Chapter 1

Combinatorial Games

The best way to get a feel for combinatorial games is to play some! Try playing these games in which two players alternate making moves.

Game 1.1 (Pick-Up-Bricks)**.** This game is played with a pile of bricks. Each move consists of removing 1 or 2 bricks from the pile. The game ends when the pile is empty, and the last player to take a brick wins.

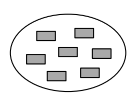

Figure 1.1. A 7-brick position in Pick-Up-Bricks

Game 1.2 (Chop)**.** Start with an $m \times n$ array viewed as a plank that is secured only at the lower left corner. On each turn, a player must either make a vertical or horizontal chop of the plank, and then any piece no longer connected to the lower left corner falls off into the water. The last player to make a move wins.

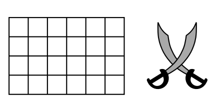

Figure 1.2. A 4×6 position in Chop

Game 1.3 (Chomp). Start with an $m \times n$ array viewed as a chocolate bar, but with the lower left corner square poisoned. On each turn, a player chooses a square and eats this square and all other squares that lie above and to the right of this one (i.e. the northeast corner). The last player to eat a nonpoison square wins.

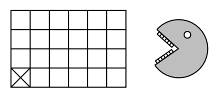

Figure 1.3. A 4×6 position in Chomp

Game 1.4 (Hex). This is a game played on the board pictured in Figure 1.4.[1] Similar to the widely familiar game Tic-Tac-Toe, on each turn a player marks an empty hexagon. One player uses $*$ as his mark, while the other uses $\circ$. If the player using $*$ can form a chain of hexagons with this mark connecting the left and right sides of the board, that player wins. The player using $\circ$ will win by forming a chain of hexagons with this mark connecting the top and bottom.

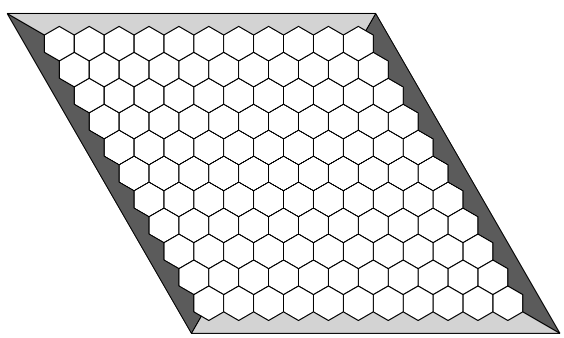

Figure 1.4. The board for Hex

Now that we have a handful of combinatorial games in mind, let's introduce some formal definitions that will allow us to treat these objects mathematically.

[1]A large version of this game board (as well as the pictured positions in Chop and Chomp) can be found at the end of this book and also online at `www.ams.org/bookpages/stml-80`.

Definition 1.5. A *combinatorial game* is a 2-player game played between **L**ouise (for **L**eft) and **R**ichard (for **R**ight).[2] The game consists of the following:

(1) A set of possible *positions*. These are the states of the game (e.g. a 2×3 board in Chop).

(2) A *move rule* indicating for each position what positions Louise can move to and what positions Richard can move to.

(3) A *win rule* indicating a set of *terminal positions* where the game ends. Each terminal position has an associated *outcome*, either Louise wins and Richard loses (denoted $+-$), Louise loses and Richard wins ($-+$), or it is a draw (00).

Observe that the definition of a combinatorial game does not indicate which player moves first, nor does it explicitly state which position is the starting position. To *play* one of these games, we choose a starting position and designate a player to move first. From then on, the players alternate making moves until a terminal position is reached and the game ends. In three of the games we have seen so far, Pick-Up-Bricks, Chop, and Chomp, the loser is the first player to have no available move. This is a common win rule, and we call a combinatorial game with this win rule a *normal-play* game.

Although numerous common games like Checkers and Chess are combinatorial games, there also exist many games that are not combinatorial games. Notably, combinatorial games have no element of randomness—so die rolls or spinners cannot be used to determine actions. Combinatorial games also require each player to have full information about the position of the game. Later in Chapter 7, we will broaden our horizons and introduce some of these variations. For now, though, we restrict ourselves to combinatorial games.

1.1. Game Trees

In this section, we will introduce a powerful tool called a game tree, which will be helpful for understanding the play of a game. We will

[2]These names are chosen in honor of Richard Guy, one of the founders of modern combinatorial game theory, and his wife, Louise.

begin by seeing how to model the play of any combinatorial game in this manner.

Modeling Play. There is a natural way to depict all possible sequences of moves in the play of a game using a tree, where each branch node models a choice point for one of the players and every terminal node indicates an outcome. This construct will prove extremely useful for us as we start to think strategically. We first introduce the simple game of Tic and consider its game tree.

Game 1.6 (Tic)**.** This is a combinatorial game similar to Tic-Tac-Toe but played on a 1×3 array. To move, Louise marks an empty square with a $\circ$ and Richard marks an empty square with a $\times$. If Richard or Louise gets two adjacent squares marked with his or her symbol, then he or she wins. If all squares get marked without this happening, the games ends in a draw.

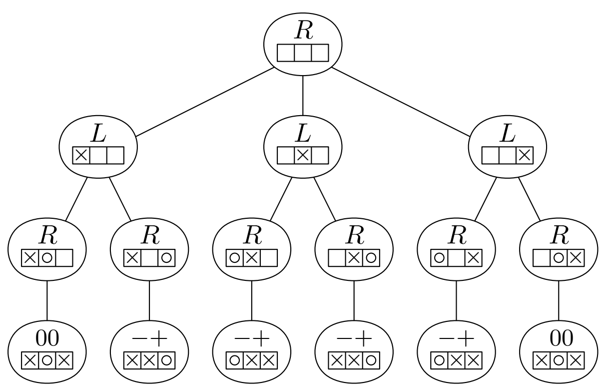

Figure 1.5. A game tree for Tic

Our game tree in Figure 1.5 depicts every possible sequence of moves starting from a blank board with Richard moving first. Observe that each node of the game tree contains the current position of the game, so we can see the positions updating as moves are made and we work downward in the tree. Each branch node also contains either an *L* or *R* to indicate that it is either Louise or Richard's turn to play. The terminal nodes indicate that the game has ended with an outcome, either a win for Richard $-+$, a win for Louise $+-$, or a draw 00. The topmost node containing the starting position is called the *root* node.

More generally, we can model any combinatorial game with this process. We will call these figures *game trees* and they will prove quite helpful for our analysis.

Procedure 1.7 (Build a Game Tree)**.** To make a game tree starting at position α with Louise moving first, we begin by making a root node containing an L (since Louise is first) and an α (since this is the starting position). If Louise can move to positions $\alpha_1, \ldots, \alpha_k$, then we join k new nodes to the root node; each one of the new nodes contains one of these α_i positions. If any of these positions is terminal, then we put the appropriate outcome (either $+-$, $-+$, or 00) in this node. For the other nonterminal positions, it will be Richard's turn to play, so all of these nodes will contain an R. Continue this process until it is complete.

W-L-D Game Trees. We introduced the game tree in Figure 1.5 as a way to model the play of Tic. However, once we have this tree in hand, we could actually play it instead of the game. Rather than starting with an empty 1×3 array and having Louise and Richard alternately mark boxes with their symbols, we could start at the root node of this game tree and descend it by having Louise choose at the nodes marked L and having Richard choose at the nodes marked R. As you can see, these two different ways of playing are essentially equivalent.

Once we are operating in this game tree model, the position information that is contained in each node is superfluous. Indeed, all that our players need to play this game tree is the information concerning which player has a choice to make at each branch node and what the outcome is at each terminal node. Ignoring the position information for the game Tic gives us the tree depicted in Figure 1.6.

This type of game tree without the position information is extremely useful, so we shall give it a name. We define a **W-L-D** *game tree* (for **W**in-**L**ose-**D**raw) to be a tree with a distinguished root node (as the starting position) in which each terminal node contains an outcome ($+-$, $-+$, or 00) and each branch node contains either an L or an R, indicating which

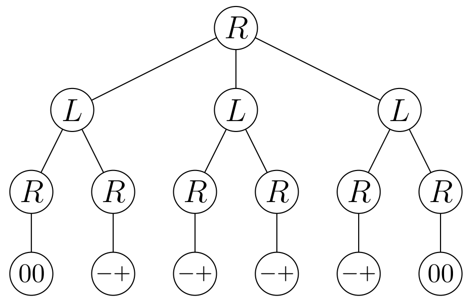

Figure 1.6. A W-L-D game tree for Tic

player moves by making a choice at that node. The nice property of W-L-D game trees is that they give us a unified way to think about the play of combinatorial games. So, instead of having to consider players taking tokens and marking squares and eating chocolate, we can always view play in terms of descending a W-L-D game tree.

It is possible to have a W-L-D game tree consisting of just a single terminal node (so neither player has any decisions to make and the outcome is already decided). This wouldn't be much fun to play in practice, but it will be convenient for our theory. There are three such trees with no moves (see Figure 1.7) and we will call them *trivial* game trees.

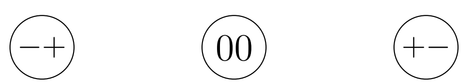

Figure 1.7. The trivial W-L-D game trees

There is another extreme to consider... it is possible that our game tree could go on forever and never end! Although the definition does not force combinatorial games to end, we will restrict our attention only to games that must end after a finite number of moves. Accordingly, we will always assume that W-L-D game trees are finite.

