Games

Bruno Salcedo*

Winter 2020

So far, we have studied the behavior of a single economic agent. For the rest of the course, we will study the behavior of groups of agents using tools from *Non-Cooperative Game Theory*. The origins of Game Theory can be traced back to Cournot (1838). It became an independent discipline after the work of von Neumann (1928) and von Neumann and Morgenstern (1944), and has become a part of the standard toolbox of Economics since then. If you are looking for further reading materials, the standard textbook references are Myerson (1991), Fudenberg and Tirole (1991), and Osborne and Rubinstein (1994).

Non-Cooperative Game Theory is loosely based on three assumptions. First, group behavior is determined by decentralized choices made by individual agents. Second, individuals are rational, i.e., their behavior is consistent with the expected utility hypothesis. Third, individuals are sophisticated and have no cognitive constraints in their ability to process information. In particular, each individual can reason about the behavior of others in order to guide their own choices.

These notes focus on a class of models called normal-form games. We will use different solution concepts to make predictions: rationality, rationalizability, and equilibrium. Each of these solution concepts requires stronger assumptions about the beliefs of the individuals, but makes more precise predictions.

Group behavior is not always determined by decentralized choices. For example, when a typical family goes to a movie theater they decide as a group which movie to watch. It is not every individual member of the family choosing a movie independently and hoping to coincide. The study of decisions made by groups falls within the domain of *Cooperative Game Theory*. Unfortunately, we will only have time to discuss non-cooperative models in class. Osborne and Rubinstein (1994, Part IV) offers a good introduction to cooperative models.

*Department of Economics, Western University · brunosalcedo.com · bsalcedo@uwo.ca

1. Normal-Form Games

These notes focus on a special class of games called *normal-form games* or *strategic-form games*. In these games each agent makes a single choice. The choices are implicitly assumed to be independent of one another, in the sense that each agent has no information about the choices of other agents at the moment of making their own.[1] Although normal-form games appear to be static, later in the course we will see that they can also model dynamic environments.

In a normal-form game, there is a nonempty set of *players*. Each player must choose an *action* from a fixed nonempty set. A vector specifying one action for each player is called an *action profile*. The preferences of the players are summarized by *utility* functions defined over action profiles. Formally, a normal-form game is a mathematical object defined as follows.

Definition 1 A *game* in normal-form is a tupple $G = (I, A, u)$ where I is a nonempty set of players, $A = \times_{i \in I} A_i$ is a nonempty set of action profiles, and $u : A \times I \to \mathbb{R}$ specifies utility functions for each player.

Unless otherwise stated, I will always assume I and A are finite. All the definitions and results can be extended to the case with A compact and u continuous. Games with an uncountable number of players can be tricky. Their analysis is well beyond the scope of the course.

Typical players are denoted by $i, j, \ldots$, typical actions by $a_i, b_i, \ldots$, and typical action profiles by $a, b, \ldots$. It is useful to express action profiles as $a = (a_i, a_{-i})$, where a_i denotes the action taken by player i, and $a_{-i} \in A_{-i} := \times_{j \neq i} A_j$ denotes the profile of actions taken by i's opponents.

When a game has only two players, it can be represented with a matrix. Each row corresponds to the action of one of the players. Each column corresponds to the action of the other player. Each entry of the matrix lists two numbers corresponding to the utility of each of the two players. The convention is to list the payoff of the row player first. Consider for example the classic Prisoner's Dilemma game.

[1] The most common way to phrase this assumption is to say that choices are made 'simultaneously and independently', and many authors use the term 'simultaneous-move games'. However, I find such language to be somewhat misleading. The assumption is neither about the timing of the choices, nor about their (statistical) independence.

	C	D
C	$1, 1$	$1+g, -l$
D	$1+g, -l$	$0, 0$

Figure 1 – Payoff matrix for the Prisoners' Dilemma, where $l, g > 0$.

