10 An Overview of Basic Concepts in Game Theory

Analysis of strategic questions—such as whether or to form an alliance, and whether to go to war in the first place—typically requires the use of game-theoretic models. We will exemplify in later lectures. Meanwhile, however, we set the scene by reviewing basic concepts in game theory.

Strategic behavior arises when the outcome of an individual's actions depends on actions taken by other individuals. For example, whether it is advantageous for drivers negotiating a 4-way junction to assume right of way depends on whether other drivers concede right of way. Likewise, whether one prospers from moral behavior depends on whether others do the right thing. If an interaction among individuals gives rise to strategic behavior and can be described mathematically, then we call this description a game, and each individual a player. Thus a game in the mathematician's sense is a model of strategic interaction, and game-theoretic modelling is the process by which such games are constructed. Correspondingly, game theory is a diverse assemblage of ideas, theorems, analytical methods and computational tools for the study of social interaction—according to Samuelson (1997, p. 4), more appropriately viewed as a language than as a theory. By using a variety of approaches, each with its own advantages and limitations, it aims to achieve a much better understanding of underlying processes and resulting patterns than is possible within a single modelling framework.

The players

A game has four key ingredients. First, there are at least two players, who may be either specific actors or individuals drawn randomly from a large population: drivers at a 4-way junction may be either neighbors or strangers. Correspondingly, the game is a either a community game or a population game. Often there are precisely two players, unsurprisingly called Player 1 and Player 2, whom we shall frequently denote by $\mathcal{P}_1$ and $\mathcal{P}_2$ for short, especially in tables. If the game is a community game, then these two players are specific individuals. If, on the other hand, the game is a population game, then Player 1 is an arbitrary focal individual, and Player 2 represents every other individual in the population.

The strategy set

Second, each player has a set of feasible plans of action—or strategies. For example, drivers at a 4-way junction can either Go or Wait; and that may be all they can do (in a model). Whenever (as in this particular example) the number of strategies is finite, we say that the game is discrete, and that each strategy is a pure strategy. If the number of possible pure strategies is small, then typically we distinguish them by letters or labels (e.g., G for Go, W for Wait); otherwise, we call them strategy 1, strategy 2, and so on. For a discrete community game (between specific actors), it makes sense to suppose that Player 1 has m_1 pure strategies and Player 2 has m_2 pure strategies, where in general $m_1 \neq m_2$; whereas, for a discrete population game, we would always set $m_1 = m_2 = m$. Although individuals may have a different number of options in a different role, it is assumed that individuals are ultimately[1] equally likely to occupy any role, and a strategy is a plan of action that prescribes what to do in every possible role.

Strategies are constrained by the information structure of an interaction. For example, drivers at a 4-way junction can condition their behavior on their lateness relative to others, but only if sufficiently aware of it. Let the latenesses of two such drivers, who wish to turn left simultaneously, be random variables X and Y taking values between 0 (unbelievably early) and 1 (desperately late); the sample space for their joint distribution is the unit square, $0 \le x, y \le 1$. Then it is possible for the first driver, or $\mathcal{P}_1$, to play a strategy u defined by "Go if $X > u$, Wait if $X \le u$" while the other driver, or $\mathcal{P}_2$, plays strategy v (Go if $Y > v$, Wait if $Y \le v$). Such strategies are often called threshold strategies. In general, we use S_i to denote the strategy set for $\mathcal{P}_i$ and D to denote the set of all feasible strategy combinations, which we call the decision set.[2] Usually, though not invariably,[3] D is a Cartesian product; in particular, for two players,

$$D \;=\; S_1 \times S_2. \tag{10.1}$$

For example, in the case of the threshold strategies discussed above, $D = [0,1] \times [0,1]$.