Strategy. When playing a game, we like to win! To do so, we will want a plan. The idea of strategy formalizes this notion of a plan. The term strategy is familiar to game players everywhere and can be used to indicate a general principle of play—in chess, for example, one may wish to "control the center"—but we will adopt a more refined and specific usage of this term. We define a *strategy* for a player in a W-L-D game tree to be a set of decisions indicating which move to make at each node where that player has a choice. In the game tree in Figure 1.8, we have depicted a strategy for Richard by boldfacing the edges indicating his choices.

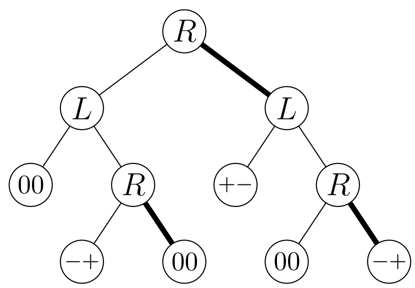

Figure 1.8. A strategy for Richard

In the play of our game, a player may *follow* a strategy by using it to make every decision. Note that the strategy indicated for Richard in Figure 1.8 is not a terribly good one since it gives Louise the opportunity to move to a terminal node with outcome +− where Richard will lose. Richard would do better to follow a strategy that makes the opposite choice at the root node since then it will be impossible for him to lose.

We will generally be interested in finding good strategies for our players. The best we could hope for is a strategy that guarantees a win for a player who follows it. We will call any such strategy a *winning* strategy. Next best would be a strategy that guarantees a player doesn't lose. A strategy with the property that following it guarantees either a win or a draw is called a *drawing* strategy. Note that a player following a drawing strategy could end up winning; the only guarantee is that this player will not lose.

Let us note that there may be some extraneous information included in a strategy as we have defined it. For example, in the strategy for Richard depicted in Figure 1.8, Richard will choose the right branch as his first move. As a result, he will never encounter the node in the lower left part of the tree labeled *R*. Since he will never encounter this node, it may seem unnecessary for him to decide what to do there. Nevertheless, our definition of strategy includes the decision Richard would make at every node labeled R. The simplicity of this definition makes strategies easier to work with.

Working Backwards. Let's now assume that our players are highly rational with perfect foresight and consider how they might play in a large W-L-D game tree. From the nodes at the top of the tree, it is not clear what choices either player would prefer since there are so many decisions still ahead. In contrast, for the nodes close to the bottom it is much easier to see how best to play. Consider the choice for Richard in the game tree depicted in Figure 1.9. It is clear that from here Richard will choose the right branch for a win.

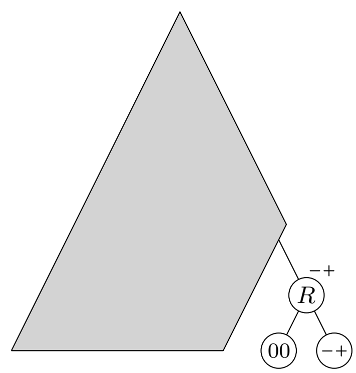

Figure 1.9. An easy decision for Richard

Since we now know that reaching Richard's decision node in Figure 1.9 will result in a win for Richard, we have written a "$-+$" next to this node. In fact, we can now think of this as a new terminal node and then apply a similar process elsewhere in the game tree. Next we formalize this approach.

Procedure 1.8 (Working Backwards)**.** Suppose one of our players has a decision to make at node N. Suppose also that we have already determined what the outcome will be under rational play for all possible nodes from node N. Then choose a best possible outcome for this player, indicate this choice by darkening this edge, and then mark the node N

with the resulting outcome. Continue this process until the root node has been marked with an outcome.

Figure 1.10 shows the result of carrying this procedure to completion on a larger game tree. Note in this figure that the root node has been labelled $-+$. This means that under rational play, Richard will win this game.

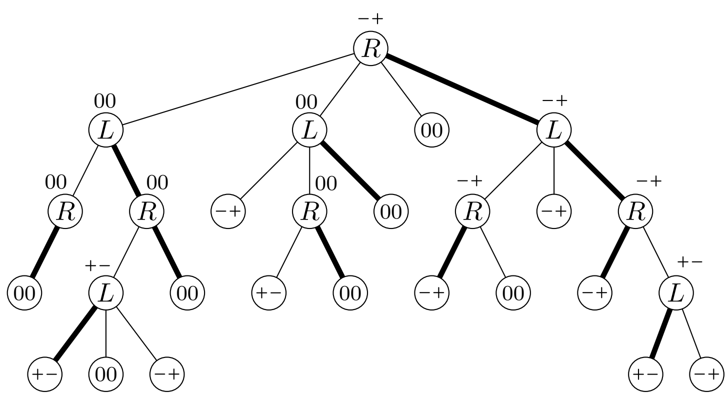

Figure 1.10. Rational decision-making

In fact, as you may readily verify, the strategy for Richard indicated in Figure 1.10 is a winning strategy. In the following section, we will show that this procedure works more generally for any W-L-D game tree. More precisely, when we apply our working backwards procedure to any W-L-D game tree, we either get a $+-$ on the root node and a winning strategy for Louise or a $-+$ on the root node and a winning strategy for Richard or a 00 on the root node and a drawing strategy for each player.

1.2. Zermelo's Theorem

In this section, we will prove a famous theorem due to Ernst Zermelo. This theorem tells us that in every combinatorial game, either Louise has a winning strategy or Richard has a winning strategy or both players have a strategy to guarantee them a draw. In the process of proving this theorem, we will develop the tools to prove that the working backwards procedure described above really does work as claimed on game trees of all sizes. The proof of Zermelo's Theorem is based on the mathematical principle of induction, so we begin with a brief discussion of this important concept.

Mathematical Induction. Induction is an extremely powerful tool for proving theorems. The simplest proofs by induction are used to prove properties of the nonnegative integers. For instance, suppose that $P(n)$ is a certain property of the number n that we wish to prove holds true for every $n \geq 0$. Since there are infinitely many nonnegative integers, it would be impossible to make a new proof for each individual one! Mathematical induction instead provides a general method to prove that $P(n)$ holds true for every $n \geq 0$.

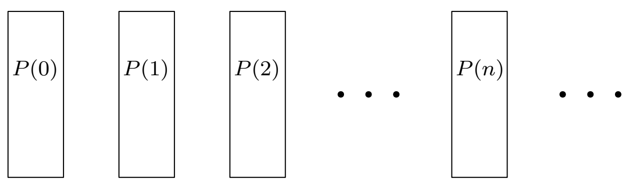

Figure 1.11. A proof by induction

The inductive approach is to view $P(0)$, $P(1)$, $P(2)$, ... as dominoes. We will think of knocking a domino over as showing that property P holds true for integer n. The proof involves two stages. We first show that the first domino falls over—that is, we prove that $P(0)$ is true. This part is called the *base case*. The second part is to prove that for every $n \geq 1$, if all the dominos before the n^{th} domino fall, then the n^{th} domino also falls. That is, we must prove that if $P(k)$ is true for all $k < n$ (this is the *inductive hypothesis*), then $P(n)$ is also true.[3] This part is called the *inductive step*. Of course, if this happens, every domino will be knocked over.

Base Case: $P(0)$ is true.

Inductive Step: If $P(k)$ is true for all $k < n$, then $P(n)$ is true.

[3]Technically speaking, we are introducing "strong induction" here since our inductive assumption is that $P(k)$ holds for all $k < n$. In "weak induction" this is replaced by the weaker assumption that $P(n-1)$ holds. These two principles are logically equivalent, but this text frequently utilizes the strong form, so that is what we will adopt throughout.

Let's consider a straightforward example that exhibits a nice property of positive integers.

Example 1.9. For every $n \geq 0$, the sum of the first n odd integers is n^2. For the proof, we proceed by induction on n. To verify the base case, observe that $0^2 = 0$ is the sum of the first 0 odd integers. For the inductive step, let $n \geq 1$ be an arbitrary integer, and assume that our formula holds true for all nonnegative integers less than n. In particular, the formula holds for $n - 1$, which means

$$1 + 3 + 5 + \cdots + (2(n-1) - 1) = (n-1)^2.$$

Starting with this equation and adding $2n - 1$ to both sides gives us

$$1 + 3 + 5 + \cdots + (2n-3) + (2n-1) = (n-1)^2 + (2n-1) = n^2$$

and this completes the proof.

Notice in this example that property $P(n)$ is never formally defined in the proof. Nevertheless, the idea is there: $P(n)$ is the property that the sum of the first n odd numbers is n^2. This proof handles the base case by verifying $P(0)$, or showing that the sum of the first 0 integers is 0^2 (as usual, the base case was the easy part). For the inductive step, we assumed $P(k)$ to be true for all $k < n$ (our inductive hypothesis) and then used this to prove that $P(n)$ holds true for every $n \geq 1$. The inductive hypothesis gave us the advantage of starting with the equation for $P(n-1)$, from which we then deduced $P(n)$.

Mathematical induction enjoys extremely wide application, well beyond proving nice properties of positive integers. Indeed, we will later see more involved instances of induction at work in our investigations of combinatorial games. The general context in which induction applies involves some property P, which we want to show holds true for infinitely many things. To proceed by induction, we will want to organize these things into different *sizes* (i.e. some things have size 0, some have size 1, etc.). Now, instead of trying to prove all at once that all of these things satisfy property P, we can proceed by induction on the size. The base case will be to prove that P is true for all things of size 0. Then, for the inductive step, we will assume (the inductive hypothesis) that P is true for all things of size less than n and use this to show that P then holds true for an arbitrary thing of size n.

In some cases where induction applies, the smallest relevant size might not be 0 and could instead be 1 or 2 or something else. Most generally, the base case involves proving the result for the smallest size that makes sense in context. The inductive step handles all things of larger size.

