers/players tend to be assumed to be risk averse, which is usually captured by assuming that $U'' < 0$.
     CPT tends to assume that the decisionmaker/player is *loss averse*. This means that people are much more sensitive to losses than to gains of the same size. This is done by making v steeper in the region of losses than the region of gains (see Barberis, p. 176).
  3. EUT assumes diminishing marginal utility over terminal wealth (U is concave); whereas, CPT assumes diminishing sensitivity more generally (v is concave over gains–like EUT–but also convex over losses–that is, people are *risk-seeking* over losses).
  4. As previously seen, EUT assumes that the decisionmaker/player weighs outcomes by the objective probabilities p, which can sometimes be seen as subjective beliefs; whereas, CPT assumes that the decisionmaker/player assigns transformed probabilities or decision weights π to outcomes. The weights are calculated with the help of a weighting function $w(.)$ whose argument is an objective probability p. In particular, in CPT the weighting function is applied to cumulative probabilities. The typical assumption is that the decisionmaker/player tends to overweight small probabilities and underweight high probabilities.

- To get a feel for EUT, refer to the appendix to GoS chapter 7 and Barberis. To get a feel for CPT, refer to Barberis and Kahneman and Tversky.

- Empirical evidence is presented by the paper by Harrison and Rutström that will be presented by one of your fellow students. The paper by Charness and Viceisza discusses different ways to elicit people's risk preferences in the field. This paper will also be discussed by one of your fellow students.

ECON/MGT 395 Behavioral Game Theory
Lecture title: "Information Updating/Processing and Anomalies"

Dr. Angelino C. G. Viceisza
Spelman College
Economics

November 1, 2013

1 Acknowledgements

This note is created by the author and is in part based on the following sources:

- The math appendix to Joel Watson's Strategy (this is posted on Moodle since the beginning of the course).
- See also the articles for the student presentations by Rabin and Vayanos (2010) and Plott and Zeiler (2005).

2 Information Updating and Processing

- Previously, when discussing information we saw the link between uncertainty, risk and information. In particular, in games of imperfect or asymmetric information the player is uncertain about another player's action or faces risk with regard to another player's type (for example, recall the case of an employer seeking to fill a position–the employee's type may be risky). So, (imperfect or incomplete/asymmetric) information tends to be joined at the hip with probabilities and random variables!
- Two key questions come to mind:
  1. *Do people understand information on probabilities (that is, expressed as probabilities) to start with?*
  2. *To the extent that they do, how do people process* new *information on probabilities?*
- The first question is really asking whether people understand random variables. The article by Rabin and Vayanos (2010) will deal with this in greater detail. See student presentation.
- The second question is really a question of "updating". Consider an example owing to Watson's appendix (pp. 375-376). Suppose there are three horses in a race: A, B, C. The first two horses are from Kentucky, whereas C is from Tennessee. Suppose further that the probability of winning is given by $p(A) = 0.5$, $p(B) = 0.25$, and $p(C) = 0.25$.
- Based on this information, the chance that horse A will win is 0.5.

- However, suppose I tell you (that is, you receive new information) that the race has been run and that I heard from a reliable source that a horse from Kentucky won. On the basis of this new information, what would you assess the probabilities of A and B to be?

- Your *updated* assessment should be A with probability 2/3 and B with probability 1/3. Why?

- Initially, you believed that A is twice as likely to win as B. Knowing that C did not win (based on the newly received reliable information), you should assign zero probability to C and partition the remaining probability of one (since $p(A) + p(B) + p(C) = 1$ (by laws of probabilitues) according to the respective weights. Knowing that $p(A) = 2p(B)$ you get: $p(A) + p(B) + p(C) = p(A) + p(B) + 0 = 3p(B) = 1$ such that $p(B) = 1/3$. Accordingly, $p(A) = 2/3$. These new probabilities are also known as "conditional probabilities". In particular, $p(A)$ now represents $p(A|Kentucky)$–the probability of A winning knowing that (conditional on the fact that) C did *not* win (the horse is from Kentucky).

- In general, conditional probability can be figured by reducing to 0 the probabilities of states that have been ruled out and scaling up the probabilities of the states that are left.