Here four remarks are in order. First, when lateness enters the picture, the strategy set becomes infinite: the game is no longer discrete. Instead we call it continuous. Second, Go (strategy 0) and Wait (strategy 1) both belong to the new continuous strategy set defined above. But we can extend a strategy set from discrete to continuous in more than one way—fortunately, because threshold strategies like u and v cannot be played by drivers who have no information about lateness. However, they still have the option of randomizing or "mixing" between their pure strategies. Specifically, it is possible for $\mathcal{P}_1$ to play a strategy p defined as selecting Go with probability p (and hence Wait with probability $1-p$), while $\mathcal{P}_2$ plays strategy q, that is, selects Go with probability q (and hence Wait with probability $1-q$). We call such strategies mixed strategies. Now we have two continuous strategy sets for our example, a mixed strategy set (in which Go is 1 and Wait is 0) and a threshold strategy set (in which Go is 0 and Wait is 1). Let us call the continuous game Crossroads I or Crossroads II, according to whether the strategies are mixed strategies or threshold strategies.

Third, any discrete game can be extended to a continuous game by mixing between the pure strategies of both players. We refer to the new continuous game as the *mixed extension* of the old discrete game. In this regard, it is convenient to define

$$s_i \;=\; m_i - 1, \quad i = 1, 2. \tag{10.2}$$

Then $\mathcal{P}_1$'s strategy set consists of probability vectors forming the s_1-dimensional simplex

$$\Delta^{s_1} \;=\; \{p \in \mathbb{R}^{m_1} | p_i \ge 0 \text{ for } i = 1, \ldots, m_1 \text{ and } \textstyle\sum_{i=1}^{m_1} p_i = 1\}. \tag{10.3a}$$

Correspondingly, $\mathcal{P}_2$'s strategy set is the s_2–dimensional simplex

$$\Delta^{s_2} \;=\; \{q \in \mathbb{R}^{m_2} | q_j \ge 0 \text{ for } j = 1, \ldots, m_2 \text{ and } \textstyle\sum_{j=1}^{m_2} q_j = 1\} \tag{10.3b}$$

[1]That is, sufficiently far back in time: You may be either male or female now, but way back in time—just before you were a twinkle in your parents' eyes—you could have turned out to be either.

[2]Simply because this phrase is so much less cumbersome than "strategy combination set."

[3]For an exception, see, e.g., Mesterton-Gibbons (2001, §1.5).

(and so the decision set is $D = \Delta^{s_1} \times \Delta^{s_2}$). Fourth, we have already implicitly adopted a convention of using consecutive letters of the alphabet for the players' strategies, p and q for mixed strategies, u and v for thresholds or other continuous strategies. We abide by this convention throughout.

Nevertheless, with regard to mixed strategies, we find it convenient to exploit the natural one-one correspondence between the simplex Δ^{s_1} and its projection onto the plane $p_{m_1} = 0$ (which forms a face of that simplex) by redefining $\mathcal{P}_1$'s strategy as a vector in $\mathbb{R}^{s_1}$. As soon as the probabilities with which $\mathcal{P}_1$ selects strategies 1 to s_1 have been specified, the probability that $\mathcal{P}_1$ selects strategy m_1 is necessarily determined as

$$p_{m_1} \;=\; 1 - \sum_{i=1}^{s_1} p_i. \tag{10.4}$$

So it suffices to keep track of the first s_1 such probabilities. Thus Player 1's strategy set becomes an s_1-dimensional subset of $\mathbb{R}^{s_1}$, as opposed to an s_1-dimensional subset of $\mathbb{R}^{m_1}$. It is convenient to adopt notation for this modified strategy set that suppresses its dependence on m_1 or s_1 (which is anyhow fixed at the outset). Accordingly, we adopt

$$\Delta_1 \;=\; \{p \in \mathbb{R}^{s_1} | p_i \geq 0 \text{ for } i = 1, \ldots, s_1 \text{ and } \textstyle\sum_{i=1}^{s_1} p_i \leq 1\}. \tag{10.5a}$$

Correspondingly, $\mathcal{P}_2$'s strategy set becomes

$$\Delta_2 \;=\; \{q \in \mathbb{R}^{s_2} | q_i \geq 0 \text{ for } i = 1, \ldots, s_2 \text{ and } \textstyle\sum_{i=1}^{s_2} q_i \leq 1\}, \tag{10.5b}$$

the probability that $\mathcal{P}_2$ selects strategy m_2 is necessarily determined as

$$q_{m_2} \;=\; 1 - \sum_{i=1}^{s_2} q_i \tag{10.6}$$

and the decision set becomes $D = \Delta_1 \times \Delta_2$. Often $s_1 = s_2 = s$, in which case we write $\Delta_1 = \Delta_2 = \Delta$ with $D = \Delta \times \Delta$. Moreover, in the important special case of only two pure strategies apiece for which $m = 2$ or $s = 1$, it is standard practice to use p and q in place of p_1 and q_1 (as for Crossroads I above).