Proof of Zermelo's Theorem. With the concept of induction in hand, we are now ready to give a proof of Zermelo's famous theorem. Our proof will rely upon an inductive argument that applies to W-L-D game trees, so we will need to decide upon a "size" for these trees. We will use a quantity called depth as the size of a game tree. The *depth* of a game tree is the maximum number of possible moves from the start to the end of the game (e.g. the tree in Figure 1.10 has depth 4). Since every game tree is finite, every game tree has some depth and that depth will always be a nonnegative integer. This sets us up to use induction on depth to prove common properties of all game trees.

To prove a property P of trees by induction on depth, we first prove the base case, that P is true for all trees of depth 0. For the inductive step, we need to prove that P holds for an arbitrary tree of depth $n > 0$ under the inductive assumption that P holds true for all trees with depth less than n. Next we'll see this idea in action in the proof of Zermelo's Theorem. This theorem introduces a new definition, the *type* of a W-L-D game tree, and it establishes that every possible W-L-D game tree has one of three types.

Theorem 1.10 (Zermelo)**.** *Every W-L-D game tree is one of*

Type	*Description*
$+-$	*Louise has a winning strategy.*
$-+$	*Richard has a winning strategy.*
00	*Both players have drawing strategies.*

Proof. We proceed by induction on the depth of the game tree. As a base case, observe that if the tree has depth 0 (i.e. it is trivial), then the game is already decided and it is either $+-$ in which case Louise has a winning strategy, $-+$ in which case Richard has a winning strategy, or 00 in which case both players have drawing strategies.

For the inductive step, let T be a W-L-D game tree of depth $n > 0$ and assume the theorem holds for every tree with smaller depth. Suppose that Richard has the first move in T (the case where Louise has the first move follows from a similar argument) and he can move to one of the nodes $N_1, N_2, \dots, N_\ell$.

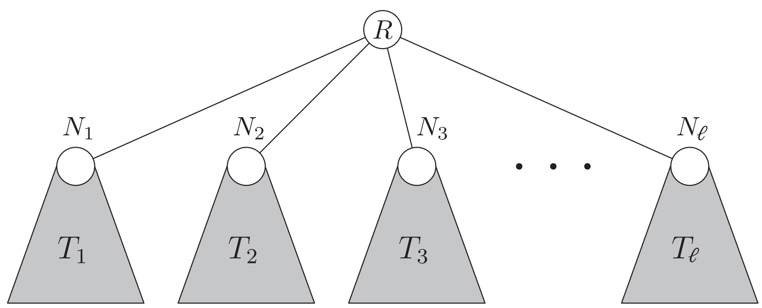

Figure 1.12. The first move in a game tree

We can consider each node N_i as the root of a new W-L-D game tree T_i, consisting of N_i and all the nodes below it. Since each T_i has depth $< n$, our inductive hypothesis tells us that all of these games must satisfy the theorem (i.e. must be type $+-$, $-+$, or 00). So, for every $1 \leq i \leq \ell$ we may choose strategies $\mathcal{L}_i$ for Louise and $\mathcal{R}_i$ for Richard in the game tree T_i with the property that either one of these strategies is winning or both are drawing. We form a strategy $\mathcal{L}$ for Louise in the original game tree by combining $\mathcal{L}_1, \dots, \mathcal{L}_\ell$. To form a strategy $\mathcal{R}$ for Richard, we will combine $\mathcal{R}_1, \dots, \mathcal{R}_\ell$ but we will also need to make a decision at the root node. Next we split into cases.

Case 1. *At least one of $T_1, \dots, T_\ell$ is type $-+$.*

Let T_i be type $-+$. Then Richard's strategy $\mathcal{R}_i$ is winning in T_i and we may form a winning strategy $\mathcal{R}$ in the original game by having Richard play to node N_i.

Case 2. *All of $T_1, \dots, T_\ell$ are type $+-$.*

In this case, every $\mathcal{L}_i$ strategy is winning, so $\mathcal{L}$ is a winning strategy for Louise.

Case 3. *None of $T_1, \dots, T_\ell$ is type $-+$, but at least one is type* 00.

Let T_i be type 00. Then Richard's strategy $\mathcal{R}_i$ is drawing and we may form a drawing strategy $\mathcal{R}$ in the original game by having Richard play to N_i. Since none of $T_1, \dots, T_\ell$ is type $-+$, each of Louise's strategies $\mathcal{L}_1, \dots, \mathcal{L}_\ell$ is drawing or winning, and it follows that $\mathcal{L}$ is a drawing strategy for Louise. □

So, Zermelo's Theorem gives us a classification of game trees into types $+-$, $-+$, and 00. Furthermore, the proof of this result implies that our "Working Backwards" technique from the previous section will always have one of the following results.

Corollary 1.11. *For every W-L-D game tree applying the Working Backwards procedure results in one of the following:*

root label	*result*
$+-$	*A winning strategy for Louise*
$-+$	*A winning strategy for Richard*
00	*Drawing strategies for both players*

For the purposes of analyzing small games, constructing a game tree is a convenient way to determine whether one of the players has a winning strategy or if they both have drawing strategies. For large games like Chess, it is theoretically possible to construct a game tree[4]. However, the number of positions in Chess has been estimated at approximately 10^{120} while the number of atoms in the universe is around 10^{80}... so our universe isn't big enough for such analysis! The fact that such a game tree does exist for Chess nevertheless means that Zermelo's Theorem applies to it. So, in Chess, either one of the two players has a winning strategy or both have drawing strategies ... we just don't know which of these it is.

1.3. Strategy

Game trees can be a useful tool to determine who will win when playing from a particular position in a game such as Chop. However, Chop boards come in infinitely many sizes, so there are infinitely many different game trees, and this approach will be limited at best! In this section

[4]Although it is possible to repeat positions in a game of chess, there are certain lesser-known rules that make it a finite game.

we will introduce a pair of useful ideas—namely symmetry and strategy stealing—that can help us to determine which player has a winning strategy without having to analyze the game tree. In particular, we will learn exactly which player has a winning strategy in every Chop position. Although these techniques do not apply to all games, they are powerful when they work.

Before we introduce these new ideas, let us pause to discuss how we represent or define a strategy. Suppose that we are interested in playing a certain position in a combinatorial game (with someone going first). To describe a strategy for a player, it is possible to construct the associated game tree and then depict the strategy there. However, this is a bit tedious, so it will be helpful for us to adopt a more relaxed treatment. Accordingly, we will generally describe a strategy for a player (in words) by giving a rule that tells this player what to do at each possible position. Given a strategy described in this manner, we could form the game tree and depict it there, but there is generally no need to do this.

Symmetry. The key to finding winning strategies in certain games is symmetry of the positions. Indeed, this simple idea is the key to understanding both Chop and Pick-Up-Bricks. By definition, every position in the game of Chop is a rectangle of the form $m \times n$. Some of these rectangles have the additional symmetry of being square (so $m = n$) and these positions are the key to understanding this game.

Proposition 1.12. *Consider an $m \times n$ position in Chop.*

(1) If $n = m$, the second player has a winning strategy.

(2) If $n \neq m$, the first player has a winning strategy.

Proof. First we prove that the second player has a winning strategy whenever the initial position is square. This winning strategy for the second player is easy to describe: On each turn, move the position to a square one. Assuming the second player does this, every time the first player has a move to make (including the first) the position will be a square and any move is to a nonsquare position. Note that the second player can always move a nonsquare position to a square one. Assuming this is done, the second player will eventually move to a 1×1 position and win the game.

A similar idea reveals a first-player winning strategy when the initial position is not square. On the first turn, the first player may move the board to a square position. From there, that player may adopt the above strategy (always moving to a square position). This will guarantee the first player a win. □

In Pick-Up-Bricks, the positions where the number of bricks is a multiple of 3 are the symmetric positions, and they play the same role as the square positions in Chop.

Proposition 1.13. *Consider a Pick-Up-Bricks position of n bricks.*

(1) If 3 *divides n, the second player has a winning strategy.*

(2) Otherwise, the first player has a winning strategy.

Proof. For the first part, the following strategy is winning for the second player: On each turn, do the opposite of the first player's move. So, if the first player picks up one brick, then the second player picks up two, and if the first player picks up two, then the second player picks up one. This ensures that after both players have played, the total number of bricks is three fewer and the new position is again a multiple of 3. Following this, the second player will eventually take the last brick and win.

The first player can win when the starting position is not a multiple of 3. To start, the first player may remove either one or two bricks to bring the position to a multiple of 3. Now the first player may adopt the second player strategy described above and win from here. □

Strategy Stealing. Strategy stealing is another tool to approach the general question of which player has a winning strategy. Much how symmetry sometimes works well to understand strategy for games, strategy stealing also is very effective when it applies. However, unlike the symmetry arguments that explicitly construct winning strategies, strategy-stealing arguments prove the existence of a winning strategy without giving any indication of what this strategy is! This is because strategy-stealing arguments employ proofs by contradiction.

To create a proof by contradiction, we begin by assuming the opposite of what we are trying to prove. We then argue deductively that a necessary consequence of that assumption is something impossible.

It follows that our assumption generating this contradiction must have been false, and therefore, the opposite is true. The strategy-stealing arguments that we will introduce here rely on somewhat subtle proofs by contradiction, which will benefit from careful contemplation. Before taking them on, we begin with a warm-up proof by contradiction in the following example.

Example 1.14. Let n be an integer and assume that n^2 is even. Then n must also be even. We will prove this (admittedly easy) fact by contradiction. So, our first step will be to assume that n is odd (i.e. assume the negation of what we are trying to prove). Since we are now assuming that n is odd, we may express it as $n = 2k + 1$ for another integer k. This gives us

$$n^2 = (2k + 1)^2 = 4k^2 + 4k + 1 = 2(2k^2 + 2k) + 1.$$

We conclude from the above equation that n^2 is odd, but this contradicts the hypothesis that n^2 is even. So, our initial assumption that n is odd has lead us to a contradiction, and thus we can conclude that n must be even, as desired.

