

## 2 Game Theory I: Simultaneous-Move Games

*You are deciding whether or not to contribute your resources to a shared computer cluster. The amount you contribute, along with the amount that others contribute, will affect the resources you can consume. What should you do?*

Game theory provides a formal framework with which to reason about situations of strategic interdependence. In introducing game theory, we start with preferences and utility theory, and then define the normal form representation of a simultaneous-move game. We introduce important solution concepts, including Nash equilibrium and dominant-strategy equilibrium, and also consider the class of potential games and congestion games.

Many of the examples that we adopt in introducing game theory are quite simple, but the techniques can be applied much more generally.

For example, game theory is a useful tool for the design and analysis of:

- Reputation systems (will buyers provide negative feedback, or worry about retaliatory negative feedback from the seller?),
- Internet security (will firms adopt new standards or reason that there will be no benefit unless others follow?), and
- Meeting scheduling systems (will users game the outcomes by submitting false information about preferences and constraints?)

For now, think about game theory as providing a mathematical way of reasoning about settings where each participant is self interested and takes actions to obtain an outcome that is the best possible for himself, given how others are acting.

### 2.1 Introduction

To fix ideas, let's consider the often told story of the *Prisoner's Dilemma*. This is a typical, simple example that nevertheless reveals some of the interesting aspects of reasoning about situations of strategic interdependence.

**Example 2.1** (Prisoner's Dilemma.). *Two people are arrested and accused of a crime. Each can cooperate  $C$  and not admit to the crime, or defect  $D$  and admit. If both cooperate, then they receive a minor charge and stay in prison for 2 years. If one person defects while the other cooperates, the defector is released (0 years in prison) while the other serves a 5-year sentence. If both defect, then they both go to prison but with early parole and serve a 4-year sentence. Figure 2.1 shows the payoff table for this game. Player one is the row player and player two the column player. In each entry, the first number represents the "payoff" to row and the second number is the "payoff" to column.*

## 2 Game Theory I: Simultaneous-Move Games

		Player 2	
		C	D
Player 1	C	-2, -2	-5, 0
	D	0, -5	-4, -4

Figure 2.1: The *Prisoner's Dilemma* Game. Each year in prison results in a payoff of -1.

In this particular game, the payoff corresponds to the number of years spent in prison: each year results in a negative payoff of -1. More generally, we can think of payoff tables as encoding the preference of participants for different outcomes, with higher payoffs indicating more preferred outcomes.

What should a player do in this game? A moment's reflection suggests an obvious answer: play *D*! If column plays *C* then row's best response is *D* (0 years in prison instead of 2). If column plays *D* then row's best response is also *D* (4 years in prison instead of 5). In particular, *D* is a *dominant* action, it is the best action whatever the action of the other player. The *dilemma* is that by both players defecting the outcome is worse for both people (namely 4 years in prison) than if they both cooperated (2 years in prison).

The *Prisoner's Dilemma* illustrates how game-theory provides a way to model a situation of strategic interdependence. We refer to the participants in a game as *agents* or *players*. Crucially, each agent is free to make its own decision about how to act. In the real-world, an agent could be a person deciding whether to leave feedback on eBay or Amazon, a firm deciding whether to install a new security protocol, or an automated bidding system such as those that caused cyclic behavior in early sponsored search auctions.

It will often make sense to model agents as selfish, for example minimizing payments in an auction or participating on a social network in order to become an influencer, and turn this influence into personal profit. However, selfish preferences are not essential to game theory and the payoffs to an individual player can also represent social (or *other-regarding*) preferences. For example, game theory can model the behavior of a user providing feedback about a hotel on a reputation platform, where the user's motivation might be to help other users or help the owners of the hotel.

## 2.2 Preferences and Utility

Economic theory suggests that an agent acts in a way that promotes outcomes that are as preferable for the agent as possible. For this, we define a *preference order* on outcomes. Let  $O$  denote a set of outcomes, and for any two outcomes, let  $o_1 \succeq o_2$  denote that an agent *weakly prefers*  $o_1$  to  $o_2$ . If both  $o_1 \succeq o_2$  and  $o_2 \succeq o_1$ , then the agent is *indifferent*, and we write  $o_1 \sim o_2$ . If  $o_1 \succeq o_2$ , and the agent is not indifferent then  $o_1$  is *strictly preferred* to  $o_2$ , and we write  $o_1 \succ o_2$ .

**Example 2.2.** *Suppose a student is trying to decide whether he prefers ( $o_1$ ) a larger apartment with plenty of light in the suburbs, ( $o_2$ ) a small, modern studio, centrally located, or ( $o_3$ ) a shared, older house, close to campus, with people that he knows reasonably well. These outcomes differ along many different dimensions.*

We insist that agents have a *complete* preference order. For every outcome pair,  $o_1, o_2$ , we

must have at least  $o_1 \succeq o_2$ , or  $o_2 \succeq o_1$ , and both if the agent is indifferent. This precludes a partial order; e.g., with  $o_1 \succeq o_3$  and  $o_2 \succeq o_3$  but no preference defined between  $o_1$  and  $o_2$ . We further require *transitivity*, which ensures that if  $o_1 \succeq o_2$  and  $o_2 \succeq o_3$  then  $o_1 \succeq o_3$ . This precludes cyclic preferences, such as  $o_1 \succeq o_2$ ,  $o_2 \succeq o_3$  and  $o_3 \succ o_1$ .

But a preference order need not provide, by itself, enough information to explain how a rational agent should behave. Returning to the example, suppose that  $o_3 \succ o_2 \succ o_1$ , and that there are two available actions:

- action 1, which leads to the shared house ( $o_3$ ) with probability 0.7 and otherwise the large apartment ( $o_1$ ).
- action 2, which leads to the studio ( $o_2$ ) with probability 1.

Given the uncertainty of the outcome of action 1, the best decision is not clear just from the preference order, because it depends on the intensity of preference for  $o_3$  over  $o_2$ , and for  $o_2$  over  $o_1$ .

In addressing this, *utility theory* associates a *utility*  $u(o)$  with each outcome  $o \in O$ , and insists that the utility is consistent with an agent's preference order. Consistency requires  $u(o_1) \geq u(o_2)$  if and only if  $o_1 \succeq o_2$ , for two outcomes  $o_1, o_2$ , and  $u(o_1) = u(o_2)$  if and only if  $o_1 \sim o_2$ . A *utility function*  $u : O \mapsto \mathfrak{R}$  (where  $\mathfrak{R}$  denotes the set of real numbers) assigns a utility to every outcome. Given this, the decision of a *rational* agent is the one that maximizes the *expected utility*. If outcome  $o_j \in O$  occurs with probability  $p_j \geq 0$ , and there are  $k$  possible outcomes, then the expected utility is  $\sum_{j=1}^k p_j u(o_j)$ .

In the example, suppose that the student's utility is  $u(o_1) = 700$ ,  $u(o_2) = 800$ , and  $u(o_3) = 1000$ , which is consistent with preference order  $o_3 \succ o_2 \succ o_1$ . The utility for an outcome represents some combination of the cost of rent and the intrinsic happiness from the living situation. Based on this, action 1 has expected utility  $0.3u(o_1) + 0.7u(o_3) = 910$ , compared to expected utility  $u(o_2) = 800$  for action 2, and the rational decision is action 1.

There is no unique utility function to 'explain' an agent's preferences, but rather a family of possible functions. If  $u(o)$  is a utility function, then  $u'(o) = a \cdot u(o) + b$ , for constants  $a \in \mathfrak{R}_{>0}$  and  $b \in \mathfrak{R}$ , is another utility function that provides the same preference order on outcomes, and the same preference order on distributions on outcomes; notation  $\mathfrak{R}_{>0}$  is the set of strictly positive real numbers. We say that utility function  $u'$  is a *positive affine* transformation of the original utility function.

In addition, it is worth emphasizing that it is not necessary that agents know their utility functions for utility theory to be a predictive theory. Rather, it is sufficient that agents act as if they are maximizing some utility function.

## 2.3 Simultaneous-Move Games

Let's now consider situations with multiple agents, where an agent's utility depends on its own actions as well as the actions of others. This brings us to game theory.

For this purpose, we assume that the outcome of a game is completely determined by the actions of all agents, called the *action profile*. Given this, we model agents as having utility functions on action profiles.

We have already seen the example of the *Prisoner's Dilemma*. This is a two player, two

action game.