- In more technical language, take a state space X, a probability distribution p and two events $K \subset X$ and $L \subset X$. Suppose you learn that the state is definitely in the set L and you want to know the updated probability that it is in K. (For the example above: L is the horse being from Kentucky; K is the horse being A) Because the state cannot be outside L, the only way it can be in K is if it is an element of BOTH sets (that is, it is an element of the intersection): $K\hat{L}$. To maintain the proportions between the states in L, you need to scale up the probabilities by the factor $p(L)$. Thus, the probability of event K conditional on event L is defined as:
$$p(K|L) \equiv \frac{p(K \cap L)}{p(L)}. \tag{1}$$

- In our example, we have:
$$p(K|L) \equiv \frac{p(K \cap L)}{p(L)} \rightarrow p(A|Kentucky) = \frac{p(A)}{p(Kentucky)} = \frac{1/2}{3/4} = 2/3. \tag{2}$$

- Note that this expression is undefined if $p(L) = 0$. That is, if you initially believe that the horse from Kentucky could never win, you cannot update your probability–it is of little guidance.

- As Watson derives, this expression can be rewritten to give the so-called *Bayes' rule*, which also says something about how people should update their probabilities as a result of new information:
$$p(K|L) = \frac{p(L|K)p(K)}{p(L)}. \tag{3}$$

- There is evidence to suggest that we–as people–do not necessarily update probabilities in this manner. Part of that may indeed be because we do not understand probabilities to begin with–see student presentation on article by Rabin and Vayanos (2010).

3 Anomalies (revisited)

- Behavioral (empirical) research over the past few years has documented several inadequacies of several models. In particular, as seen when discussing time preferences, the DU model has been critiqued for not being a proper descriptive model of behavior.
- These empirical regularities, which are at odds with the DU (and other models), are typically referred to as *anomalies*. The reason for returning to anomalies here is because the fact that we do not "Bayesian update" (so to speak) can be considered an anomaly of Bayes' rule as a descriptive model.
- As stated in the note on time preferences, there are several anomalies and as a result, "better" models may need to be considered.
- The "endowment effect" discussed in one of the student presentations is one type of anomaly that needs to be accounted for.
- See the time preference note and Frederick et al. for other types of anomalies and models that may be able to account for them.

4 Neuroeconomics

- Some have argued that neurological data being combined with economics data (e.g. from experiments) can improve our understanding of behavior and thus enable us to describe and predict people better. This new, evolving, and controversial field is also known as "*neuroeconomics*".
- Those who are interested feel free to discuss with me further under separate cover. One useful starting point is a collection of papers by Caplin and Schotter (2008) entitled "The Foundations of Positive and Normative Economics".

ECON/MGT 395 Behavioral Game Theory
Lecture title: "Miscellaneous Topics"

Dr. Angelino C. G. Viceisza
Spelman College
Economics

September 18, 2013

1 Acknowledgements

This note is created by the author and is in part based on the following sources:

- The course textbook–Games of strategy by Dixit et al; chapters 3 and 6.
- Parts of Behavioral game theory by Camerer such as chapter 2 on dictator, trust, and ultimatum games and chapter 4 on bargaining games (this is related to the student presentations).
- The article by Gibbons–in particular, the section on dynamic games of complete information.
- The article by Weibull–Testing Game Theory.
- The following article: Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri. 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications" *Journal of Economic Literature*, 51(1): 5-62 (see Moodle readings).

2 Three typical sequential-move games (BGT, chapters 2 and 4)

- Three standard sequential-move games that tend to be used quite a bit in experimental/behavioral economics/game theory are:
  1. The trust game (sometimes also referred to as the investment game–its more general counterpart). This game was first introduced by Berg at al. in 1995 in the article "Trust, Reciprocity, and Social History" in the journal called *Games and Economic Behavior*. See also discussions by BGT (in particular, the student presentations–chapter 2) and Gibbons. We will see this game again in the experimental section of the course!
  2. The dictator game. See also discussions by BGT (again, the student presentations–chapter 2). We will also see this game again in the experimental section of the course!
  3. The ultimatum game (sometimes also referred to as the bargaining game). See also discussions by BGT (again, the student presentations–chapter 4).
- What do each of these games look like and how are they related?