Example 1.1 Prisoner's Dilemma. Two prisoners are suspected of a crime. The district attorney (DA) has enough evidence to convict them of a misdemeanor, but would need a confession to convict them of the crime they allegedly committed. The DA offers each prisoner a sentence reduction in exchange of a confession. Each prisoner must choose to cooperate with her partner (C), or to defect (D) by accepting the DA's deal.

If none of the prisoners confess, then both of them would serve short sentences. If only one prisoner confesses, she would be free while her accomplice would serve a long sentence. However, if both prisoner's confess, then both of them would serve an intermediate sentence. Assuming that the prisoner's preferences depend only on the amount of time they serve, these preferences can be represented by the utility functions in Figure 1.

1.1. Randomization

Probability measures over actions play an important role for at least three reasons. First, recall from our discussion of expected utility theory that dominance by randomized actions is useful to characterize rational behavior. Second, treating a_{-i} as a random variable captures the idea that i is uncertain about the behavior of their opponents—whether they are actually randomizing or not. Third, there are situations where players actually randomize. For example, Walker and Wooders (2001) and Chiappori et al. (2002) find evidence of randomization by professional athletes. Consider also the following anecdote about Kenneth Arrow.

During World War II, Arrow was assigned to a team of statisticians to produce weather forecasts. The forecasts were used to make strategic decisions, such as the timing of bombarding attacks. At some point, Arrow determined that their forecasts were not very good. He wrote a report showing that using today's weather as a forecast of tomorrow's weather would be more accurate that the forecasts his team was producing. His commanding general responded saying that he was well aware that the forecasts were no good. However, he needed them for planning purposes. Basing the deployment of troops on the current weather would make the attacks predictable, which would give an advantage to the opposing forces. Part of the goal of producing complex inaccurate forecasts was to serve as a randomization device that induced unpredictable behavior.[2]

Let us introduce some notation and terminology for random objects. A*mixed action* for player i is a distribution $\alpha_i \in \Delta A_i$. I will sometimes refer to elements of A as *pure actions* to emphasize the distinction with mixed actions. With slight abuse of notation, we identify each pure action with the degenerated mixed action that assigns full probability to it. The *beliefs* of player i are captured by distributions $\alpha_{-i} \in \Delta A_{-i}$. A *correlated action profile* is a joint distribution $\alpha \in \Delta A$. Given $\alpha_i \in \Delta A_i$ and $\alpha_{-i} \in \Delta A_{-i}$, let (α_i, α_{-i}) denote the correlated action profile α given by $\alpha(a_i, a_{-i}) = \alpha_i(a) \cdot \alpha_{-i}(a_{-i})$.

The utility that agent i receives from action profile a is denoted by $u_i(a)$. Expected utility is denoted by

$$U_i(\alpha) = \mathbb{E}_\alpha\left[\, u(\mathbf{a}) \,\right] = \sum_{a \in A} \alpha(a) u_i(a).$$

1.2. Rationality

From the perspective of a player, a normal-form game is nothing more than a decision problem. Let us assume that all players are rational, in the sense that they make choices to maximize their expected utility. As we learned during our discussion of decision theory, rationality alone can rule out some forms of behavior that are never optimal. For example, defecting always yields a higher expected utility than cooperating for each prisoner in the Prisoner's Dilemma. This is regardless of her beliefs about her accomplice's behavior. Hence, assuming that players are rational allows us to conclude that they will both defect. This section restates some of the definitions and results we learned before using the language of normal-form games.

[2] I heard this version of the anecdote from my adviser. People often quote a different—less exciting—version that was recorded in the minutes of a meeting of the Federal Reserve Board (Board of Governors of the Federal Reserve System, 2005, pp. 11).