The rewards

The third key ingredient of any game is a well defined reward to each player from every potential strategy combination. For a 2-player discrete community game, the rewards are given by two $m_1 \times m_2$ payoff matrices A and B, where a_{ij} and b_{ij} are the payoffs to $\mathcal{P}_1$ and $\mathcal{P}_2$, respectively, from pure strategy combination (i, j); for illustration, see Tables 10.1 and 10.2 below. For a 2-player discrete population game, however, a single payoff matrix A suffices, because $B = A^T$ by symmetry (as in Tables 10.1 and 10.2 with $\tau_1 = \tau_2$).

For a 2-player continuous community game, on the other hand, the rewards are given by a pair of functions defined on the set of all feasible strategy combinations. For the mixed extension of a discrete community game, the reward to $\mathcal{P}_i$ from strategy combination (p, q) is denoted by $f_i(p, q)$. Because $\mathcal{P}_1$'s reward from pure strategy i when $\mathcal{P}_2$ adopts strategy j is a_{ij}, and because $\mathcal{P}_2$ adopts strategy j with probability q_j, the unconditional reward to $\mathcal{P}_1$ from strategy i when $\mathcal{P}_2$ adopts mixed strategy $q = (q_1, \ldots, q_{s_2})$ is

$$\sum_{j=1}^{m_2} a_{ij} q_j \tag{10.7}$$

with q_{m_2} determined by (10.4). Player 1 obtains the above reward with probability p_i when adopting mixed strategy $p = (p_1, \ldots, p_{s_1})$, with p_{m_1} defined by (10.4), and so the unconditional reward to $\mathcal{P}_1$ from mixed strategy p when $\mathcal{P}_2$ adopts mixed strategy q is

$$f_1(p,q) \;=\; \sum_{i=1}^{m_1} \left\{ \sum_{j=1}^{m_2} a_{ij} q_j \right\} p_i \;=\; \sum_{i=1}^{m_1} \sum_{j=1}^{m_2} a_{ij} p_i q_j \;=\; (p, p_{m_1}) A (q, q_{m_2})^T \tag{10.8a}$$

where a superscripted T denotes transpose and we are using (p, p_{m_1}) and (q, q_{m_2}) to denote $(p_1, p_2, \ldots, p_{s_1}, p_{m_1}) \in \mathbb{R}^{m_1}$ and $(q_1, q_2, \ldots, q_{s_2}, q_{m_2}) \in \mathbb{R}^{m_2}$, respectively. Correspondingly, the unconditional reward to $\mathcal{P}_2$ from mixed strategy combination (p, q) is

$$f_2(p,q) \;=\; \sum_{j=1}^{m_2} \left\{ \sum_{i=1}^{m_1} b_{ij} p_i \right\} q_j \;=\; \sum_{i=1}^{m_1} \sum_{j=1}^{m_2} b_{ij} p_i q_j \;=\; (p, p_{m_1}) B (q, q_{m_2})^T. \tag{10.8b}$$

For any other continuous game, however, we instead use $f_i(u, v)$ to denote the reward to Player i from strategy combination (u, v).[4] Moreover, for a 2-player continuous population game, a single reward function suffices because $f_2(u, v) = f_1(v, u)$ by symmetry, and it is usual to denote this function simply by f (as opposed to f_1).