We next use proof by contradiction to determine which player has a winning strategy in Chomp and in Hex. These proofs are called strategy-stealing arguments because we will assume (for a contradiction) that one of the players has a winning strategy, and then the other player will try to steal this strategy to win. Note again that both of these arguments will only prove the existence of a winning strategy. The proofs give no information about the specific moves that could be used to win the game.

Proposition 1.15. *For every rectangular position in Chomp except* 1×1*, the first player has a winning strategy.*

Proof. Chomp cannot end in a draw, so it follows from Zermelo's Theorem that one of the two players has a winning strategy. To prove that it is the first player who has a winning strategy, we will employ a proof by contradiction. For our proof, we will assume the first player does *not* have a winning strategy. It then follows from Zermelo's Theorem that the second player does have a winning strategy, which we will call $\mathcal{S}$. Now, consider what would happen if the first player removed just the upper rightmost square and then the second player chose her move according to $\mathcal{S}$. The resulting position must be some version of one of the shapes in Figure 1.13.

Figure 1.13. Simple positions in Chomp

This position must be one from which the second player to move has a winning strategy (since the second player followed the winning strategy $\mathcal{S}$ to get here). However, this very same position could also be reached in one step! The first player could have moved the board to this position on his or her first move. Then the first player would be second to play from this position, and this means that the first player has a winning strategy. This contradicts our assumption that the first player did *not* have a winning strategy. We conclude that our assumption is false, which means this game does indeed have a winning strategy for the first player. □

In Chapter 9, Exercise (11), we will show that the game of Hex cannot end in a draw. Our next theorem uses this property together with a strategy-stealing argument to show that the first player has a winning strategy.

Proposition 1.16. *The first player has a winning strategy in Hex (starting from an empty board).*

Proof. Assuming Hex cannot end in a draw (this fact is proved in Exercise (11) of Chapter 9), Zermelo's Theorem tells us that one of the two players must have a winning strategy. Let's assume (for the sake of contradiction) that the first player does not have a winning strategy. In this case we may choose a winning strategy $\mathcal{S}$ for the second player.

Now we will take control of the first player and we'll take advantage of the strategy $\mathcal{S}$ to win (this will give us a contradiction, thus proving $\mathcal{S}$ cannot exist). On our first move, we choose an arbitrary hexagon h and mark it with our symbol $*$. However, we will pretend that the hexagon h

has no mark and that we are the second player. This allows us to adopt the strategy $\mathcal{S}$ to make our moves. By following this winning strategy, we are guaranteed to win in the pretend version of our game. Since the true game just has one extra hexagon marked with $*$, this also gives us a win in the true game.

There is one slight complication we ignored in the above argument, but it isn't hard to fix. Namely, it might be the case that at some point the strategy $\mathcal{S}$ we are following in the pretend game instructs us to mark the hexagon h with $*$. Since h is already occupied by a $*$, this is not a possible move for us in the true game. In this case, we just choose another unoccupied hexagon h', mark it with $*$ and now pretend that h' is empty. This permits us to keep following the strategy $\mathcal{S}$ in our pretend game. We may later end up with other pretend-empty hexagons h'', h''', and so on. But in any case, it follows from the assumption that $\mathcal{S}$ is a winning strategy that we win our pretend game. Since any win in the pretend game guarantees us a win in the true game, we have constructed a winning strategy for the first player. This contradiction shows that the second player does not have a winning strategy, so the first player does. □

Exercises

(1) For each game below, construct a game tree with Louise moving first from the indicated starting position.
  (a) Tic starting from a blank board.
  (b) Pick-Up-Bricks starting with 4 bricks.
  (c) Chop starting from a 2×3 board.

(2) In each W-L-D game tree, use Procedure 1.8 to find a winning strategy for one of the players or a drawing strategy for both.

(a)

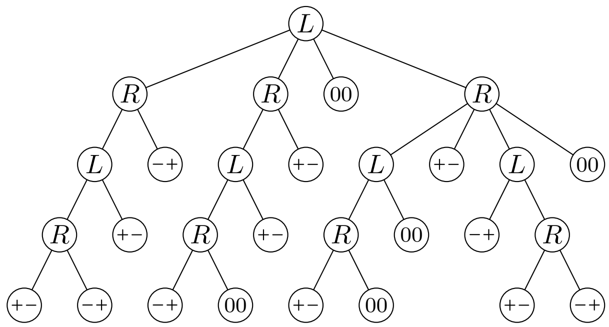

(b)

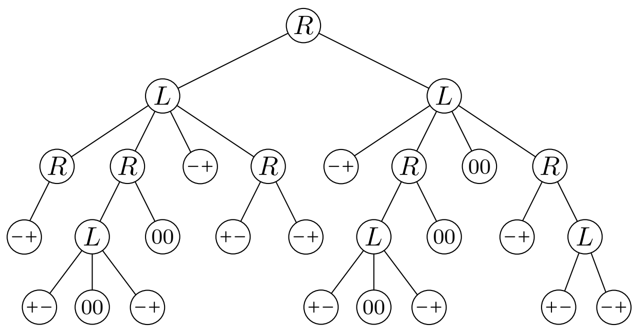

(3) For each position below, construct a game tree with Louise moving first. Then use Procedure 1.8 to find a winning strategy for one of the players.
  (a) A 5-brick position in Pick-Up-Bricks.
  (b) The Chomp position 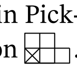.

(4) Consider the labeled W-L-D game tree below. To describe a strategy for Louise, we can indicate the label corresponding to each of her choices. For instance, $AB'C$ and $A''BC'$ are strategies for Louise. Similarly, XY' and $X''Y$ are strategies for Richard.
  (a) Find all of Louise's strategies.
  (b) Find all of Richard's strategies.
  (c) Find a pair of strategies, one for Louise and one for Richard, that will produce outcome $+-$. Repeat for the outcomes $-+$ and 00.

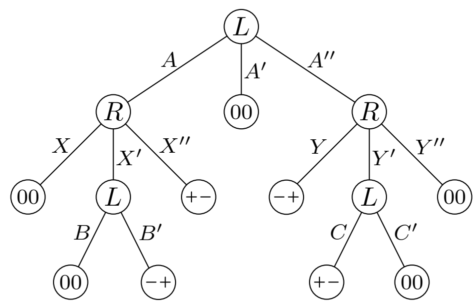

(5) As in the previous problem, we may describe strategies in the game tree below using labels so $AB'C$ is a strategy for Louise and XY'' is a strategy for Richard.
  (a) Find all of Louise's strategies.
  (b) Find all of Richard's strategies.
  (c) Which of Louise's strategies are drawing?
  (d) Which of Richard's strategies are drawing?

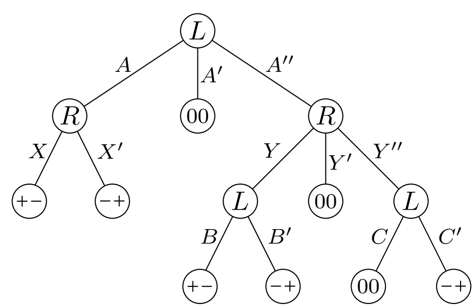

(6) Determine if each of the following games is a combinatorial game. If a given game is not a combinatorial game, explain why not.
  (a) Connect-Four.
  (b) Battleship.
  (c) Backgammon.
  (d) Go.

(7) For every pair m, n of positive integers, determine the depth of the game tree for
  (a) an m-brick position in Pick-Up-Bricks,

(b) an $m \times n$ position in Chop, and
(c) an $m \times n$ position in Chomp.

(8) Prove the following formulas by induction on n:
(a) $1 + 2 + 3 + \cdots + n = \frac{n(n+1)}{2}$.
(b) $1^2 + 2^2 + 3^2 + \cdots + n^2 = \frac{n(n+1)(2n+1)}{6}$.

(9) For every positive integer n let α_n be the two-row Chomp position where the top row has a single square and the bottom row has n squares (e.g. α_5 is 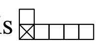). Let β_n be the position in Chop given by a $2 \times n$ array (e.g. β_5 is 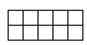). Prove that the game trees for the positions α_n and β_n have the same number of nodes for every $n \geq 1$.

(10) For any pair m, n of positive integers, define $\alpha_{m,n}$ to be the "L" shaped position in Chomp consisting of a column of m squares with the poison square at the bottom and a row of n squares with the poison square on the left (e.g. $\alpha_{3,6}$ is 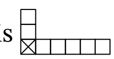). For all possible values of m and n find a winning strategy for either the first or second player.

(11) Let T be a W-L-D game tree with the property that every time a player has a choice to make there are exactly two options. If t is the number of terminal nodes and b is the number of branch nodes (i.e. nonterminal nodes), prove that $t = b + 1$.

(12) Consider a game tree T where Louise has n_2 nodes where she has a choice between 2 options, n_3 nodes where she has a choice between 3 options, ..., n_k nodes where she has a choice between k options (and no nodes where she chooses among more than k options). What is the formula for the total number of strategies that Louise has in T?

(13) Let T be a W-L-D game tree and assume that every time Louise has a decision to make in T she has exactly two options, and also assume that none of Richard's choices can bring the game to a terminal position. Let ℓ be the number of nodes marked L and let n be the total number of nodes. Prove that $n = 3\ell$ if Louise makes a choice at the root node, and otherwise $n = 3\ell + 1$.

(14) If N and N' are two nodes of a game tree and a player with a choice at N can move to N', then N and N' are connected by an edge (in our drawings, this edge is realized by a line segment). Prove every game tree has exactly one more node than edge.