1. In the simplest version of the trust game, there are two players and each of them has x to start the game. Player 1–the first mover–has the option to send all of x to Player 2–the second mover–or not. If she sends all of x, she is left with nothing and Player 2 is left holding all the stakes. In particular, if Player 1 sends x, the amount Player 2 gets is actually its triple–that is, Player 2 gets $3x$ in addition to the x she already had. So, if Player 1 sends x, she is left with zero and Player 2 has $4x$. Now, Player 2 has a move to make. Player 2 can either choose to return $2x$ to Player 1 and keep $2x$ for herself (that is, she can split the pie), or she can keep all $4x$. Recall that Player 1 could also choose not to send x to start with. If so, then the game would end right there. She would keep x and so would Player 2. This game is typically called the trust game and its extensive form representation is given by the following tree:

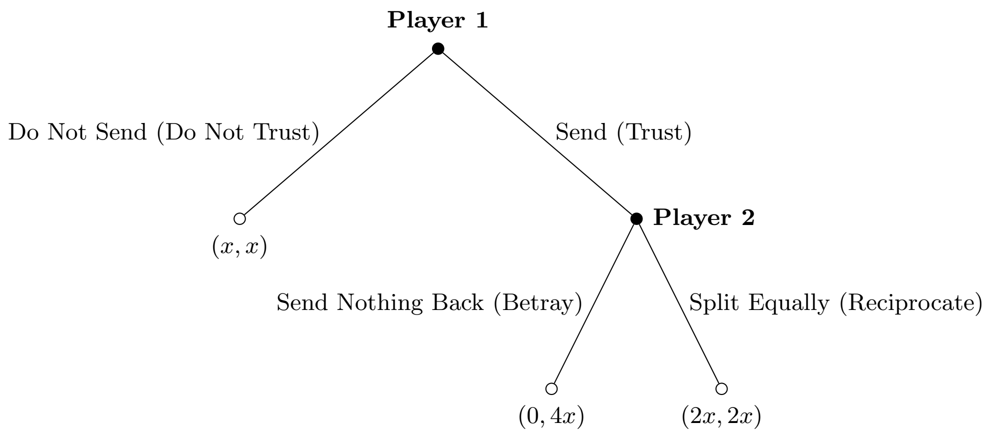

It is called the trust game because Player 1 takes a risk by sending all to Player 2. Once she does, there is no reason why Player 2 should send anything back. She is "trusting" Player 2! In turn, Player 2 has no reason to send anything back if Player 1 sends x. So, her act to send back and split the pie equally is typically called "reciprocity". She is returning Player 1's kind and trusting act by reciprocation! *Does this make sense?*

2. What we said previously is not quite correct! We said that Player 1 has no reason to send x because Player 2 has no reason to send back. Is this true? Well, that depends. It is true that IF you assume that players are economically rational in the *traditional* sense (meaning that they are *selfish* or *self-regarding*, with no concern for others), then Player 1 should expect nothing back. However, what if Player 1 believes that Player 2 is someone who cares about others, either because she is altruistic or because social norms dictate that one must share! Then, the situation is different. If we assume that Player 2 cares about Player 1's utility or happiness–that is, Player 2 has so-called social preferences/altruism–then Player 1 may send in expectation that Player 2 will return. Also, Player 1 may send because she is also altruistic. To disentangle Player 1's social preference (altruism) from her 'trust' motive and similarly Player 2's social preference (altruism) from her 'reciprocal' motive, we can use the so-called dictator game.[1] The dictator game is very simple. A player has x and she can choose to split it or not with another player. Whatever she decides is the final allocation. The dictator game can thus be seen as a measure of altruism since the one player has no reason–other than some form of caring–to share with the other.[2]

[1] This is further discussed in an article by James C. Cox called "How to Identify Trust and Reciprocity?" in 2004 in *Games and Economic Behavior*.

[2] Some research has raised concerns with the dictator game suggesting that its findings are a mere artefact of

3. Finally, the ultimatum game in its simplest form adds one additional step to the dictator game. The other player now has a move also. Once the first mover decides whether or not to share x, the second mover can decide whether or not to accept that allocation. If she does, it becomes binding; however, if she does not, they both get zero. Of course, the assumption is that $x > 0$ and therefore, the split is also greater than zero.
4. These three games can be combined to disentangle people's motives in the trust game! The dictator and ultimatum games can also be used independently. The former to measure social preferences; the latter to measure bargaining situations. In particular, the so-called centipede game is a series of combined ultimatum games, which can be seen as alternating offers. See discussions by Gibbons and Camerer.