**Definition 2.1** (Simultaneous-move game). A simultaneous-move game  $(N, A, u)$  has

- $N = \{1, \dots, n\}$  agents, indexed by  $i$ .
- $A = A_1 \times \dots \times A_n$ , where  $A_i$  is a set of actions available to agent  $i$ , and  $a = (a_1, \dots, a_n) \in A$  denotes an action profile.
- $u = (u_1, \dots, u_n)$ , where  $u_i : A \mapsto \mathfrak{R}$  is a utility function (or payoff function) for agent  $i$ , and assigns a utility (or payoff) to every action profile  $a \in A$ .

The action sets in the *Prisoner's Dilemma* are  $A_1 = A_2 = \{C, D\}$ , there are two agents  $N = \{1, 2\}$ , and the utility functions  $u_1(a_1, a_2)$  and  $u_2(a_1, a_2)$  are as defined in Figure 2.1. For example,  $u_1(C, C) = -2, u_1(C, D) = -5, u_1(D, C) = 0, u_1(D, D) = -4$ . The joint action set  $A = A_1 \times A_2 = \{C, D\} \times \{C, D\}$  and examples of action profiles are  $a = (C, C)$  and  $a = (D, C)$ .

A *finite* game has a finite number of players, each with a finite set of actions. However, in general, the action sets can be infinite (e.g., the action might represent the fraction of a resource demanded by a player) and we might even want to model games with an infinite number of players, capturing limiting behavior as the number of players becomes very large.

By assumption, every agent knows the available actions and all utility functions, knows that every agent knows this, knows that every agent knows that every agent knows this, and so on, *ad nauseum*. Formally, the actions and utilities are said to be *common knowledge*.

On the other hand, an action is selected by each agent without knowledge about the actions selected by other agents. This is what makes this a simultaneous-move game. It is not important that the actions are taken exactly at the same time, only that one agent does not know the action selected by another agent when selecting its own action.

### 2.3.1 The Normal form representation

The standard representation for (finite) simultaneous-move games is the *normal form*. In a normal form representation, the actions, and payoffs of each agent for every action profile, are explicitly represented. We see this in the tabular form of Figure 2.1, which is also called the *payoff matrix*. An action profile  $a \in A$  identifies an element of the payoff matrix, with  $u_i(a)$  corresponding to the  $i$ th value of the entry.

Although the normal form is most typically adopted for simultaneous-move games, and most useful for simultaneous-move games, the normal form can also represent more complex situations. For example, by associating the 'action' in the normal form with a description of the way the player will act given observations of earlier play, the normal form can also model sequential-move games. But the normal form representation has some disadvantages for richer contexts such as sequential-move games. We delay any further discussion of this until Chapter 3.

One concern about the normal form is that because it is explicit in the way it represents actions and payoffs, it quickly becomes unmanageable for modeling large games. This is because a game with  $n$  agents, and  $m$  actions per agent, has  $m^n$  action profiles. The explicit representation of payoffs for all agents requires  $n \cdot m^n$  numbers; i.e., one payoff for each agent for each action profile. The number of action profiles grows exponentially in the number of agents. See Figure 2.2 for an illustration of the representation for a 3 agent, 2-action game.




Figure 2.2: The normal form representation for a game with 3 agents and 2 actions per agent.

The number number of actions available to agents can also grow exponentially in the size of a natural description of the game. This occurs, for example, in a network flow game such as that used to illustrate Braess' Paradox in Chapter 1. In that setting, there is an action for every possible path from start to end.

Certainly we can imagine more succinct representations of games, and we will see examples in the model of *congestion games* in Section 2.7 and also in *graphical games* in Chapter 5.

### 2.3.2 Pareto optimality

We have suggested that the players in the *Prisoner's Dilemma* will both play defect, and that this is unfortunate because both-cooperate would be a preferable outcome for both players. In particular, action profile  $(D, D)$  is Pareto dominated by  $(C, C)$ :

**Definition 2.2** (Pareto dominated). *An action profile  $a \in A$  is Pareto dominated by action profile  $a' \in A$  if  $u_i(a') \geq u_i(a)$  for all agents  $i \in N$  and  $u_i(a') > u_i(a)$  for some agent  $i \in N$ .*

**Definition 2.3** (Pareto optimality). *An action profile  $a \in A$  is Pareto optimal if there is no action profile  $a' \in A$  that Pareto dominates  $a$ .*

In particular, an action profile is Pareto optimal if no agent can be made better off without making some other agent worse off. For example, outcomes  $(C, C)$ ,  $(D, C)$  and  $(C, D)$  are all Pareto optimal in the *Prisoner's Dilemma*. Pareto optimality cannot be used to make predictions regarding how agents will behave. Rather, Pareto optimality provides a minimal criterion for whether or not an outcome is good from the perspective of social welfare.

Pareto optimality also extends to distributions on action profiles. A distribution on action profiles is Pareto optimal if there is no other distribution that provides one agent with strictly greater expected utility without giving another agent strictly less.

**Example 2.3.** *A distribution on action profiles in the Prisoner's Dilemma where agents play  $(C, C)$  with probability 0.5 and  $(D, C)$  with probability 0.5 is Pareto optimal. To prove this, we must show that any other distribution provides strictly less expected utility to at*

least one player. The current expected utility is -1 to player 1 and -3.5 to player 2. First, any increase in probability on  $(C, C)$  is worse for player 1, and any increase in probability on  $(D, C)$  is worse for player 2. Second, any shift in probability from  $(C, C)$  or  $(D, C)$  to one or both of  $(C, D)$  or  $(D, D)$  is worse for player 1, since player 1's payoffs are -2 and 0 for  $(C, C)$  and  $(D, C)$  but -5 and -4 for  $(C, D)$  and  $(D, D)$ .

We will soon return to randomized play in games, but for now we will continue assuming that each agent just picks a particular action.

## 2.4 Dominant-Strategy Equilibrium

Game theory seeks to predict how players will behave in a situation of strategic interdependence. An especially simple case occurs when the game has a *dominant-strategy equilibrium*, which is the situation in the *Prisoner's Dilemma*.

Let notation  $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n) \in A_{-i}$  denote the actions chosen by other agents than  $i$ . Here,  $A_{-i}$  is the joint action set not including agent  $i$ . Given this, we adopt  $u_i(a_i, a_{-i})$  to denote the payoff to agent  $i$  when its action is  $a_i$  and the other agents play  $a_{-i}$ .

**Definition 2.4** (Dominant-strategy equilibrium). *Action profile  $a^* = (a_1^*, \dots, a_n^*)$  is a dominant-strategy equilibrium of a simultaneous-move game  $(N, A, u)$  if, for all  $i$ ,*

$$u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i}), \quad \text{for all } a_i \in A_i, a_{-i} \in A_{-i}. \quad (2.1)$$

We adopt the convention that equilibrium action profiles are denoted with  $*$ . In words, an action profile  $a^*$  is a dominant-strategy equilibrium if every agent  $i$  maximizes its utility with its action  $a_i^*$  whatever the other agents do. For example, the action profile all-defect is a dominant-strategy equilibrium in the *Prisoner's Dilemma* because each player prefers action  $D$  over  $C$  whatever the action of the other player.

## 2.5 Nash Equilibrium

Games with dominant-strategy equilibria are easy to analyze because each agent has the same best (or *dominant*) action, whatever the behavior of other agents. There is no need for an agent to reason about how other agents will behave when deciding how to act. But the more typical case is that there is no dominant-strategy equilibrium.

Here's an example of such a game.

**Example 2.4.** *Figure 2.3 depicts a two player, three action game. Row plays  $\{U, M, D\}$  (up, middle, down) and column  $\{L, M, R\}$  (left, middle, right). A moment's inspection reveals there is no dominant-strategy equilibrium. But, there is another way to understand how rational players should act. Action  $M$  is dominated by action  $R$  for the column player: whatever row does,  $R$  is a better response for column than  $M$ . Based on this, action  $M$  can be eliminated from consideration, and now row can reason that action  $U$  dominates action  $M$  and  $D$  as long as column only selects  $L$  or  $R$ . Finally, if row will play  $U$  then column's best action is  $L$  and we identify  $(U, L)$  as the predicted outcome. But outcome  $(U, L)$  is not*

		Player 2		
		L	M	R
Player 1	U	4, 3	5, 1	6, 2
	M	2, 1	8, 4	3, 6
	D	3, 0	9, 6	2, 8

Figure 2.3: A game without a dominant strategy equilibrium but solvable by iterated elimination of strictly-dominated actions.

a dominant-strategy equilibrium. Certainly, if column plays M then row would not want to play U.