(15) Let a and b be positive integers and let T be a W-L-D game tree with the property that whenever Louise has a choice to make she has exactly a options and whenever Richard has a choice to make he has exactly b options. Suppose that n is the total number of nodes, ℓ is the number of nodes marked L, and r is the number of nodes marked R. Prove that $a\ell + br + 1 = n$.

(16) For every integer $n \geq 1$ let c_n be the number of nodes in the game tree for a $1 \times n$ position in Chop. Find (and prove) a formula for c_n for every $n \geq 1$.

(17) The Fibonacci Sequence is an infinite sequence that starts off 0, 1, 1, 2, 3, 5, 8, 13, 21, ... It is defined by the following rules: $f_0 = 0$, $f_1 = 1$, and $f_{n+2} = f_{n+1} + f_n$ for every $n \geq 0$.

(a) Let φ, ψ be the solutions to the equation $x^2 = x + 1$ where $\varphi > 0$ (here φ is the Golden Ratio). Prove by induction that $f_n = \frac{\varphi^n - \psi^n}{\sqrt{5}}$ for every $n \geq 0$.

(b) For every real number x, let $[x]$ denote the closest integer to x. Prove that $f_n = \left[\frac{\varphi^n}{\sqrt{5}}\right]$ holds for every $n \geq 0$.

(c) Let p_n be the number of nodes in a game tree for a Pick-Up-Bricks position with n bricks. Find (and prove) a formula for p_n.

(18) Consider a Chomp position that consists of just two rows, with the bottom row of length n (including the poison square) and the other row of length m and assume $m \leq n$. Determine which player has a winning strategy for all possible values of m and n and prove your answer. (Hint: Consider the special case $n = m + 1$.)

(19) The game Kayles is played with a $1 \times n$ array, each square of which is either empty or is marked by a P indicating that it contains a bowling pin. On each turn we think of a player as bowling a ball toward this line-up of bowling pins. Each player has perfect aim and may choose to knock down any one pin or any two consecutive pins (i.e. erase the P from either one box or from two adjacent boxes). The first player to clear the board wins. For every positive integer n find a winning strategy for either the first or second player when the starting position is a $1 \times n$ array where every square contains a P.

(20) Consider playing d-dimensional Tic-Tac-Toe on a board with dimensions $\underbrace{n \times n \times \cdots \times n}_{d}$. To win, a player needs to get all n boxes in some line marked with his or her symbol ($\times$ or $\circ$). Prove that for every $n, d \geq 1$ the second player does not have a winning strategy.

(21*) The game SOS is played on a $1 \times n$ array. On his turn a player may choose any empty square and mark either an "S" or an "O" in it. If a player manages to get the letters "SOS" appearing in three consecutive squares and he is the first one to achieve this, then he wins. If the game ends without this arrangement, it is a draw. Prove that from each of the following starting positions, the first player has a winning strategy.

(a) A 1×4 array where the leftmost square is marked "S" and the other three squares are empty.

(b) A blank 1×7 array.

(c) A blank $1 \times n$ array where $n \geq 7$ is odd.

Chapter 9

Nash's Equilibrium Theorem

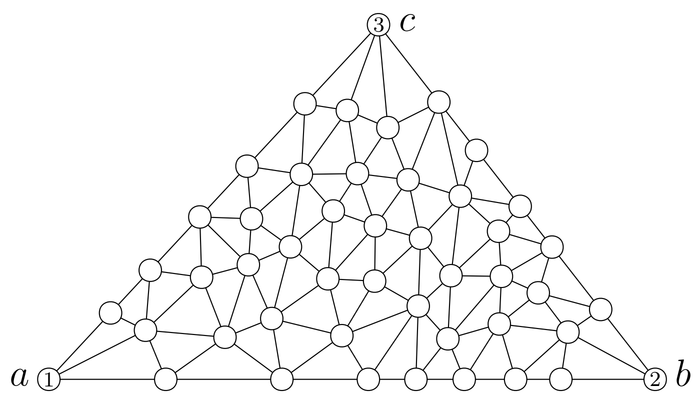

Figure 9.1. Triangle Solitaire

This chapter presents a proof of Nash's Equilibrium Theorem. The intricate proof of this deep result even involves the game of Triangle Solitaire shown in Figure 9.1! For accessibility, the argument here focuses on the special case of 2×2 games. Each result cleanly generalizes to higher dimensions and all those details appear in Appendix C.

We now know how to calculate a Nash equilibrium in a 2×2 matrix game, and it's easy to hope that higher-dimensional cases could be handled similarly, perhaps with more complicated algebraic expressions. This is simply not the case. Analogously, the familiar quadratic formula finds roots of a second-degree polynomial and there are similar formulae for finding roots of polynomials of degrees three and four. However, there is no such formula for a fifth-degree polynomial ... the situation is too complex to admit simple algebraic solutions. It is nevertheless still possible to reason and prove things about the roots of these polynomials.[1] Likewise, there are explicit formulae for finding Nash equilibria in small cases, but such conveniences do not exist for higher dimensions. Hence, we take a less direct approach to proving the existence of a Nash equilibrium. Figure 9.2 outlines the major pieces of the argument.

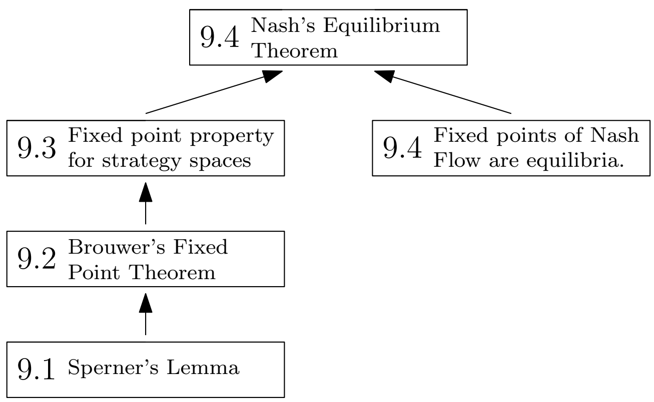

Figure 9.2. Steps in proving Nash's Theorem

We begin with the following one-player game.

Game 9.1 (Triangle Solitaire)**.** This game uses the triangle in Figure 9.1.[2] To play, write a 1, 2, or 3 in each of the empty circles according to the rule that every circle on the side between the corners labeled i and j must get either an i or a j. So, for instance, every circle on the bottom of the big triangle must be filled with either a 1 or a 2, but those

[1]For instance, there is an easy proof that a polynomial of degree 5 must have at least one root. This follows from the fact that a polynomial $f(x)$ of degree 5 either has $\lim_{x\to\infty} f(x) = \infty$ and $\lim_{x\to-\infty} f(x) = -\infty$ or has $\lim_{x\to\infty} f(x) = -\infty$ and $\lim_{x\to-\infty} f(x) = \infty$. In either case, we can deduce that f has both positive and negative values, so by the Intermediate Value Theorem it must have a root.

[2]A larger version of this game board can be found at the end of the book and also online at `www.ams.org/bookpages/stml-80`.

in the interior can be labeled 1, 2, or 3. The goal in this game is to minimize the number of little triangles with three differently labeled corners. How well can you do?

9.1. Sperner's Lemma

This section builds to a proof of a beautiful general result due to Emanuel Sperner. Despite its significance, this result is traditionally called a lemma. We begin with a one-dimensional version that involves a subdivided line segment. The two-dimensional version calls on the triangle in Figure 9.1.

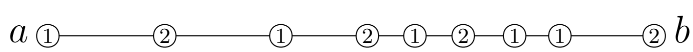

Figure 9.3. Sperner's Lemma in 1D

Lemma 9.2 (Sperner 1D)**.** *Let $\overline{ab}$ be a line segment that is subdivided into edges by adding some new vertices. Assume that each vertex is labeled according to the following rules:*

- *a is labeled 1 and b is labeled 2.*
- *Every other vertex is labeled 1 or 2.*

Then there exist an odd number of edges with endpoints of both numbers.

Proof. Imagine starting at a and walking along the line segment to b. Starting at a vertex labeled 1 and walking to one labeled 2 means switching numbers an odd number of times. Thus, there are an odd number of edges with ends of different labels. □

This one-dimensional result figures into the following proof that the game of Triangle Solitaire will always have at least one small triangle with all three labels.

Lemma 9.3 (Sperner 2D)**.** *Let $\triangle abc$ be subdivided into small triangles by adding new vertices and edges. Assume that each vertex is given a label according to the following rules:*

- *The labels on a, b, c are (respectively) 1, 2, 3.*
- *Every vertex on $\overline{ab}$ is labeled 1 or 2.*

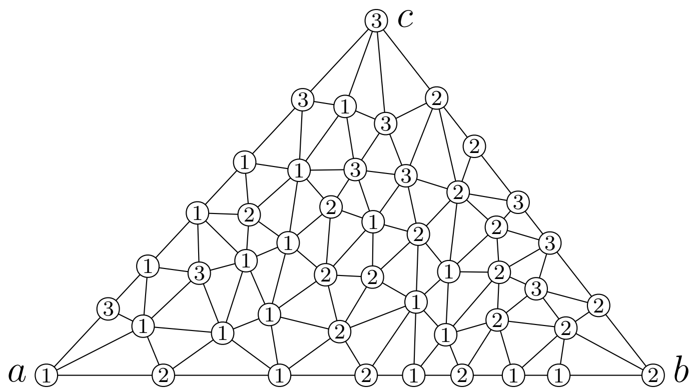

Figure 9.4. Sperner in 2D

- *Every vertex on $\overline{bc}$ is labeled 2 or 3.*
- *Every vertex on $\overline{ac}$ is labeled 1 or 3.*
- *Every vertex inside $\triangle abc$ is labeled 1, 2, or 3.*

Then there are an odd number of small triangles with vertices of all three labels.