An action $a_i \in A_i$ is said to be a *best response* to belief $\alpha_{-i} \in \Delta A_{-i}$ if and only if $U(a_i, \alpha_{-i}) \geq U(a'_i, \alpha_{-i})$ for every $a'_i \in A_i$. The set of best responses to α_{-i} is denoted by $\mathrm{BR}_i(\alpha_{-i})$. Given a game G, an action $a_i \in A_i$ is *rational* for player i if there exists some $\alpha_{-i} \in \Delta A_{-i}$ such that $a_i \in \mathrm{BR}_i(\alpha_{-i})$. Rational players would never choose actions that are not rational. Rational actions can be characterized using strict dominance. A pure action $a_i \in A_i$ is *strictly dominated* by a pure or mixed action $\alpha_i \in \Delta A_i$ if and only if $u_i(a_i, a_{-i}) < U_i(\alpha_i, a_{-i})$ for every $a_{-i} \in A_{-i}$. Actions that are not strictly dominated by any pure or mixed action are called *undominated.*

Theorem 1.1 *A pure action is rational if and only if it is undominated.*

The proof of Theorem 1.1 and the proofs of the following two lemmas are in the lecture notes on Dominance and Rationality.

Lemma 1.2 *If a mixed action assigns positive probability to a dominated pure action, then it is dominated.*

Lemma 1.3 *Every dominated action is strictly dominated by a mixed action supported on the set of undominated actions.*

Every solution concept we will analyze for the rest of these notes will assume that players are rational. However, even this assumption is not without controversy. In the Prisoner's Dilemma, rationality implies that both prisoners must defect, which is problematic for two reasons. First, mutual cooperation is a Pareto improvement relative to mutual defection. It is peculiar that our definition of rationality implies behavior that is unambiguously irrational from the perspective of the group. Second, the existing experimental and field evidence shows that people often cooperate even when doing so is dominated (e.g., Dawes and Thaler (1988)). Hence, the rationality assumption can be questioned both from a normative and from a positive perspective.

There is a vast literature attempting to explain cooperation in the Prisoner's Dilemma and other forms of pro-social behavior. Cooperation could be driven by different factors including psychological aspects (Geanakoplos et al., 1989, Rabin, 1993), altruism or preferences for fairness (Andreoni, 1990, Fehr and Schmidt, 1999), dynamic incentives (Kandori, 1992, Ellison, 1994), information incentives

(Nishihara, 1997), or Kantian reasoning (Roemer, 2010, 2019). The most common approach in economics (including this class) is to assume that agents are rational, have selfish preferences, and make independent choices. However, it is worth remembering that this assumption is not without loss of generality. It might lead to poor approximations of reality, at least in some specific contexts.

2. Rationalizability

Assuming that all players are rational is enough to predict a specific outcome in the prisoner's dilemma. In most games, rationality alone has limited predictive power. We can make more precise predictions by making further assumptions that restrict the beliefs of the players. In this section, we assume that all players know that all players are rational, all players know that all players know that all players are rational, all players know that all players know that all players know that all players are rational, and so on. This assumption is called *common knowledge of rationality*. The corresponding solution concept is called *rationalizability* and was introduced by Pearce (1984) and Bernheim (1984).[3]

2.1. Common Knowledge

The notion of common knowledge was introduced by Lewis (1969) and formalized by Aumann (1976). Unfortunately, we don't have enough time in class to go over the rigorous details. We will have to make do with an informal definition and an example.

Definition 2 A fact is *mutually known* among a group of individuals if everybody knows it. It is *commonly known* if everybody knows it, and, in addition, everybody knows that everybody knows it, everybody knows that everybody knows that

[3]The solution concept that I call rationalizability is what other authors call *correlated rationalizability*. It differs from the original definition of rationalizability, which require that the players beliefs about their opponents are product beliefs. For example, if $I = \{1, 2, 3\}$, then player 1 must believe that the actions of 2 and 3 are statistically independent. I believe such restriction was imposed for historical reasons, and I see no good reasons to maintain it.