The procedure of iterated elimination of strictly-dominated actions is illustrated in Algorithm 2.1. This procedure is the first algorithmic tool we introduce for the strategic analysis of games. There are other variations, including iterated elimination of weakly-dominated actions. The procedure, its properties, and properties of variations are developed in Exercise 2.1.

---

**Algorithm 2.1:** Iterated elimination of strictly-dominated actions.

---

**Input:** Simultaneous-move game  $G = (N, A, u)$   
**Variables:**  $R_i$  (for each agent  $i$ ): the set of undominated actions for agent  $i$   
**begin**  
  **foreach** agent  $i \in N$  **do**  
    ⌊ Initialize  $R_i := A_i$   
    **bool** DominatedActionFound: = true  
    **while** *DominatedActionFound* **do**  
      **if** there exists some agent  $i$ , some action  $a_i \in R_i$ , and some action  $a'_i \in R_i$   
      such that  $u_i(a_i, a_{-i}) < u_i(a'_i, a_{-i})$  for all  $a_{-i} \in R_{-i}$  **then**  
        ⌊ remove action  $a_i$  from  $R_i$   
      **else**  
        ⌊ DominatedActionFound: = false  
    output  $(R_1, \dots, R_n)$

---

Unfortunately, iterated elimination will not always terminate with a single action profile. Indeed, the same exercise asks for an example where it does not eliminate even a single action. For this reason, this iterated elimination procedure does not provide a way to predict the behavior in general simultaneous-move games.

For this, we need the concept of a *Nash equilibrium*. For now we focus on pure strategies, in which agents act without randomizing over actions.

**Definition 2.5** (Pure strategy Nash equilibrium). *Action profile*  $a^* = (a_1^*, \dots, a_n^*)$  is a *pure-strategy Nash equilibrium (PSNE)* of simultaneous-move game  $(N, A, u)$  if, for all  $i$ ,

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*), \quad \text{for all } a_i \in A_i. \quad (2.2)$$

In words, an action profile is a Nash equilibrium if every agent maximizes its utility given that the other agents behave according to the action profile. Every agent is best-responding to the behavior of every other agent, and no agent has a useful deviation. The crucial distinction from a dominant-strategy equilibrium is that each agent’s action is only sure to be a best response when the other agents play the equilibrium.

For example, action profile  $(U, L)$  is a Nash equilibrium of the game in Figure 2.3. Each player is best-responding to the other player. In contrast to a dominant-strategy equilibrium, for a Nash equilibrium to be a sensible prediction of behavior in a game, agents must believe that other agents are rational. For example, if column was to play  $M$  then row certainly wouldn’t want to play  $U$ . In fact, each player needs to believe that the other players believe that all players are rational, and so on. Why else should a rational column player play  $L$ ? For Nash equilibrium to make sense there must be common knowledge of rationality.

Another difficulty with the concept of Nash equilibrium is that there can be multiple Nash equilibria. We see this in the game of *Chicken* in Section 2.6.3. In games with multiple Nash equilibria it is often unclear how agents should reason about which equilibria will be played.

## 2.6 Mixed-Strategy Nash Equilibrium

Let’s look at another example.

**Example 2.5** (Matching Pennies Game). *In the Matching Pennies game, each player places a penny ‘heads up’ or ‘tails up.’ Row tries to match: if the actions of the two players are the same, then row takes both pennies, and gains a penny while row loses a penny. Column tries not to match: if the actions are different, column takes both pennies and gains a penny. See Figure 2.4. There is no Nash equilibrium without randomized actions: from every action profile there is an agent with a useful deviation. What should the players do?*

		Player 2	
		H	T
Player 1	H	1, -1	-1, 1
	T	-1, 1	1, -1

Figure 2.4: Matching Pennies Game.

The *Matching Pennies* game illustrates that there may not exist a pure-strategy Nash equilibria. But what if we allow agents to randomize over actions and adopt a *mixed strategy*?

**Definition 2.6** (Mixed strategy). *A mixed strategy  $s_i : A_i \mapsto [0, 1]$  for agent  $i$  assigns a probability  $s_i(a_i) \geq 0$  to each action  $a_i \in A_i$ , with the sum  $\sum_{a_i \in A_i} s_i(a_i) = 1$ , so that  $s_i$  is a well defined probability distribution on actions.*

In words, a mixed strategy  $s_i$  assigns a probability  $s_i(a_i)$  to each action  $a_i$ . For example, in *Matching Pennies*, a mixed strategy for agent 1 is  $s_1(H) = 0.4, s_1(T) = 0.6$ , such that the agent plays  $H$  with probability 0.4 and  $T$  with probability 0.6. We can represent this



strategy through a vector of probabilities, writing  $s_1 = (0.4, 0.6)$ . Mixed strategies include pure strategies as a special case; e.g.,  $s_2 = (1, 0)$  is a strategy for agent 2 that plays action  $H$  with probability 1.

Given mixed strategies  $s_1 = (0.4, 0.6)$  for agent 1 and  $s_2 = (1, 0)$  for agent 2, the expected utility to agent 1 is:

$$\begin{aligned} u_1(s_1, s_2) &= p(H, H)u_1(H, H) + p(H, T)u_1(H, T) + p(T, H)u_1(T, H) + p(T, T)u_1(T, T) \\ &= (0.4)(1)(1) + (0.4)(0)(-1) + (0.6)(1)(-1) + (0.6)(0)(1) = -0.2, \end{aligned}$$

where  $p(a_1, a_2)$  is the probability that action profile  $(a_1, a_2)$  is played for strategies  $(s_1, s_2)$ . The probability of an action profile such as  $(H, H)$ , where both players 1 and 2 play  $H$ , is given by the product  $p(H, H) = s_1(H)s_2(H) = (0.4)(1) = 0.4$ .

Given strategy profile  $s = (s_1, \dots, s_n)$ , let

$$u_i(s) = \sum_{(a_1, \dots, a_n) \in A} p(a_1, \dots, a_n) u_i(a_1, \dots, a_n), \quad (2.3)$$

denote the expected utility to agent  $i$ , with probability  $p(a_1, \dots, a_n) = s_1(a_1) \cdot s_2(a_2) \dots \cdot s_n(a_n)$  for each action profile  $(a_1, \dots, a_n)$ . A Nash equilibrium can now be defined for mixed strategies:

**Definition 2.7** (Mixed-strategy Nash equilibrium). *Strategy profile  $s^* = (s_1^*, \dots, s_n^*)$  is a mixed-strategy Nash equilibrium in game  $(N, A, u)$  if, for all  $i$ ,*

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \quad \text{for all mixed strategies } s_i. \quad (2.4)$$

In words, every agent  $i$  maximizes its expected utility by adopting strategy  $s_i^*$ , given that the other agents play their mixed strategies  $s_{-i}^*$ .

The following theorem, due to John Nash in 1951, provides the main theoretical grounding for game theory.

**Theorem 2.1** (Existence of Mixed-Strategy Nash Equilibrium). *Every finite simultaneous-move game  $(N, A, u)$  has at least one mixed-strategy Nash equilibrium.*

The proof is beyond the scope of this book, but references are provided in the chapter notes. Given this seminal result, we can model agents as best-responding to the play of other agents and know that this is always possible, in the sense that such a strategy profile always exists.

Note: An agent's preferences, even on distributions of outcomes, are invariant to positive affine transformations of utility (see Section 2.2). Because of this, the Nash equilibria of a game are unchanged under these transformations. Multiplying any player's payoffs by a positive number, and adjusting them up or down by a constant, leaves the equilibria of the game unchanged.

### 2.6.1 Best-response Analysis

Let's return to *Matching Pennies*. To identify an equilibrium, we define the *best-response correspondence* of each player. A correspondence is a function that maps to sets.




Figure 2.5: Best-response correspondences: (a) *Prisoner's Dilemma*. (b) *Matching Pennies*.

For this, let  $p$  denote the probability with which row (player 1) plays  $H$ , and  $q$  denote the probability with which column (player 2) plays  $H$ .