Proof. Imagine Figure 9.4 as a floor plan for a house, so each little triangle is a room with three walls. Then add a door along every wall (edge) that has one end labeled 1 and the other labeled 2. Now consider the possibilities. A room with at least one door must have one vertex of label 1 and one of label 2. If the third vertex has label 3, then the room has just one door. Otherwise, the third vertex is labeled either 1 or 2, and, in either case, the room will have exactly two doors. No room has more than two doors, so a person walking through the house must walk only forward or backwards along a pathway. Figure 9.5 highlights the pathways in this labeling.

Some rooms have no doors and some pathways form cycles—these are irrelevant. Focus on the pathways that have a beginning and an end. The first and last door on any such pathway must either be a door to the outside (along $\overline{ab}$) or a door into a room with just one door. Since each such pathway has two ends, the total number of doors to the outside plus the number of rooms with exactly one door is even. The number of

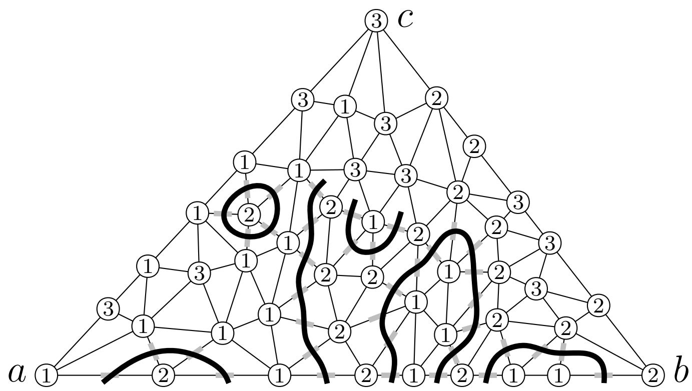

Figure 9.5. Pathways

doors to the outside must be odd by Sperner 1D. Therefore, the number of rooms with exactly one door must also be odd. In particular, there is at least one room with just one door. This is a small triangle with vertices of all three labels! □

Higher Dimensions. To extend Sperner's Lemma to higher dimensions requires sets in $\mathbb{R}^n$ that behave like line segments in $\mathbb{R}^1$ and triangles in $\mathbb{R}^2$. The notions of hyperplane and convex hull that we saw with von Neumann's Theorem in Chapter 6 recur here. One way to describe a triangle is as the convex hull of 3 points in $\mathbb{R}^2$ that do not lie on a common line. This idea generalizes to the following definition of an n-dimensional simplex.

Definition 9.4. An *n-simplex* is the convex hull of $n + 1$ points in $\mathbb{R}^n$ that do not lie on a common hyperplane.

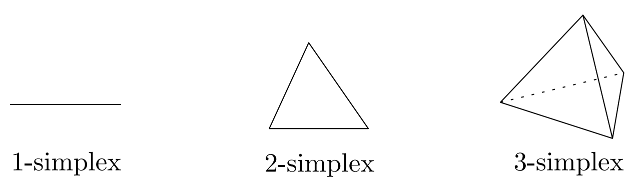

Figure 9.6. Small-dimensional simplexes

Observe that the 1D Sperner Lemma involved a 1-simplex that was subdivided and had vertices labeled 1 and 2. The 2D Sperner Lemma involved a 2-simplex that was subdivided with vertices labeled 1, 2, and 3. More generally, the n-dimensional Sperner Lemma uses an n-simplex that has vertices labeled $1, 2, \ldots, n+1$.

Lemma 9.5 (Sperner)**.** *Consider an n-simplex given as the convex hull of* $\mathbf{x}_1, \ldots, \mathbf{x}_{n+1}$ *that has been subdivided into small simplexes. Suppose that each vertex is given a label from* $1, 2, \ldots, n+1$ *satisfying the following rule:*

- *If a vertex has label i, then it does not lie in the convex hull of the points* $\mathbf{x}_1, \ldots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \ldots, \mathbf{x}_{n+1}$.

Then the number of small simplexes with all $n+1$ *labels is odd.*

The proof of this n-dimensional result is a generalization of the proof for the 2D version and appears explicitly in Appendix C.

9.2. Brouwer's Fixed Point Theorem

Some subsets of $\mathbb{R}^n$ exhibit a special property called the fixed point property. Perhaps surprisingly, Sperner's Lemma features prominently in the proof of a topological result called Brouwer's Fixed Point Theorem concerning this characteristic. This result is essential in proving Nash's Equilibrium Theorem.

Fixed Points. This section focuses on subsets $X \subseteq \mathbb{R}^n$ and functions of the form $f : X \to X$. Define a *fixed point* of such a function f to be a point $x \in X$ for which $f(x) = x$ (i.e. it is a point that is fixed by the function).

Examples 9.6.

(1) The function $f_1 : [0,1] \to [0,1]$ given by $f_1(x) = 1 - x$ has $1/2$ as a fixed point since $f_1(1/2) = 1/2$.

(2) The function $f_2 : \mathbb{R}^2 \to \mathbb{R}^2$ given by $f_2(x, y) = (y+1, x-1)$ has $(1, 0)$ as a fixed point since $f_2(1, 0) = (1, 0)$.

(3) The function $f_3 : \mathbb{R} \to \mathbb{R}$ given by $f_3(x) = e^x$ has no fixed point since there is no solution to the equation $x = e^x$.

Moving forward, it will be useful to know if *all* continuous functions on a particular set have a fixed point.

Definition 9.7. We say that a set $X \subseteq \mathbb{R}^n$ has the *fixed point property* if every continuous function $f : X \to X$ has a fixed point.

Note that this definition only concerns continuous functions. To prove that a set X does *not* have the fixed point property, simply find one continuous function from X to itself that has no fixed points.

Examples 9.8.

(1) Consider a circle $C \subseteq \mathbb{R}^2$. Now, choose an angle $0 < \theta < 2\pi$ and define a function $f : C \to C$ by the rule that f rotates each point around the circle by an angle of θ. Since this is a continuous function with no fixed point, we conclude that C does not have the fixed point property.

(2) Consider the set consisting of the entire real number line $\mathbb{R}$. The function $f : \mathbb{R} \to \mathbb{R}$ given by $f(x) = x + 1$ is a continuous function with no fixed point. Therefore, $\mathbb{R}$ does not have the fixed point property.

It is considerably more difficult to prove that a set X does have the fixed point property since this means showing that *every* continuous function $f : X \to X$ has a fixed point. In the case when X is a closed interval, we have good tools to solve this problem.

Theorem 9.9 (Brouwer 1D)**.** *The closed interval* $[0, 1]$ *has the fixed point property.*

Proof. We must show that every continuous function from the interval $[0, 1]$ to itself has a fixed point. Let $f : [0, 1] \to [0, 1]$ be such a function and define the new function $g : [0, 1] \to \mathbb{R}$ by the rule $g(x) = x - f(x)$. Now g is also a continuous function and $g(0) = -f(0) \leq 0$ while $g(1) = 1 - f(1) \geq 0$. It follows from the Intermediate Value Theorem that there exists a point c in $[0, 1]$ so that $g(c) = 0$. This point c satisfies $0 = g(c) = c - f(c)$, so it is a fixed point of the function f. Since this holds for every continuous function $f : [0, 1] \to [0, 1]$, we have established that $[0, 1]$ has the fixed point property, as desired. □

Next, we will prove that a certain triangle also has the fixed point property.

Triangles. Define Δ_2 to be the solid planar triangle with vertices $(0, 0)$, $(0, 1)$, and $(1, 0)$. More formally,

$$\Delta_2 = \{(x, y) \in \mathbb{R}^2 \mid x, y \geq 0 \text{ and } x + y \leq 1\}.$$

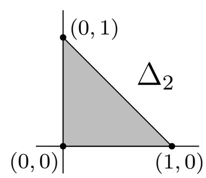

Figure 9.7. The triangle Δ_2

To prove that Δ_2 has the fixed point property, we need to consider continuous functions $f : \Delta_2 \to \Delta_2$. While it's common to visualize functions $f : \mathbb{R} \to \mathbb{R}$ using a graph, this technique is not so helpful for picturing a function $f : \Delta_2 \to \Delta_2$. Instead, draw Δ_2 and draw a collection of arrows to indicate the output of particular points under f (an arrow indicates that the initial point (x, y) is mapped to the terminal point $f(x, y)$). Since only continuous functions concern us, these arrows give a good idea of how f acts on nearby points, too.

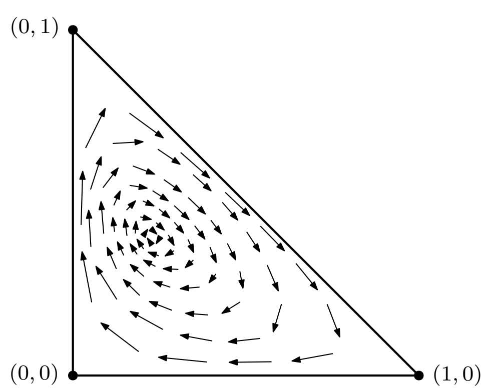

Figure 9.8. A continuous function $f : \Delta_2 \to \Delta_2$

Let (x, y) be a point in the triangle Δ_2 and suppose that $f(x, y) = (x', y')$. Define $(x', y') - (x, y) = (x' - x, y' - y)$ to be the *direction of* (x, y). Observe that in Figure 9.8, each arrow indicates the direction of its initial point.

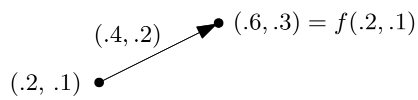

Figure 9.9. The direction of $(.2, .1)$ is $(.4, .2)$.