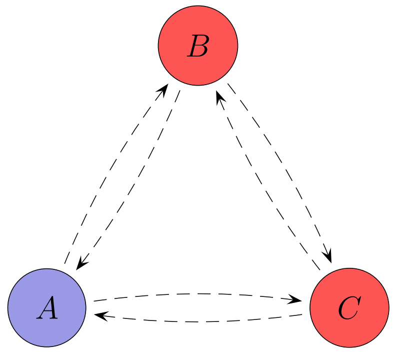

Figure 2 – Three logicians wearing hats

everybody knows it, and so on and so forth.

The following example illustrates that there is a big difference between mutual knowledge and common knowledge. Anna Bob and Charlie are sitting in opposite corners of a room with no mirrors. Each one of them is wearing either a blue hat or a red hat. Each one of them can see color of the hat of the other two people in the room, but not the color of his/her own hat. For instance, Anna can see that Bob and Charlie have red hats, but she cannot tell whether her own hat is red or blue. We assume that all this information is common knowledge. Anna is wearing a blue hat and Bob and Charlie are wearing red hats. See Figure 2.

2.1.a. *Mutual knowledge.*– Suppose that Daniel enters the room and announces that everybody is wearing either blue or red hats. Then he proceeds to ask them one by one "*Which color is your hat?*" First he asks Anna, then Bob and then Charlie. In every case the answer is the same "*I don't know*".

Note that it is common knowledge that everybody is wearing either a blue hat or a red hat. This is because this fact was publicly announced and everybody noticed that everybody heard it. This however does not imply that there has to be either a red hat or a blue hat. It could very well be the case (given the information that Anna, Bob and Charlie have) that all hats are blue or all hats are red.

Anna knows that there are at least two red hats, because she can see Bob and Charlie's hats. But this does not imply any information about her own hat. Similarly, Bob and Charlie know that there is at least one red hat, and at least one blue hat. But they cannot infer anything about the color of their own hats.

Hence nobody is able to provide a definitive answer to Daniel's question. *Notice that it is mutual knowledge (everybody knows) that there is at least one red hat.*

2.1.b. *Common knowledge.*– Now suppose that Daniel announces that everybody is wearing either a blue or a red hat. In addition, he also announces that there is at least one red hat in the room. The he proceeds as before asking them one by one "*Which color is your hat?*". Anna and Bob answers as before: "*I don't know*". However, Charlie answers triumphant: "*My hat is red!*"

The only difference between the two scenarios is that, in the second one, Daniel made the additional announcement that there is at least one red hat. However, this is something that *everybody already knew*. The difference is that, by making the announcement public, the existence of at least one red hat went from being mutual knowledge to being common knowledge. After the announcement, everybody knew that everybody knew that there was at least one red hat. This is what allowed Charlie to deduce that her hat was red. Let's see how.

Charlie knew that Bob knew that there was at least one red hat. Hence, if he had seen only blue hats, he would have known that his own hat had to be the red one. Since he did not know the color of his own hat, it had to be the case that he was already seeing at least one red hat. That is, either Anna or Charlie hat to be wearing a red hat. Since Charlie could see that Anna's hat was blue, this meant that her own hat had to be red. This line of thought was only possible because she knew that Bob knew that there was at least one red hat.

This cartoon about three logicians in a bar tells a simpler version of the story.

2.2. Definition of Rationalizability

To understand how common knowledge of rationality can help make finer predictions, let us start with an example.

Example 2.1 Consider the game from Figure 3. The only dominated action is d, which is dominated by b. Action w is rational, because it is a best response to b. However, if Anna knows that Bob is rational, then she would assign zero probability to Bob playing d. For any such belief, y gives a higher expected utility that w. Hence, if Anna is rational and knows that Bob is rational, she would never play w.