For player 1, the best-response  $p \in f(q)$  is a utility-maximizing probability for action  $H$ , given that 2 plays  $H$  with probability  $q$ . For player 2, the best-response  $q \in g(p)$  is a utility-maximizing probability for action  $H$ , given that 1 plays  $H$  with probability  $p$ .

For a (mixed) Nash equilibrium, we require probabilities  $(p^*, q^*)$  such that,

$$p^* \in f(q^*), \quad q^* \in g(p^*), \tag{2.5}$$

so that each player is best-responding to the other player.

There is a simple graphical approach to find such an equilibrium. We vary  $p$  on the x-axis and plot  $g(p)$  on the y-axis, and vary  $q$  on the y-axis and plot  $f(q)$  on the x-axis. Where these two intersect on some  $(p^*, q^*)$ , we have  $q^* \in g(p^*)$  from the plot of  $g(p)$  and  $p^* \in f(q^*)$  from the plot of  $f(q)$ .

This is illustrated for *Prisoner's Dilemma* and *Matching Pennies* in the next example. We see in Figure 2.5 (b) that the best-response correspondence  $g(p)$  returns any probability  $q \in [0, 1]$  for  $p = 0.5$ . This illustrates the indifference that can occur across actions.

**Example 2.6.** Consider Figure 2.5 (a), for the *Prisoner's Dilemma*. This plots player 2's best response  $g(p)$  on the y-axis and player 1's best response  $f(q)$  on the x-axis. The lines intersect at  $(p^*, q^*) = (0, 0)$ , corresponding to  $(D, D)$ . Indeed, this is the unique Nash equilibrium of *Prisoner's Dilemma*. Since each player has a dominant strategy, the best-response correspondences take on the same value for all strategies of the other player.

Consider Figure 2.5 (b), for the game of *Matching Pennies*. In this case, we see one intersection at  $(p^*, q^*) = (0.5, 0.5)$ , corresponding to each player mixing 50:50 over  $H$  and

*T. This is the unique Nash equilibrium. When player 1 plays  $p = 0.5$  then player 2 is indifferent between  $H$  and  $T$  and has as a best-response  $q = 0.5$ . Similarly for player 2. In particular, each player's mixed strategy makes the other player indifferent over both pure actions, and thus willing to play each action with some probability.*

A well known real-world example to motivate mixed strategies comes from penalty kicks in soccer, where the goal-keeper dives left or right and the kicker simultaneously kicks left or right. The goal-keeper is like the row player in *Matching Pennies* and wants to match, the kicker like the column player. Any fixed action could be anticipated and exploited by the other player. By randomizing, neither player can exploit knowledge of the strategy adopted by the other player.

### 2.6.2 The Support of a Mixed Strategy

A useful observation about mixed-strategy Nash equilibria comes from reasoning about the support of each player's strategy:

**Definition 2.8** (Support). *The support of mixed strategy  $s_i$ ,  $\sigma(s_i) = \{a_i : s_i(a_i) > 0, a_i \in A_i\}$ , is all actions played with strictly positive probability.*

Given this, a strategy profile  $s^*$  is a mixed-strategy Nash equilibrium if, and only if, for all agents  $i$ ,

$$u_i(a_i, s_{-i}^*) = u_i(a'_i, s_{-i}^*) \geq u_i(a''_i, s_{-i}^*), \quad (2.6)$$

for all  $a_i, a'_i \in \sigma(s_i)$  and all  $a''_i \notin \sigma(s_i)$ , where  $u_i(a_i, s_{-i}^*)$  is the expected utility of agent  $i$  for action  $a_i$  given the mixed strategy of other agents.

For sufficiency, note that every agent  $i$  is indifferent across all actions it is mixing over (the actions in its support) and weakly prefers them to other actions. Because of this, the agent is best-responding.

To see why this condition is necessary for a Nash equilibrium, suppose that player 1 had both  $H$  and  $T$  in the support of its strategy in *Matching Pennies*, but  $u_1(H, s_2) > u_1(T, s_2)$ . Player 1 is not best-responding by putting some probability on both  $H$  and  $T$  when the utility from  $H$  is greater than that from  $T$ . Mixing across actions is only a best-response if: (1) the player is indifferent across these actions, and (2) they're as good as all other actions.

**Example 2.7.** *Looking for a mixed-strategy Nash equilibrium in Matching Pennies in which both players mix over both actions, we need a probability  $p$  for player 1 such that player 2 is indifferent across  $H$  and  $T$ . This is  $p = 0.5$ . Similarly, we need to find a probability  $q$  for player 2 such that player 1 is indifferent across  $H$  and  $T$ . This is  $q = 0.5$ . We can conclude that  $(p^*, q^*) = (0.5, 0.5)$  is a mixed-strategy Nash equilibrium. Each player is indifferent across its two actions given the strategy of the other player, and thus both players are best-responding.*

We will make extensive use of this concept of the support of a strategy and this definition of a mixed-strategy Nash equilibrium in Chapter 5 when discussing algorithmic approaches to finding the equilibrium of games.

### 2.6.3 Multiplicity of Equilibria

Many games have multiple equilibria. Consider the following simple example:

**Example 2.8** (Game of Chicken). *In the game of Chicken, two drivers drive up to an intersection and each can either yield (i.e., stop) or continue going straight. If both yield (Y, Y) then both wait and their payoff is 0. If one goes straight and the other yields, (S, Y) or (Y, S), then one has payoff 2 and one has zero payoff. If both continue straight then there is a collision and both have payoff -4. See Figure 2.6. There are two pure-strategy Nash equilibria: (S, Y) or (Y, S). There is also a mixed-strategy Nash equilibrium, with strategies (2/3, 1/3) and (1/3, 2/3) for players 1 and 2. See Exercise 2.2.*

		Player 2	
		Y	S
Player 1	Y	0, 0	0, 2
	S	2, 0	-4, -4

Figure 2.6: Game of *Chicken*.

The existence of multiple equilibria can make it difficult to predict how players will act in a game. Certainly, when every player has an action that strictly dominates every other action (as in the *Prisoner's Dilemma*), then there is a unique Nash equilibrium. Similarly, when iterated elimination of strictly-dominated actions yields a single action profile (as in the example in Figure 2.3), then there is a unique Nash equilibrium. But many games have multiple equilibria.

Approaches to reconcile this difficulty include identifying an equilibrium that seems more likely because it Pareto dominates other equilibria, or is more stable to small, random mistakes by other agents. But these details are beyond the scope of this book.

### 2.6.4 Two-player zero-sum games

The *Matching Pennies* game is an example of a two-player *zero-sum* game.

**Definition 2.9** (Two-Player Zero-Sum Game). *A two-player simultaneous-move game  $(\{1, 2\}, (A_1, A_2), (u_1, u_2))$  is a zero-sum game if, for every action profile  $a = (a_1, a_2) \in A_1 \times A_2$ , the total utility is  $u_1(a) + u_2(a) = 0$ .*

In a zero-sum game, the outcomes most preferred by player 1 are the outcomes least preferred by player 2. In addition to modeling well-known games such as chess and poker, zero-sum games can be used to model resource allocation problems. For example, the problem facing two users who share the bandwidth of a wireless base station can be modeled as a zero-sum game. We return to zero sum games in Chapter 5, where we study the computational problem of finding a Nash equilibrium.

## 2.7 Congestion Games

In this section we discuss a special class of games that have a succinct representation, can be used to model many interesting domains, and for which a pure strategy Nash equilibrium

always exists.

A *congestion game* is a simultaneous-move game in which there are resources, each agent selects one or more resources, and the utility of an agent is the negated total congestion cost associated with the resources that it selects. Given resources  $E$ , the power set  $2^E$  is the set of all possible subsets; e.g., if  $E = \{1, 2\}$  then  $2^E = \{\emptyset, \{1\}, \{2\}, \{1, 2\}\}$ .

**Definition 2.10** (Congestion game). *A congestion game  $(N, A, E, c)$  has:*

- $N = \{1, \dots, n\}$  agents, indexed by  $i$
- $E = \{1, \dots, m\}$  resources, indexed by  $j$
- joint action set  $A = A_1 \times \dots \times A_n$ , where  $A_i$  is a set of actions available to agent  $i$  and  $A_i \subseteq 2^E$ , where  $2^E$  is the power set on the set of resources. Action  $a_i \in A_i$  selects a subset of resources.
- cost function  $c_e(x) \in \mathfrak{R}$  for resource  $e$  which depends on the number of agents  $x$  that select the resource

Let  $x_e$  be the total number of agents that select resource  $e$  given action profile  $a$ . The cost to agent  $i$ , given action profile  $a \in A$ , is

$$c_i(a) = \sum_{e \in a_i} c_e(x_e), \quad (2.7)$$

where the summation is taken over all resources  $a_i \subseteq E$  selected by agent  $i$ . The utility to agent  $i$  for action profile  $a$  is just the negated cost:  $u_i(a) = -c_i(a)$ .