Now for an unusual move. Instead of considering all possible directions for a point, divide the directions into three groups: West, Southeast, and Northeast. This divides the points in the triangle into three sets that give a very rough indication of where the points go when we apply the function f. Formally, assign each point in Δ_2 a label 1 (for West), 2 (for Southeast), and 3 (for Northeast) according to the following rule: If $(x', y') = f(x, y)$, then

$$(x, y) \text{ has label} \begin{cases} 1 & \text{if } x' < x, \\ 2 & \text{if } x' \geq x \text{ and } y' < y, \\ 3 & \text{if } x' \geq x \text{ and } y' \geq y. \end{cases}$$

Figure 9.10 helps visualize these directions. In words, when the direction associated with a point has an angle of θ, the label will be a 1 if $\frac{\pi}{2} < \theta \leq \frac{3\pi}{2}$, a 2 if $\frac{3\pi}{2} < \theta < 2\pi$, and a 3 if $0 \leq \theta \leq \frac{\pi}{2}$.

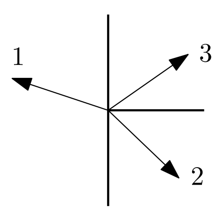

Figure 9.10. Assigning labels to directions

Consider a point (x, y) that is not a fixed point, and suppose it has a direction with angle $\frac{\pi}{4}$. By the above labeling scheme, this point will have label 3. Because f is continuous, every point near (x, y) will have a similar direction. So all points sufficiently close to (x, y) also have label 3.

Next suppose that (x, y) has direction with angle $\frac{\pi}{2}$ and is thus assigned label 3. This point would be on the boundary between points labeled 1 and 3, so all points sufficiently close would be labeled 1 or 3. Similarly if (x, y) has direction with angle $-\frac{\pi}{2}$, all points sufficiently close would be labeled 1 or 2, and if the direction has angle 0, all points sufficently close would be labeled 2 or 3. The only way for (x, y) to have points of all three labels arbitrarily close to it is for (x, y) to be a fixed point.

The following proof utilizes exactly this feature of the labeling. For any continuous function f, we ignore everything except the associated labeling of the points. We show that there exists a point (x, y) with points of all three labels arbitrarily close to it, and from this we deduce that f has (x, y) as a fixed point.

Theorem 9.10 (Brouwer 2D)**.** *The triangle Δ_2 has the fixed point property.*

Proof. Let $f : \Delta_2 \to \Delta_2$ be continuous, and label the points of Δ_2 in accordance with the above description.

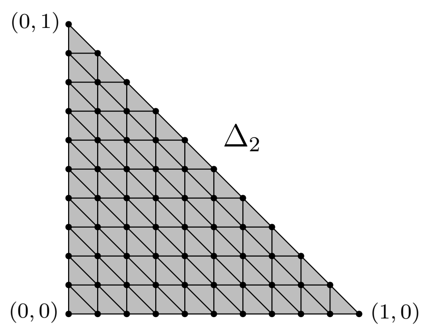

Figure 9.11. A subdivision of Δ_2

Claim. *For every $\ell \geq 0$ either there is a fixed point or a point in Δ_2 with distance $\leq \frac{1}{2^\ell}$ to points of all three labels.*

Subdivide the triangle Δ_2 into small triangles using a fine mesh (as in Figure 9.11) so that the center of each small triangle is a distance at most $\frac{1}{2^\ell}$ from each of its three vertices. Now consider all of the vertices of this subdivision. If one of them is a fixed point, then there is nothing left to prove, so assume no vertex is a fixed point. Consider the 1, 2, 3 labeling of these vertices and observe that this labeling satisfies the assumptions of Sperner 2D (with $a = (1,0)$, $b = (0,1)$, and $c = (0,0)$). The point $(0,0)$, for instance, must get a label of 3, the point $(1,0)$ must get a label of 1, and the points on the line segment between these two all have the form $(x,0)$, so they will be labeled 1 or 3. It follows that there is a small triangle with vertices of all three labels. By construction, the center of that triangle is a distance $\le \frac{1}{2^\ell}$ from points of all three labels.

By applying the above claim for $\ell = 1,2,3,\ldots$ we either find a fixed point (thus completing the proof) or we generate a sequence of points in the triangle $(x_1,y_1),(x_2,y_2),\ldots$ so that (x_i,y_i) is a distance $< 1/2^i$ from points of all three labels. It follows[3] that there is a particular point (x^*,y^*) in the triangle that has points of all three labels arbitrarily close to it, so (x^*,y^*) is a fixed point. □

Higher Dimensions. To generalize this to higher dimensions we introduce some special n-simplexes defined as

$$\Delta_n = \{(x_1,\ldots,x_n) \in \mathbb{R}^n \mid x_1,\ldots,x_n \ge 0 \text{ and } x_1 + \cdots + x_n \le 1\}.$$

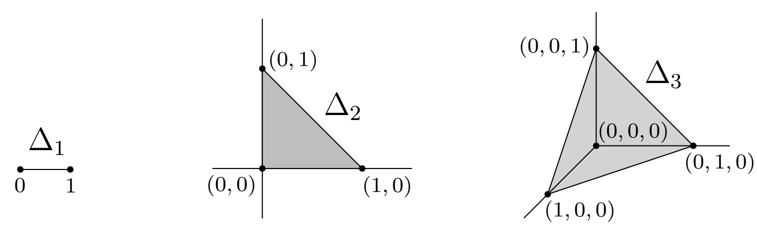

Figure 9.12. More simplexes

So Δ_1 is the line segment $[0,1]$ and Theorem 9.9 showed that it has the fixed point property. Theorem 9.10 proved that Δ_2 also has the fixed point property. The Brouwer Fixed Point Theorem asserts that this holds true in general.

Theorem 9.11 (Brouwer). *The simplex Δ_n has the fixed point property for every $n \ge 1$.*

[3]Since Δ_2 is compact, this sequence has a convergent subsequence.

The proof of this more general result is a straightforward extension of the proof of Theorem 9.10 that calls upon the general Sperner Lemma. See Appendix C for details.

9.3. Strategy Spaces

So far in this chapter, we have no hint whatsoever of a matrix game! This is about to change: We will use certain subsets of $\mathbb{R}^n$ to describe all possible pairs of mixed strategies for Rose and Colin and then apply the Brouwer Fixed Point Theorem to prove that these subsets have the fixed point property. This constitutes the main step in the anticipated proof of Nash's Equilibrium Theorem.

2×2 matrices. Suppose that Rose and Colin are playing a 2×2 matrix game A. Recall that a mixed strategy for Rose is a vector of the form $\mathbf{p} = \begin{bmatrix} p & 1-p \end{bmatrix}$ where $0 \leq p \leq 1$. Although there are two entries in this vector, the single real number p entirely determines $\mathbf{p}$. Likewise, a mixed strategy for Colin is a vector $\mathbf{q} = \begin{bmatrix} q \\ 1-q \end{bmatrix}$, but just the single number q is enough to describe this strategy completely. So we may identify Rose's mixed strategies with numbers $0 \leq p \leq 1$ and Colin's mixed strategies with numbers $0 \leq q \leq 1$.

How can we simultaneously describe a mixed strategy for Rose and one for Colin? In the case of a 2×2 matrix game as above, a pair of real numbers (p, q) with $0 \leq p \leq 1$ and $0 \leq q \leq 1$ is enough. This set of ordered pairs makes up what is formally called a strategy space.

Definition 9.12. The *strategy space* of a 2×2 matrix game is

$$S_{2,2} = \{(p, q) \in \mathbb{R}^2 \mid 0 \leq p \leq 1 \text{ and } 0 \leq q \leq 1\}.$$

The strategy space $S_{2,2}$ is the familiar set of points that make up the unit square. Associating each point (p, q) in $S_{2,2}$ with the pair of mixed strategies $\begin{bmatrix} p & 1-p \end{bmatrix}$ for Rose and $\begin{bmatrix} q \\ 1-q \end{bmatrix}$ for Colin equips us with a nice geometric interpretation of all possible pairs of strategies.

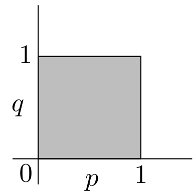

Figure 9.13. The strategy space of a 2×2 game

Strategy spaces will reappear with Nash flow in the following section. For now, recognize that the strategy space $S_{2,2}$ is a square that encodes pairs of mixed strategies for the players. The desired result here is that the strategy space of every matrix game has the fixed point property. We will need another concept from the world of topology to achieve this.

Topological Equivalence. Thanks to the Brouwer Fixed Point Theorem, we know that the triangle Δ_2 has the fixed point property. To deduce from this the fact that square $S_{2,2}$ also has the fixed point property will involve showing that Δ_2 and $S_{2,2}$ have a certain kind of equivalence—topological equivalence—and that any topologically equivalent sets either both have the fixed point property or neither does.

Central to this key notion of equivalence is the definition of a bijection. A function $f : X \to Y$ is a bijection if every $y \in Y$ is the image of exactly one point $x \in X$. So, a bijection gives a correspondence that pairs up the points between X and Y. Assuming f is a bijection, define an inverse function $f^{-1} : Y \to X$ by the rule that $f^{-1}(y) = x$ where x is the unique point in X for which $f(x) = y$. Bijections are precisely those functions that have inverses.

Definition 9.13. We say that two sets $X, Y \subseteq \mathbb{R}^n$ are *topologically equivalent* if there is a bijection $g : X \to Y$ with the property that both g and g^{-1} are continuous.