		Bob			
		a	b	c	d
	w	$0,0$	$0,1$	$0,0$	$1,0$
	x	$3,0$	$1,1$	$0,3$	$0,0$
Anna	y	$1,0$	$10,10$	$1,0$	$0,0$
	z	$0,3$	$1,1$	$3,0$	$0,0$

Figure 3 – A 4×4 game. The rationalizable space is $A^* = \{x, y, z\} \times \{a, b, c\}$.

Given a game G, action sub-space is a set $B = \times_{i \in I} B_i$ such that $B_i \subseteq A_i$ for all $i \in I$. Action a_i is said to be rationalizable with respect to an action subspace B if and only if it is a best response to some belief that assigns full probability to $B_{-i} := \times_{j \neq i} B_j$. Let $R_i(B) \subseteq A_i$ denote the space of actions in B_i that are rationalizable with respect to B, that is,

$$R_i(B) = \left\{ b_i \in B_i \middle| \quad \exists \beta_{-i} \in \Delta B_{-i}, \forall a_i \in A_i,\ U_i(b_i, \beta_{-i}) \geq U_i(a_i, \beta_{-i}) \right\}.$$

Also, let $R(B) = \times_{i \in I} R_i(B)$.

Definition 3 Given a game G, an action subspace B is said to be *self-rationalizable* if every action $b_i \in B_i$ is rationalizable with respect to B, i.e., if $R(B) = B$. An action a_i is *rationalizable* if and only it belongs to a self-rationalizable subspace.

Example 2.1 (continued). In our example, the subspace $\{(y, b)\}$ is self-rationalizable because y is a best response to b and vice versa. The sub-space $\{x, z\} \times \{a, c\}$ is a also self-rationalizable. This is because x is a best response to a, z is a best response to b, a is a best response to z, and c is a best response to x. The union ob both subspaces, i.e., $\{x, y, z\} \times \{a, b, c\}$ is also self-rationalizable.

Moreover, these are the only self-rationalizable subspaces. To see why, first note that a self-rationalizable space cannot contain d, because it is not rational. Therefore, it also cannot contain w because w cannot be a best response to beliefs that assign zero probability to d. By a similar reason, if a self-rationalizable subspace contains x then it must also contain a. If it contains a, then it must also contain z. If it contains z, then it must also contain b. And if it contains b, then it must also contain x.

We can thus conclude that the rationalizable actions are a, b, and c for Anna, and x, y and z for Bob.

Let us try to understand the connection between our definition of rationalizability and common knowledge of rationality. Suppose that an action a_i is consistent with common knowledge of rationality. If i is rational, a_i must be a response to some belief. If, in addition, i knows that all players are rational, then a_i must be a best response to a belief over actions that are in turn best responses. It other words, it must be a best response to best responses. If , in addition, i knows that all players know that all players are rational, then a_i must be a best response to best responses to best responses. If all players are rational, know that all players are rational, and know that all players know that all players are rational, then they would choose actions that are best responses to best responses to best responses.

This line of reasoning can be repeated an infinite number of times. Now, consider the subspace consisting of all actions that have positive probability at least once in the sequence of beliefs. This subspace must be self-rationalizable. Hence, if a choosing action is consistent with common knowledge of rationality, then it must be rationalizable according to Definition 3. The connection between self-rationalizable sets and commonknwledge of rationality can be formalized using epistemic models, but doing so is beyond the scope of this course. Instead, I will only state the result without a proof. This result was first proven by Brandenburger and Dekel (1987) and Tan and Werlang (1988). For further reading on this topic, I recommend Brandenburger (1992).

Proposition 2.1 *Given a game G, an action a_i is rationalizable if and only if choosing it is consistent with common knowledge of rationality.*

An action profile is rationalizable if it consists of rationalziable actions. Denote the set of all rationalizable actions for player i by A_i^*, and the set of rationalizable action profiles by $A^* = \times_{i \in I} A_i^*$. We can show that A^* is the largest self-rationalizable set, in that it is self-rationalizable and contains every other self-rationalizable set.