In words, each player selects some subset of resources, this induces congestion on each resource, and the total cost experienced by a player is the sum over the cost on the resources he selects. It is often natural for the cost function  $c_e(x_e)$  to be non-decreasing and positive, but neither restriction is necessary.

In regard to succinctness, recall that the normal-form representation is exponential in the number of agents. In comparison, congestion games have a succinct representation because the cost (or negated utility) depends only on the number of players who select each resource and not the particular subset of players.

To see the modeling power of congestion games, let's consider two illustrative examples. The first is the network flow problem that illustrated Braess' Paradox in Chapter 1.

**Example 2.9** (Network flow). *See Figure 2.7. There are  $n = 2000$  agents, and resources  $E = \{12, 13, 23, 24, 34\}$  corresponding to the edges in the network. Each edge has a cost function, with  $c_{12}(x) = c_{34}(x) = x/100$ ,  $c_{13}(x) = c_{24}(x) = 25$ , and  $c_{23}(x) = 0$ . Each agent's available actions are:  $\{\{12, 24\}, \{12, 23, 34\}, \{13, 34\}\}$ , and correspond to the three possible paths. The cost function of agent  $i$  is  $c_i(a_i, a_{-i}) = \sum_{e \in a_i} c_e(x_e)$ , where  $e \in a_i$  enumerates the edges on its selected path and  $x_e$  is the number of agents that select edge  $e$  given action profile  $a = (a_i, a_{-i})$ . Each agent's utility is its negated cost. From Chapter 1, we know the unique pure-strategy Nash equilibrium (in fact a dominant strategy equilibrium) is for every agent to select action  $\{12, 23, 34\}$ , and take the path that includes zero-cost edge 2-3. This has cost  $c_i(a) = c_{12}(x_{12}) + c_{23}(x_{23}) + c_{34}(x_{34}) = 2000/100 + 0 + 2000/100 = 40$ . In comparison, the social optimal flow has 1000 agents taking path 1-2-4 and 1000 agents taking path 1-3-4, with cost 35 to each agent.*




Figure 2.7: Network flow problem: there are 2000 units to flow from location 1 to 4 over edges, each of which has an associated cost function.




Figure 2.8: Network connection problem:  $n$  agents need to connect location 1 with 2 and can select edge  $T$  or  $B$ .

Compared with the succinct congestion game formulation, a normal form representation for this network flow game would require enumerating the payoffs for each player for each of the  $3^{2000}$  possible action profiles.

**Example 2.10** (Network connection game). *See Figure 2.8. Consider a connection game, where each of  $n$  agents must choose to connect locations 1 and 2 by edge  $T$  or edge  $B$ . The agents that select  $T$  share cost  $n$  and the agents that select  $B$  share cost  $1 + \epsilon$  for some  $0 < \epsilon < 1$ . For example, the setting could be multiple firms each choosing a mode of transport that their employees will share to get across a city. The social optimal outcome is that everyone uses connection  $B$  with total cost  $1 + \epsilon$ .*

*Modeling this as a congestion game, the resources are  $\{T, B\}$  and the cost functions are  $c_T(x_T) = \frac{n}{x_T}$  and  $c_B(x_B) = \frac{1+\epsilon}{x_B}$ , where  $x_T$  and  $x_B$  are the number of agents who select  $T$  and  $B$  respectively. The action set is  $\{T, B\}$  for each agent. One Nash equilibrium is “all  $B$ ,” because an agent’s cost  $\frac{1+\epsilon}{n} < n$ , which is the cost to deviate. Another Nash equilibrium is “all  $T$ ,” because an agent’s cost  $\frac{n}{n} = 1 < 1 + \epsilon$ , which is the cost to deviate. One equilibrium*

is socially optimal and one is not.

These examples illustrate that a Nash equilibrium in a congestion game need not be socially optimal. In Chapter ?? we will study the inefficiency of both worst-case and best-case equilibria in these and other games.

In addition to providing succinct representations of many interesting settings, congestion games have the following property:

**Theorem 2.2.** *Every congestion game has a pure-strategy Nash equilibrium.*

This is an important property because pure-strategy Nash equilibria are more natural than mixed-strategy Nash equilibria, and when they exist they may better predict game play. To prove this result, we show in the next section that a congestion game is an example of a *potential game*, and that a pure-strategy Nash equilibrium always exists in a potential game.

## 2.8 Potential Games

In a potential game, a single value can be assigned to every action profile that captures what is essential about the strategic structure of the game.

**Definition 2.11** (Potential game). *A simultaneous-move game  $(N, A, u)$  is a potential game if there exists a function  $P : A \rightarrow \mathbb{R}$  from action profiles to reals, such that for every agent  $i$ , all actions  $a_{-i}$  chosen by the other players except  $i$ , and all actions  $a_i, a'_i$  of agent  $i$ , we have:*

$$u_i(a'_i, a_{-i}) - u_i(a_i, a_{-i}) = Pot(a'_i, a_{-i}) - Pot(a_i, a_{-i}) \quad (2.8)$$

In words, a game is a potential game if there is a function (the *potential function*) such that the difference in potential between any two action profiles, that differ only in the action of a single agent, is exactly the difference in utility to the agent whose action changes between the profiles. Potential functions are not unique: an arbitrary constant can always be added to the potential value of every action profile.

**Example 2.11.** *In Figure 2.9 (a) we provide a potential function for Prisoner's Dilemma. To check this, just verify that the difference in potential for all action profiles satisfy the potential property. For example, going from  $(C, C)$  to  $(D, C)$ , the action that changes is that of agent 1, and agent 1's utility increases by 2, which is exactly the change in potential function  $Pot(D, C) - Pot(C, C) = 2$ .*

**Example 2.12.** *Matching Pennies is not a potential game. In attempting to construct a potential function in Figure 2.9 (b), we begin with 0 in the top-left and work clockwise, defining the next potential value to correctly capture the change in utility. At the bottom-left we require a potential of 6, but then in moving from bottom-left to top-left, the difference in potential  $Pot(H, H) - Pot(T, H) = -6 \neq u_1(H, H) - u_1(T, H) = 2$ . For a potential to exist in this game, we need  $Pot(H, H) < Pot(H, T) < Pot(T, T) < Pot(T, H) < Pot(H, H)$ , which is impossible.*

2 Game Theory I: Simultaneous-Move Games

		Player 2		
		C	D	
(a) Player 1	C	-2, -2	-5, 0	$P = \begin{pmatrix} 0 & 2 \\ 2 & 3 \end{pmatrix}$
	D	0, -5	-4, -4	

		Player 2		
		H	T	
(b) Player 1	H	1, -1	-1, 1	$P = \begin{pmatrix} 0 & 2 \\ 6 & 4 \end{pmatrix}$
	T	-1, 1	1, -1	

Figure 2.9: (a) The *Prisoner's Dilemma* game and a potential function for the game. (b) The *Matching Pennies* game, and an attempt to construct a potential function (the bottom-left and top-left entries are incorrect for a deviation by player 1.)




Figure 2.10: An illustration of the potential function in a potential game and the action profile with the maximum potential.

See Figure 2.10 for an illustration of the potential function in a potential game, illustrated here for an arbitrary ordering on action profiles. There is no reason to expect the potential to vary smoothly as suggested in the figure. What is significant is the existence of an action profile  $a^*$  with maximum potential. From the *Prisoner's Dilemma* example, we see that  $(D, D)$  has maximum potential, and corresponds to the pure-strategy Nash equilibrium in the game.

This property holds generally in potential games:

**Theorem 2.3.** *Every potential game has a pure-strategy Nash equilibrium.*

*Proof.* Consider action profile  $a^* \in \arg \max_{a \in A} Pot(a)$ . By construction, for any other action profile  $(a'_i, a_{-i})$ , then  $Pot(a'_i, a_{-i}) \leq Pot(a)$  and therefore  $u_i(a'_i, a_{-i}) - u_i(a) = Pot(a'_i, a_{-i}) - Pot(a) \leq 0$ , and there can be no beneficial unilateral deviation. We conclude that action profile  $a^*$  is a pure-strategy Nash equilibrium.  $\square$

In particular, every congestion game is a potential game, and thus has a pure-strategy Nash equilibrium.