3 Testing game theory (based on Weibull, Camerer, GoS)

- There are additional noteworthy aspects:
  1. In all these sequential games, we can use rollback/backward induction and/or SPNE.
  2. Thus far, we have been silent on a key aspect! When deriving solutions to games (be it, IEDS, NE, rollback/backward induction, or SPNE), we always assumed that players were rational. In particular, we assumed that players care about maximizing their own payoffs in the game. The payoffs constructed theoretically have all assumed that we have correctly captured ALL factors that are important to players. This is easy to do in theory, because it is a mere assumption. As soon as we step away from theory and start asking how real people *actually* behave, we need to be very careful with these assumptions. In particular, if real players were to play an actual trust game, an actual PD game, or any other real game, the payoffs they typically get will be *monetary* payoffs. These are different from game-theoretic payoffs because monetary payoffs are individualistic. So, if a player is altruistic, she will care about her payoff, but also the other player's payoff! Recall when I said that we may define utility as follows: $U(x_i, x_{-i})$ where x_i is the monetary payoff that comes to player i and x_{-i} is the monetary payoff that comes to the other player $-i$.
  3. This is related to the student presentation on BGT, chapter 2. Recall that the theory section was discussing Fehr-Schmidt, Bolton-Ockenfels, and so on. These theories were trying to model how people form social preferences such as altruism. Why is this relevant? Because as we saw, the rollback equilibrium and in fact any other solution to a game, "in theory" depends on the assumptions that are placed on people's utility functions/preferences. So, in order to derive testable hypotheses for when real games as conducted in the discipline of "behavioral" game theory, we must adjust our theories of preferences.
  4. You will see several of these aspects discussed while reading the paper by Weibull. To help you link these aspects, he calls "real" games, game protocols. So, a theoretical trust game is a "game". A trust game conducted with real players (for example in a classroom experiment) is a "game protocol".
  5. Weibull also discusses the very related concept of *context dependence*. When solving a game through backward induction/rollback, as we saw, we (1) start at the end of

the experiment and not necessarily a measure of altruism. Be cautious when constructing theory and designing games/experiments! See for example John List's discussion in "On the interpretation of giving in dictator games" that appeared in 2007 in the *Journal of Political Economy*.

the game; (2) solve successive subgames; and (3) simplify the game accordingly till we get to one or more outcomes. As he points out, while this is a valid approach, it is important to factor in whether a subgame was preceded by another player's move or not. Namely, if so, this subgame is context-dependent and real players are likely to perceive the move differently. As such, their preferences need to account for that. Otherwise, the theoretical solution to the game is likely to be misguided. As such, he distinguishes between the context-dependent definition of a subgame and the isolated definition of a subgame. As an example, consider the second player's move in a trust game. If you view this as an isolated subgame, it is as if the player were playing a dictator game. However, viewing it as a subgame of the larger game (that is viewing it as context-dependent from a preference standpoint), we know that it is NOT like a dictator game because it was preceded by another player's move. So, the second mover treats this differently than a dictator game. This is in fact the rationale for why Cox (2004) is able to distinguish reciprocity from altruism!

4 Empirical Evidence: Equilibrium or Not? (GoS, chapter 6; BGT; Crawford et al.)

- Thus far, we have talked about several solution concepts, in particular, equilibrium in the form of NE and SPNE. These are based on best-responses and backward induction/rollback, combined with assumptions on preferences.
- As Crawford et al. argue (see also BGT), despite these powerful concepts, equilibrium does not always exist or is not confirmed by the empirical evidence, at least when real people play games for the first time (equilibrium may however emerge through *learning*).
- In particular, several recent experimental and empirical studies suggest that people's initial responses to games often deviate systematically from equilibrium, and that structural nonequilibrium "level-k" or "cognitive hierarchy" models often out-predict equilibrium.
- To illustrate this concept, let us do a classroom experiment: A guessing game! (really, a version of a simultaneous-move game–relate back to presentation on coordination games based on BGT, chapter 7)