Proposition 2.2 *A^* is the largest self-rationalizable set.*

Proof. The definition of A^* directly implies that if an action subspace B is self-rationalizable, then $B \subseteq A^*$. Hence, we only have to show that A^* is self-rationalizable. Take any player i and any action $a_i \in A_i^*$. There exists some self rationalizable subspace $B = \times_{i \in I} B_i$ such that $a_i \in B_i$. Since $B = R(B)$, there exists some $\alpha_{-i} \in \Delta B_{-i}$ such that $a_i \in \mathrm{BR}_i(\alpha_{-i})$. Since $B \subseteq A^*$, it follows that $\alpha_{-i} \in \Delta A^*_{-i}$. Hence, a_i is rationalizable with respect to A^*. Since i and a_i were arbitrary, it follows that $A^* = R(A^*)$. ∎

References

Andreoni, J. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *The Economic Journal*, 100(401):464–477.

Aumann, R. J. (1976). Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239.

Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028.

Board of Governors of the Federal Reserve System (2005). Meeting of the Federal Open Market Committee on September 20, 2005. Retrieved on February 5th, 2020 from https://www.federalreserve.gov/monetarypolicy/files/FOMC20050920meeting.pdf.

Brandenburger, A. (1992). Knowledge and equilibrium in games. *Journal of Economic Perspectives*, 6(4):83–101.

Brandenburger, A. and Dekel, E. (1987). Rationalizability and correlated equilibria. *Econometrica*, 55(6):1391–1402.

Chiappori, P.-A., Levitt, S., and Groseclose, T. (2002). Testing mixed-strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer. *American Economic Review*, 92(4):1138–1151.

Cournot, A. (1897/1838). *Researches into the Principles of the Theory of Wealth [Recherches sur les principes mathématiques de la théorie des richesses]*. The Macmillan Company. bated from French by Bacon, N. T.

Dawes, R. M. and Thaler, R. H. (1988). Anomalies: cooperation. *Journal of Economic Perspectives*, 2(3):187–197.

Ellison, G. (1994). Cooperation in the prisoner's dilemma with anonymous random matching. *The Review of Economic Studies*, 61(3):567–588.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation.

The Quarterly Journal of Economics, 114(3):817–868.

Fudenberg, D. and Tirole, J. (1991). *Game Theory*. The MIT Press.

Geanakoplos, J., Pearce, D., and Stacchetti, E. (1989). Psychological games and sequential rationality. *Games and Economic Behavior*, 1(1):60–79.

Kandori, M. (1992). Social norms and community enforcement. *The Review of Economic Studies*, 59(1):63–80.

Lewis, D. (1969). *Convention: A Philosophical Study*. Harvard University Press.

Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press.

Nishihara, K. (1997). A resolution of N-person prisoners' dilemma. *Economic Theory*, 10(3):531–540.

Osborne, M. J. and Rubinstein, A. (1994). *A course in Game Theory*. The MIT Press.

Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050.

Rabin, M. (1993). Incorporating fairness into Game Theory and Economics. *The American Economic Review*, 83(5):1281–1302.

Roemer, J. E. (2010). Kantian equilibrium. *Scandinavian Journal of Economics*, 112(1):1–24.

Roemer, J. E. (2019). *How We Cooperate: A Theory of Kantian Optimization*. Yale University Press.

Tan, T. C. and Werlang, S. R. (1988). The Bayesian foundations of solution concepts of games. *Journal of Economic Theory*, 45(2):370–391.

von Neumann, J. (1959/1928). On the theoty of games of strategy [Zur theorie der gesellschaftsspiele]. In Tucker, A. W. and Luce, R. D., editors, *Contributions to the Theory of Games*, volume 4. Princeton University Press. Translated from German by Bargmann, S.

von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.

Walker, M. and Wooders, J. (2001). Minimax play at Wimbledon. *American Economic Review*, 91(5):1521–1538.