**Theorem 2.4.** *Every congestion game is a potential game.*



*Proof.* Given a congestion game, we construct a potential function, and show that the congestion game is a potential game. For this, consider the following potential function:

$$Pot(a) = - \sum_{e \in E} \sum_{j=1}^{x_e} c_e(j), \quad (2.9)$$

where  $x_e$  is the number of agents who select resource  $e$  given action profile  $a$ , and the second summation is zero when  $x_e = 0$ . Fix any actions  $a_{-i}$  chosen by all agents except  $i$ , and consider the change in potential from action  $a_i$  to action  $a'_i$ :

$$Pot(a'_i, a_{-i}) - Pot(a_i, a_{-i}) = \sum_{e \in E} \sum_{j=1}^{x_e} c_e(j) - \sum_{e \in E} \sum_{j=1}^{x'_e} c_e(j) \quad (2.10)$$

$$= \sum_{e \in E} \left( \sum_{j=1}^{x_e} c_e(j) - \sum_{j=1}^{x'_e} c_e(j) \right) = \sum_{e \in a_i \setminus a'_i} c_e(x_e) - \sum_{e \in a'_i \setminus a_i} c_e(x_e + 1) \quad (2.11)$$

$$= u_i(a'_i, a_{-i}) - u_i(a_i, a_{-i}), \quad (2.12)$$

where  $x_e$  and  $x'_e$  denote the count on resource  $e$  at action profile  $a$  and  $(a'_i, a_{-i})$  respectively. The third equality follows by recognizing that the sums for a resource  $e$  that is in both  $a_i$  and  $a'_i$  cancel (since  $x_e = x'_e$ ). For resources  $e \in a_i \setminus a'_i$  selected in  $a_i$  but not  $a'_i$ , there is an additional term in the first summation. For resources  $e \in a'_i \setminus a_i$  selected in  $a'_i$  but not  $a_i$ , there is an additional term in the second summation. The final equality holds because the increase in utility is the decrease in cost to agent  $i$ , which is exactly (2.11).  $\square$

A sequence of action profiles form a *path* if the sequence has the *single-deviation property*, such that only one agent changes its action at each step. A path  $a^{(0)}, a^{(1)}, a^{(2)}, \dots$ , is *improving* if  $u_i(a_i^{(k+1)}, a_{-i}^{(k)}) > u_i(a_i^{(k)}, a_{-i}^{(k)})$ , where agent  $i$ 's action changes in step  $k$ . Potential games have the finite-improvement property:

**Theorem 2.5** (Finite-improvement property). *Any improving path on action profiles in a potential game with a finite number of actions terminates in a finite number of steps with a pure-strategy Nash equilibrium.*

*Proof.* Consider an improving path on action profiles  $a^{(0)}, a^{(1)}, a^{(2)}, \dots$ . The potential  $Pot(a^{(k+1)}) > Pot(a^{(k)})$  for all steps  $k$ , and thus no action profile is repeated, and the path must terminate after a finite number of steps because there is a finite number of actions and thus a finite number of action profiles. Upon termination the action profile is a pure-strategy Nash equilibrium because no improvement is possible, and thus every agent is simultaneously maximizing its utility.  $\square$

An improving path need not reach the action profile with maximum potential. Rather, it can terminate at a local maxima in the potential landscape; i.e., an action profile where no deviation by a single agent can increase the potential.

The finite-improvement property suggests a natural *better-response* dynamic for finding a Nash equilibrium, in which players continually select improving actions given the actions of others. However, one caution is that this is not guaranteed to find a Nash equilibrium in a small number of steps.




Figure 2.11: A 2-by-2 cycle involving agents 1 and 2.

What is required of the payoff matrix for a game to be a potential game? For this, define the *value* of a path as the total change in utility, summed over the change incurred by the agent that changes its action at each step on the path. A *cycle* is a path that starts and ends at the same action profile.

**Example 2.13.** *The value of the cycle  $(C, C), (C, D), (D, D), (D, C), (C, C)$  in the Prisoner’s dilemma is,  $(u_2(C, D) - u_2(C, C)) + (u_1(D, D) - u_1(C, D)) + (u_2(D, C) - u_2(D, D)) + (u_1(C, C) - u_1(D, C)) = 0$ .*

Certainly, the value of all cycles in a potential game must be zero. We have seen this idea in Example 2.12. Say that a cycle is a *2-by-2 cycle* if it involves 2 agents, each of which changes its action twice. For example, Figure 2.11 illustrates a 2-by-2 cycle involving agents 1 and 2 and actions  $a_1, a_1'$  and  $a_2, a_2'$ . Exercise 2.5 establishes that it is sufficient that all 2-by-2 cycles have zero value for a game to be a potential game.

## 2.9 Notes

For a detailed introduction to game theory, a comprehensive reference is provided by “A Course in Game Theory” (Osborne and Rubinstein, MIT Press 2001). Gibbons “Game Theory for Applied Economists” (Princeton University Press 1992) provides a more accessible introduction. For an advanced reference, Fudenberg and Tirole’s “Game Theory” (MIT Press, 1991) is recommended. A large number of refinements have been proposed to the basic equilibrium concept, each of which imposes additional requirements on the outcome and seeks to identify a particular equilibrium prediction. We will see an example of such a refinement, in the context of games with sequential moves, in Chapter 3.

Chapters 1 and 17-20 in “Algorithmic Game Theory” (Nisan, Roughgarden, Tardos and Vazirani, eds, Cambridge University Press 2007) expands on some of the themes related to representational issues, as well as congestion games and potential games. “Essentials of Game Theory: A Concise, Multidisciplinary Introduction” (Leyton-Brown and Shoham, Morgan Claypool 2008) provides an accessible proof of the existence of a mixed-strategy Nash equilibrium in finite games, and develops utility theory within the von Neumann-Morgenstern axiomatic framework.

Congestion games were introduced by R. W. Rosenthal “A class of games possessing pure-strategy Nash equilibria” *Int. J. Game Theory* **2** (1973), 65-67. Later, D. Monderer and L.

S. Shapley “Potential games” *Games and Economic Behavior* **14**: 124-143, 1996 formalized the equilibrium properties of congestion games from the viewpoint of potential functions. In fact, every finite potential game is a congestion game. The development of potential games in Exercise 2.5 follows Monderer and Shapley. See T. Roughgarden “Computing Equilibria: A Computational Complexity Perspective” *Economic Theory* **42** 193-236 (2010) for a discussion of the complexity of finding equilibrium in congestion games. See N. Nisan, M. Schapira and A. Zohar “Asynchronous Best-Reply Dynamics” Proc. WINE 2008 for an example of a potential game in which best-response dynamics need not converge to a Nash equilibrium when players move at the same time and perhaps with delayed information about earlier moves.

Example 2.10, the network connection game, is introduced in E. Anshelevich, A. Dasgupta, J. Kleinberg, Tardos, T. Wexler, and T. Roughgarden, “The price of stability for network design with fair cost allocation”, Proceedings of IEEE Symposium on Foundations of Computer Science, 2004, pp. 295-304. Exercise 2.6 is developed from material in B. Awerbuch, Y. Azar, and A. Epstein “The price of routing unsplittable flow” in Proc. 37th ACM Sympos. on Theory of Computing, ACM Press, New York, 2005, pp. 57-66.

The existence of pure-strategy Nash equilibria in weighted congestion games (see Exercise 2.7) is due to D. Fotakis, S. Kontogiannis, and P. Spirakis “Selfish unsplittable flows,” Proc. 31st ICALP, LNCS 3142, Springer-Verlag, Berlin, 2004, pp. 593-605. The proof follows the same approach as that of Theorem 2.2, adopting a modified potential function.

The 3-player normal form game in Exercise 2.2 is from CS 224 (Stanford) Homework #1 (game theory). The scheduling game in Exercise 2.3 was introduced in Y. Azar, K. Jain and V. Mirrokni “(Almost) Optimal Coordination Mechanisms for Unrelated Machine Scheduling” Proc. Annual ACM-SIAM Symp. on Discrete Algorithms (2008). The agenda of designing coordination mechanisms (such as shortest-first precedence orders) for selfish scheduling was introduced by G. Christodoulou, E. Koutsoupias, and A. Nanavati “Coordination mechanisms”, Proc. 31st International Colloquium on Automata, Languages and Programming, pages 345-357 (2004).