To illustrate: For Crossroads II regarded as a population game, let τ be the time it takes a driver to traverse the junction, and suppose that drivers discount this delay by a fraction η of their earliness. Thus if $X \le u$ and $Y > v$, then the delay of τ as $\mathcal{P}_2$ traverses the junction is experienced as $-\tau\{1 - \eta(1 - X)\}$ by $\mathcal{P}_1$; whereas if $X > u$ and $Y \le v$, then the delay of τ is experienced as $-\tau\{1-\eta(1-Y)\}$ by $\mathcal{P}_2$, but is 0 for $\mathcal{P}_1$. If both drivers either Go or Wait in the first instance, then there is an additional delay of δ or ϵ, respectively, with $\delta > \epsilon$, as they sort out who will subsequently drive away first. Assuming that it is equally likely to be either driver and that delays should be as short as possible, so that negatives of delays serve as payoffs, $\mathcal{P}_1$'s payoff is the random variable

$$F_1(X, Y, u, v) \;=\; \begin{cases} -\tau\{1 - \eta(1 - X)\} & \text{if } X \le u, Y > v \\ -\left(\delta + \frac{1}{2}\tau\{1 - \eta(1 - X)\}\right) & \text{if } X > u, Y > v \\ -\left(\epsilon + \frac{1}{2}\tau\{1 - \eta(1 - X)\}\right) & \text{if } X \le u, Y \le v \\ 0 & \text{if } X > u, Y \le v. \end{cases} \tag{10.9}$$

$\mathcal{P}_1$'s reward from the strategy combination (u, v) is the expected value of F_1, which we denote by $f_1(u, v)$. That is, using E to denote expectation,

$$f_1(u,v) \;=\; \mathrm{E}[F_1(X, Y, u, v)] \;=\; \int_0^1 \int_0^1 F_1(x, y, u, v) g(x) g(y) \, dx \, dy \tag{10.10}$$

[4]These functions are typically continuous, although in some instances they have isolated discontinuities (that is, are discontinuous across a set of measure zero).

where g is the probability density function of X and Y's common distribution. By symmetry, $\mathcal{P}_2$'s reward is $f_2(u,v) = f_1(v,u)$.

For Crossroads I regarded as a community game, we must replace (10.9) by

$$F_1 = \begin{cases} -\delta - \frac{1}{2}\tau_2 & \text{if } \Pi_1 = G, \Pi_2 = G \\ 0 & \text{if } \Pi_1 = G, \Pi_2 = W \\ -\tau_2 & \text{if } \Pi_1 = W, \Pi_2 = G \\ -\epsilon - \frac{1}{2}\tau_2 & \text{if } \Pi_1 = W, \Pi_2 = W \end{cases}$$

and (10.10) by

$$\begin{aligned} f_1(p,q) = \mathrm{E}[F_1] &= (-\delta - \tfrac{1}{2}\tau_2) \cdot \mathrm{Prob}(F_1 = -\delta - \tfrac{1}{2}\tau_2) + 0 \cdot \mathrm{Prob}(F_1 = 0) \\ &\quad + (-\tau_2) \cdot \mathrm{Prob}(F_1 = -\tau_2) + (-\epsilon - \tfrac{1}{2}\tau_2) \cdot \mathrm{Prob}(F_1 = -\epsilon - \tfrac{1}{2}\tau_2) \\ &= \left(\epsilon + \tfrac{1}{2}\tau_2 - \{\delta + \epsilon\}q\right) p + \left(\epsilon - \tfrac{1}{2}\tau_2\right) q - \epsilon - \tfrac{1}{2}\tau_2 \end{aligned} \tag{10.11}$$

where τ_2 is how long it takes $\mathcal{P}_2$ to traverse the junction and Π_i is the pure strategy selected by Player i, a random variable with binomial distribution having parameter p or q, so that $\mathrm{Prob}(\Pi_1 = G, \Pi_2 = G) = p \cdot q$, and so on: we assume that strategies are chosen independently. Likewise, the reward to $\mathcal{P}_2$ from strategy combination (p,q) is

$$f_2(p,q) = \left(\epsilon + \tfrac{1}{2}\tau_1 - \{\delta + \epsilon\}p\right) q + \left(\epsilon - \tfrac{1}{2}\tau_1\right) p - \epsilon - \tfrac{1}{2}\tau_1 \tag{10.12}$$

where τ_1 is how long it takes $\mathcal{P}_2$ to traverse the junction. Note that if $\tau_1 = \tau_2$ (as when Crossroads I is instead regarded as a population game), then $f_2(p,q) = f_1(q,p)$.