The auction game in Exercise 2.4 (b) and (c) are based on A. Hassidim, H. Kaplan, M. Mansour, and N. Nisan. “Non-price equilibria in markets of discrete goods,” Proc. 12th ACM Conference on Electronic Commerce (EC), pages 295-296, 2011. The second-price auction game in Exercise 2.4 (d) is from K. Bhawalkar and T. Roughgarden, “Welfare Guarantees for Combinatorial Auctions with Item Bidding,” *Proc. SODA* (2011).

The load balancing game in Exercise 2.7 is from Chapter 20 “Selfish Load Balancing” by B. Vöcking in “Algorithmic Game Theory” (Nisan, Roughgarden, Tardos and Vazirani, eds, CUP 2007), which also provides an extensive discussion of this and related problems. The study of load balancing in Nash equilibrium was introduced in an influential paper by E. Koutsoupias and C. Papadimitriou, “Worst-case equilibria” in *Proc. 16th Sympos. on Theoretical Aspects of Computer Science*, 404-413 (1999). The load balancing game can be interpreted as a selfish routing game where the underlying network consists of two nodes, a source and a sink, and there are a set of parallel links from the source to the sink. Each machine corresponds to a link, and each task to a flow of a different size. The effect of selfish behavior on social welfare in the worst-case equilibria of games was later coined the *Price of Anarchy* by C. H. Papadimitriou in “Algorithms, games, and the Internet,” Proc. 33rd Annual ACM Symposium on Theory of Computing (STOC), pp. 749-753, 2001.

## 2.10 Comprehension Questions and Exercises

### 2.10.1 Comprehension Questions

- c2.1 What is the dilemma in the Prisoner's Dilemma and what does this illustrate more generally about Nash equilibria?
- c2.2 Why is there no pure-strategy Nash Equilibrium in the Matching Pennies game?
- c2.3 Why is it important to have succinct game representations?
- c2.4 Why must all actions in the support of a mixed strategy that is part of a Nash equilibrium have the same expected utility?
- c2.5 What do you see as two fundamental challenges in the application of game theory?

### 2.10.2 Exercises

#### 2.1 Iterated elimination of dominated actions

- (a) Prove that iterated elimination of strictly-dominated actions never removes an action that is part of any mixed-strategy Nash equilibrium, and that the set of equilibria in the reduced game is equal to that in the original game.
- (b) Give an example of a game where no action can be eliminated by iterated elimination of strictly-dominated actions.
- (c) What is the complexity of iterated elimination of strictly-dominated actions?
- (d) Consider a variation of iterated elimination that will remove an action  $a_i \in R_i$  if there is *weak* dominance, with some  $a'_i \in R_i$  such that  $u_i(a_i, a_{-i}) \leq u_i(a'_i, a_{-i})$  for all  $a_{-i} \in R_{-i}$  and  $u_i(a_i, a_{-i}) < u_i(a'_i, a_{-i})$  for at least one  $a_{-i} \in R_{-i}$ .
  - (i) Construct an example that shows that the order of elimination affects the set of eliminated actions under this notion of weak dominance.
  - (ii) Construct an example that shows that a Nash equilibrium can be eliminated.
  - (iii) Prove that any equilibrium of the game that results from iterated elimination of weakly dominated strategies will be an equilibrium of the original game.

#### 2.2 Pareto optimality

- (a) Prove that the mixed-strategy Nash equilibrium in which each player plays  $H$  with probability 0.5 is the only Nash equilibrium of *Matching Pennies*, and that every mixed-strategy profile is Pareto optimal in the game of *Matching Pennies*.
- (b) Prove that there always exists at least one Pareto Optimal action profile in a finite simultaneous-move game.
- (c) By plotting best-responses confirm the three Nash equilibria in the game of *Chicken* are  $(S, Y)$ ,  $(Y, S)$  and mixed with each player yielding with probability  $2/3$ . Is the distribution on action profiles in the mixed-strategy Nash equilibrium Pareto optimal? If not, provide a distribution on outcomes in *Chicken* that is Pareto optimal.

2.10 Comprehension Questions and Exercises

		Player 2	
		Stag	Hare
Player 1	Stag	400, 400	0, 100
	Hare	100, 0	100, 100

Figure 2.12: The game of *Stag*.

		Player 2		Player 2		
		L	R	L	R	
Player 1	T	(5, 5, 5)	(2, 6, 2)	T	(2, 2, 6)	(-1, 3, 3)
	B	(6, 2, 2)	(3, 3, -1)	B	(3, -1, 3)	(0, 0, 0)
		N		F		

Figure 2.13: A three-player normal form game. Player 3 plays  $N$  or  $F$ , and players 1 and 2 select  $T$  or  $B$  and  $L$  or  $R$  respectively.

- (d) In the game of *Stag* there are two hunters, and they can decide to hunt for a Stag or a Hare. The Stag is hard to catch and they both need to agree, while the Hare is less valuable. The payoff matrix is in Figure 2.12. Plot the best-responses and identify the Nash equilibria of the game.
- (e) Consider the 3-player normal form game in Figure 2.13. Each player has two actions:  $(T, B)$  for player 1,  $(L, R)$  for player 2 and  $(N, F)$  for player 3. Player 3 gets to select the left or right payoff matrix, player 2 the column and player 1 the row. For example, if they play  $(T, L, F)$  the payoffs are 2, 2, 6 to players 1, 2 and 3 respectively. List all of the pure strategy Nash equilibria and list all the Pareto optimal outcomes of the game.

2.3 Scheduling game

Consider a scheduling game with two machines and three agents. Each agent  $i$  has a task with cost  $c_{ij} > 0$  for machine  $j \in \{1, 2\}$ , representing the time the task takes to complete on the machine. Machine 2 is faster than machine 1 for jobs 1 and 2, but not for job 3 and the costs are  $c_{11} = 12, c_{12} = 10, c_{21} = 16, c_{22} = 10, c_{31} = 2, c_{32} = 16$ . Each agent selects a machine, with the tasks scheduled on the selected machine according to a *precedence order*. An agent's cost is the time when its own task completes, and it seeks to minimize this cost. Each machine adopts a *shortest-first* precedence order, preferring tasks that are shorter and breaking ties in favor of agents with a lower index.

- (a) What is the precedence order on tasks for each machine?
- (b) Give a pure-strategy Nash equilibrium for this game and argue why it is an equilibrium.
- (c) Explain without enumerating all possible action profiles why the Nash equilibrium is unique.

## 2 Game Theory I: Simultaneous-Move Games

- (d) The make-span is the time that the last task is completed. What is the make-span in the Nash equilibrium? What is the socially optimal assignment; i.e., the one that minimizes the make-span?
- (e) Provide a precedence ordering for machine 1 and 2 in the scheduling game for which there is no pure-strategy Nash equilibrium. Explain.

### 2.4 Auction games

- (a) Consider a first-price auction game with two agents and a single item to allocate. Agent 1's value is \$1 and agent 2's value is \$2. Agent 1 can bid  $x \in [0, 10]$  and agent 2 can bid  $y \in [0, 10]$ . The agent with the highest bid wins the item and pays its bid amount. Ties are broken in favor of agent 2. If allocated, an agent's utility is its value for the item minus its payment. Provide a pure-strategy Nash equilibrium in this game and prove that it is unique.
- (b) Now suppose there are two items  $A$  and  $B$ . Agent 1 needs both items and has a value \$4 for both together. Agent 2 has value \$3 for either  $A$  or  $B$ , and the same value for both  $A$  and  $B$  together (agent 2 only needs one item.) Agent 1 can bid  $x \geq 0$  for item  $A$  and the same amount  $x$  for item  $B$ . Agent 2 can bid  $y \geq 0$  for either item  $A$  or  $B$ , and must pick which. Each item is assigned to the agent with the highest bid, at a price equal to the bid on the item, and ties are broken in favor of agent 1. An agent's utility is its value minus its payment.

Prove that there is no pure strategy Nash equilibrium of this game.

- (c) A mixed-strategy Nash equilibrium of the game has agent 1 bid  $x$  in interval  $[0, 2]$  according to cumulative distribution function  $F(x) = \frac{1}{3-x}$  ( $\Pr(x \leq z) = F(z)$ ). Agent 2 bids on item  $A$  or  $B$  with equal probability, and bids  $y$  in interval  $[0, 2]$  according to cumulative distribution function  $G(y) = \frac{y}{4-y}$ .

Sketch or plot the distributions on bid  $x$  and bid  $y$ . Prove that agent 1 is best-responding to agent 2 and agent 2 is best-responding to agent 1, and thus this is a mixed-strategy Nash equilibrium.

- (d) Consider a second-price auction game. There are two items to allocate,  $A$  and  $B$  and two agents, 1 and 2. Let  $\gamma > 0$  denote a constant. Agent 1's value is 1 for  $A$ ,  $1 + \gamma$  for  $B$  and  $1 + \gamma$  for both items. Agent 2's value is  $1 + \gamma$  for  $A$ , 1 for  $B$ , and  $1 + \gamma$  for both items. The agent with the highest bid on an item wins the item and pays the bid amount of the *other* agent. Ties are broken in favor of agent 2. The social value of an assignment is the total value of the allocation. For example, assigning both items to agent 1 has social value  $1 + \gamma$ . Construct a pure-strategy Nash equilibrium for which the social value of the assignment in equilibrium is a vanishingly small fraction of the value of the socially-optimal assignment as  $\gamma$  increases.

### 2.5 Potential games

- (a) Show that the game of *Chicken* is a potential game, and construct a slight variation that illustrates that asymmetric games with 2 players and 2 actions can be potential games.

- (b) Construct a 2 player, 2 action game that has the finite-improvement property but is not a potential game. (Hint: it will need to have a pure-strategy Nash equilibrium in which an agent is indifferent between playing the equilibrium and deviating.)
- (c) Construct a 2 player game with a dominant-strategy equilibrium that is not a potential game.
- (d) Consider a potential game  $G$ , and a game  $G'$  in which every player's utility is  $u_i(a) = Pot(a)$ , and equal to the potential function in game  $G$ . Do games  $G$  and  $G'$  have the same set of Nash equilibria? Why or why not?
- (e) Recall that a path on action profiles satisfies the single-deviation property, and is a cycle if it starts and ends on the same action profile. A cycle may pass through the same action profile more than once. The value of a path is the total incremental change in utility along the path. Prove that if all cycles in a game have zero value then the game is a potential game. [Hint: fix any action profile  $z$ , and as a first step, establish by the zero-value-cycle property that any two paths from  $z$  to an action profile  $a \neq z$  have the same value. Second, show that  $Pot(a) = I(z \rightarrow a)$ , where  $I(z \rightarrow a)$  is the value of any path from  $z$  to  $a$ , satisfies potential property (2.8).]
- (f) Prove that if all 2-by-2 cycles (see Figure 2.11) have zero value then all cycles have zero value. [Hint: Assume for contradiction that there is a cycle  $\gamma = (a^{(0)}, a^{(1)}, \dots, a^{(\ell-1)}, a^{(\ell)})$ , where  $a^{(\ell)} = a^{(0)}$ , of length  $\ell \geq 5$ , with value  $I(\gamma) \neq 0$ , and that this positive-value cycle is minimal, in that all cycles of length  $< \ell$  have zero value. Assume WLOG that agent 1 moves in step 0, and let  $j$  denote another step in which 1 must move (this is required for it to be a cycle). First, argue by minimality (or the 2-by-2 assumption if  $\ell = 5$ ) that  $j$  is not step 1 or  $\ell - 1$ . WLOG, suppose agent 2 moves in step  $j - 1$ . Now consider cycle  $\gamma'$ , which differs from  $\gamma$  only in that agent 1 now deviates in step  $j - 1$  and agent 2 in step  $j$ ; i.e., steps  $a^{(j-1)}, a^{(j)}, a^{(j+1)}$  in  $\gamma$  become steps  $a^{(j-1)}, z^{(j)}, a^{(j+1)}$  in  $\gamma'$ , where  $z^{(j)}$  is obtained from  $a^{(j-1)}$  by agent 1's deviation. Second, argue by the zero-value-2-by-2 property that  $I(\gamma) = I(\gamma')$ . Third, by considering minimality, and recognizing  $I(\gamma') \neq 0$ , complete the proof.]

## 2.6 Network routing game

Consider a network routing game where each agent has to route a unit flow on a directed graph from one node to another. Each edge has a delay that depends on the total flow on the edge. Each agent wants to minimize cost, which is the total delay on the edges on its selected route. See Figure 2.14. There are four players, with start and end nodes as indicated. Each edge is annotated with its cost function, either  $c_e(x_e) = 0$  or  $c_e(x_e) = x_e$ .

- (a) Formulate this as a congestion game.
- (b) Identify two pure strategy Nash equilibria of the game. Argue why there are equilibria.
- (c) What is the socially optimal flow, i.e. the flow that minimizes the total cost to all agents?

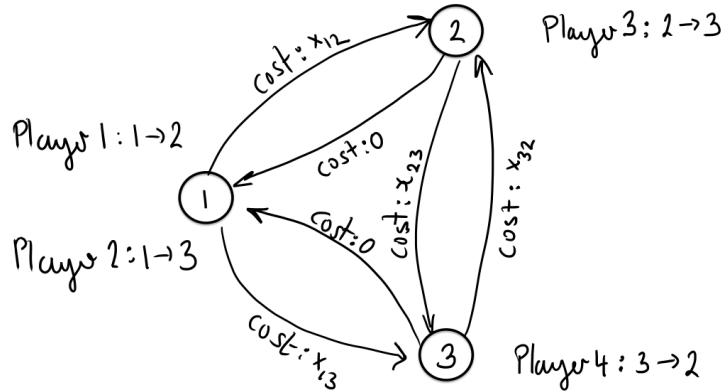


Figure 2.14: A network routing game, indicating the origin-destination pairs for each player and the delay functions on each edge.

- (d) Following the approach in Theorem 2.4, formulate the potential function and determine which of the two Nash equilibrium corresponds to the maximum potential in the game.

2.7 Load balancing game

In a *weighted* congestion game each agent has a weight  $w_i > 0$  and the congestion on resource  $e$  is  $x_e = \sum_{i: e \in a_i} w_i$  (the total weight of agents who selected the resource). A player's cost is  $c_i(a) = \sum_{e \in a_i} c_e(x_e)$ , as is standard.<sup>1</sup> If cost functions are linear, and  $c_e(x_e) = a_e x_e + b_e$  for  $a_e, b_e \geq 0$ , then every weighted congestion game has a pure-strategy Nash equilibrium.

Consider a load balancing game with two identical machines and four agents. Each agent  $i$  has a task with size  $w_i > 0$ , representing the size of the task. The sizes are  $w_1 = w_2 = 2$  and  $w_3 = w_4 = 1$ . Two large tasks and two small tasks. Each machine's speed is 1 unit per second and completes all assigned tasks at the same time. Each agent selects a machine, and incurs a cost equal to the time for the machine to complete.

- (a) Explain why the load balancing game is a weighted congestion game.
- (b) Verify that (i) agents 1 and 3 on machine 1 and agents 2 and 4 on machine 2 and (ii) agents 1 and 2 on machine 1 and agents 3 and 4 on machine 2 are pure-strategy Nash equilibria.
- (c) Explain why no assignment in which the maximum time for a machine to complete is 5 or larger can be a Nash equilibrium.

<sup>1</sup>It is also possible to define a player's cost as  $c_i(a) = w_i \sum_{e \in a_i} c_e(x_e)$  in a weighted congestion game. In fact, this does not change the pure-strategy or mixed-strategy Nash equilibria of the game.



## 2.10 Comprehension Questions and Exercises

- (d) What is the socially optimal assignment; i.e., the one that minimizes the maximum completion time across both machines (the “make-span”)?
- (e) What do you observe about the minimum and maximum ratio of make-span in a pure-strategy Nash equilibrium to the socially optimal make-span in this example?

### 2.8 Risk preferences

Consider a setting with two choices, 1 and 2. Choice 1 provides certainty of winning \$1 million, choice 2 provides a 50% chance of winning \$2 million. Most people will choose 1. Now consider a different choice where choice 1 provides a certainty of winning \$1, and choice 2 provides a 50% chance of winning \$2, then many people will choose 2.

Provide a utility function on money such that the choice that maximizes expected utility is 1 when the quantities are in millions and 2 when the quantities are in single dollars.