Nash equilibrium: A solution concept

The last main ingredient of any game is a solution concept. An appropriate solution concept for community games is that of Nash equilibrium, a strategy combination from which no individual has a unilateral incentive to depart. Equivalently, a Nash equilibrium is a combination of mutual best replies.

For a 2-player discrete community game in which $\mathcal{P}_1$ is the row player and $\mathcal{P}_2$ is the column player, strategy combination (i,j) is a Nash equilibrium if

$$a_{ij} \geq a_{kj} \tag{10.13a}$$

for all $k = 1, \ldots, m_1$ and

$$b_{ij} \geq b_{il} \tag{10.13b}$$

for all $l = 1, \ldots, m_2$. If (10.13a) is satisfied with strict inequality for all $k \neq i$ and (10.13b) is satisfied with strict inequality for all $l \neq j$ (a total of $m_1 + m_2 - 2$ conditions), then the Nash equilibrium is said to be strong; otherwise (that is, if even a single one of the $m_1 + m_2 - 2$ conditions is satisfied with equality), the Nash equilibrium is said to be weak. A special case of a strong Nash equilibrium occurs when both strategy i is a "strongly dominant" strategy for $\mathcal{P}_1$ and strategy j is a strongly dominant strategy for $\mathcal{P}_2$, that is, when both $a_{il} > a_{kl}$ for all $k \neq i$ for all $l = 1, \ldots, m_2$ (not only $l = j$) and $b_{kj} > a_{kl}$ for all $l \neq j$ for

Table 10.1: Payoff matrix A

	$\mathcal{P}_2$	
	G	W
$\mathcal{P}_1$ G	$-\delta-\frac{1}{2}\tau_2$	0
$\mathcal{P}_1$ W	$-\tau_2$	$-\epsilon-\frac{1}{2}\tau_2$

Table 10.2: Payoff matrix B

	$\mathcal{P}_2$	
	G	W
$\mathcal{P}_1$ G	$-\delta-\frac{1}{2}\tau_1$	$-\tau_1$
$\mathcal{P}_1$ W	0	$-\epsilon-\frac{1}{2}\tau_1$

all $k = 1, \ldots, m_1$ (not only $k = i$). For example, if we restrict Crossroads I to the two pure strategies G and W, then G is a strongly dominant strategy for $\mathcal{P}_1$ when $\delta < \frac{1}{2}\tau_2$ because $a_{11} > a_{21}$ and $a_{12} > a_{22}$; see Table 10.1. Likewise, G is a strongly dominant strategy for $\mathcal{P}_2$ when $\delta < \frac{1}{2}\tau_1$ because $b_{11} > b_{12}$ and $b_{21} > b_{22}$ (Table 10.2). Thus strategy combination (G, G) is a strong Nash equilibrium for $\delta < \frac{1}{2}\min(\tau_1, \tau_2)$.

Correspondingly, for the mixed extension with m_i pure strategies for Player i, strategy combination $(p^*, q^*) \in \Delta_1 \times \Delta_2$ is a Nash equilibrium if

$$f_1(p^*, q^*) \;\geq\; f_1(p, q^*) \tag{10.14a}$$

for all $p \in \Delta_1$ and

$$f_2(p^*, q^*) \;\geq\; f_2(p^*, q) \tag{10.14b}$$

for all $q \in \Delta_2$, where Δ_1, Δ_2, f_1 and f_2 are defined by (10.5) and (10.8). If (10.14a) is satisfied with strict inequality for all $p \neq p^*$ and (10.14b) is satisfied with strict inequality for all $q \neq q^*$, then the Nash equilibrium is said to be strong; otherwise it is said to be weak. At least one Nash equilibrium always exists for such a game (Nash, 1951), so existence is never an issue—but uniqueness frequently is. Even for moderate values of m_1 and m_2, there can exist numerous Nash equilibria, and often it is not considered useful to find and characterize them all; rather, which ones are worth examining is determined by the particular question for which the game was constructed in the first place.