Example 9.14. Consider the intervals $[0, 1]$ and $[0, 2]$ in $\mathbb{R}$. The function $g : [0, 1] \to [0, 2]$ given by $g(x) = 2x$ is a bijection between $[0, 1]$ and $[0, 2]$. Note that $g^{-1} : [0, 2] \to [0, 1]$ is given by $g^{-1}(x) = \frac{1}{2}x$. Since both g and g^{-1} are continuous, $[0, 1]$ and $[0, 2]$ are topologically equivalent.

Notably, the following proposition shows that whenever one of two topologically equivalent sets has the fixed point property, the other does, too. For example, Theorem 9.9 proved that the interval $[0, 1]$ has the fixed point property, so the topological equivalence from the previous example means that $[0, 2]$ also has the fixed point property.

Proposition 9.15. *If X and Y are topologically equivalent and X has the fixed point property, then Y has the fixed point property.*

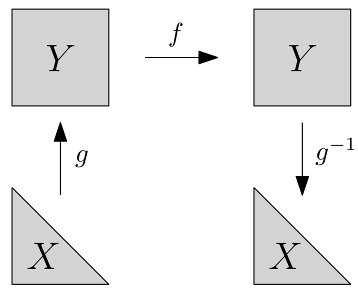

Figure 9.14. Inheriting the fixed point property

Proof. Assume that X and Y are topologically equivalent and also that X has the fixed point property. To prove that Y has the fixed point property, let $f : Y \to Y$ be an arbitrary continuous function. Since X and Y are topologically equivalent, there exists a continuous bijection $g : X \to Y$ so that g^{-1} is also continuous. Now combine the functions g, f, and g^{-1} as in Figure 9.14. More precisely, construct a new function from X to itself given by the rule

$$x \to g^{-1}(f(g(x))).$$

Since this function is continuous and X has the fixed point property, it must have a fixed point. So, we may choose a point $x \in X$ for which $x = g^{-1}(f(g(x)))$. Next apply the function g to both sides of this equation to get

$$g(x) = f(g(x)).$$

Now set $y = g(x)$ and observe that $y \in Y$ and $f(y) = y$. Thus f has a fixed point and, since f was an arbitrary continuous function, we conclude that Y has the fixed point property. □

With these tools in hand, we are now ready to prove the main theorem for this section.

Corollary 9.16. *The strategy space $S_{2,2}$ has the fixed point property.*

Proof. With Theorem 9.10 and the above proposition, we can prove that $S_{2,2}$ has the fixed point property by showing that $S_{2,2}$ is topologically equivalent to Δ_2. This follows from the continuous functions $g : \Delta_2 \to S_{2,2}$ and $g^{-1} : S_{2,2} \to \Delta_2$ defined by

$$g(x,y) = \left(x + \tfrac{x+y-|x-y|}{2}, y + \tfrac{x+y-|x-y|}{2}\right),$$
$$g^{-1}(x,y) = \left(x - \tfrac{x+y-|x-y|}{4}, y - \tfrac{x+y-|x-y|}{4}\right). \qquad \square$$

There is nothing particularly special about this triangle and square. It is tangential to our investigations here, but, in fact, every solid polygon and every circle plus its interior are topologically equivalent to one another and to many more shapes as well.[4]

Higher Dimensions. In a 2×2 matrix game, a mixed strategy for Rose has the form $\begin{bmatrix} p & 1-p \end{bmatrix}$ so it can be described with a single real number p. More generally, in an $m \times n$ matrix game, Rose's mixed strategies will have the form $\mathbf{p} = \begin{bmatrix} p_1 & \cdots & p_m \end{bmatrix}$. Here again there is some redundancy. Since $p_1 + \cdots + p_m = 1$, the last coordinate of Rose's mixed strategy can be deduced from the earlier ones: $p_m = 1 - (p_1 + \cdots + p_{m-1})$. Therefore, each possible mixed strategy for Rose can be associated with a point $(p_1, \ldots, p_{m-1})$ where $p_1, \ldots, p_{m-1} \geq 0$ and $p_1 + \cdots + p_{m-1} \leq 1$. We will adopt this convenient description of Rose's mixed strategies. Correspondingly, each mixed strategy $\mathbf{q} = \begin{bmatrix} q_1 \\ \vdots \\ q_n \end{bmatrix}$ for Colin can be associated with the point $(q_1, \ldots, q_{n-1})$, which will satisfy $q_1, \ldots, q_{n-1} \geq 0$ and $q_1 + \cdots + q_{n-1} \leq 1$.

With this interpretation, define the *strategy space* of a general $m \times n$ matrix game A to be

$$S_{m,n} = \{(p_1, \ldots, p_{m-1}, q_1, \ldots, q_{n-1}) \in \mathbb{R}^{m+n-2} \mid$$
$$p_1, \ldots, p_{m-1}, q_1, \ldots, q_{n-1} \geq 0, \textstyle\sum_{i=1}^{m-1} p_i \leq 1, \text{ and } \sum_{j=1}^{n-1} q_j \leq 1\}.$$

Just as in the 2×2 case, each point in the strategy space corresponds to a pair of strategies, one for Rose and one for Colin. In the same way that

[4]More generally, any two closed convex sets in $\mathbb{R}^n$ with nonzero volume are topologically equivalent.

we used the 2D version of the Brouwer Fixed Point Theorem to prove that $S_{2,2}$ has the fixed point property, the general version of Brouwer's Theorem can be used to prove that $S_{m,n}$ has the fixed point property. Full details of this similar argument appear in Appendix C.

Lemma 9.17. *For every pair of positive integers m, n, the strategy space $S_{m,n}$ has the fixed point property.*

9.4. Nash Flow and the Proof

This section finally concludes our proof of Nash's Equilibrium Theorem. We have already done the hard work of proving that the strategy space has the fixed point property. What remains is to introduce a continuous function called Nash flow on the strategy space and use the fixed point property to locate a Nash equilibrium.

Nash Flow. Imagine that Rose and Colin are playing a game of Ping-Pong. Rose notices that she is somewhat more successful today when she hits to Colin's backhand rather than his forehand. It would be silly for Rose to respond by hitting every single ball to Colin's backhand—he would quickly realize and exploit her strategy. It would make more sense for Rose instead to adjust and play a slightly higher percentage of her shots to his backhand.

This situation has a very natural and important game-theoretic analogue. Suppose that Rose and Colin are playing a 2×2 matrix game, with Rose using the strategy $\mathbf{p} = \begin{bmatrix} p & 1-p \end{bmatrix}$ and Colin using the strategy $\mathbf{q} = \begin{bmatrix} q \\ 1-q \end{bmatrix}$. If Rose observes that she does better playing the second row than playing the first against strategy $\mathbf{q}$, she might decide to modify her strategy $\mathbf{p}$ to play row 2 more often. As in the Ping-Pong game above, a subtle adaptation makes more sense than an abrupt change to playing the second row 100% of the time. Nash flow exhibits this idea of making a small adjustment to improve.

For a 2×2 matrix game A, Nash flow is a function, denoted f_A, that maps a point (p, q) in the strategy space to another point in the strategy space $f_A(p, q) = (p', q')$. Formally, then, $f_A : S_{2,2} \to S_{2,2}$. We think of this as updating Rose's and Colins' strategies $\begin{bmatrix} p & 1-p \end{bmatrix}$ and $\begin{bmatrix} q \\ 1-q \end{bmatrix}$

to new strategies $\begin{bmatrix} p' & 1-p' \end{bmatrix}$ and $\begin{bmatrix} q' \\ 1-q' \end{bmatrix}$. Each player modifies his or her initial strategy to do better against the other player.

Example 9.18. Figure 9.15 depicts the Nash flow function f_A for the matrix $A = \begin{bmatrix} (1,2) & (2,3) \\ (0,3) & (3,1) \end{bmatrix}$.[5]

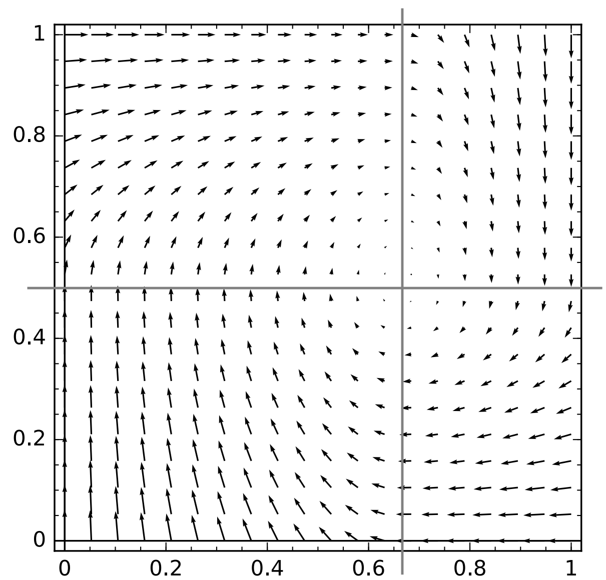

Figure 9.15. Nash flow for a 2×2 matrix

In Figure 9.15, arrows point to the right in the region where $q > 1/2$. This corresponds to Rose increasing the probability that she plays the first row. Rose indeed gets a better payoff if she plays the first row instead of the second row when $q > 1/2$. On the other hand, in the region where $q < 1/2$, the vectors are all directed to the left, indicating Rose's adjustment to decrease the probability that she chooses the first row. When $q < 1/2$, Rose in fact gets a better payoff playing the second row instead of the first. Near the point $(2/3, 1/2)$, the arrows shink to length zero. This

[5]This plot was created using the computing program Sage. See `www.ams.org/bookpages/stml-80` for a link to the code.

suggests that Rose has little incentive to modify her strategy in the region. Indeed (2/3, 1/2) is the unique fixed point of the Nash flow f_A, corresponding to the fact that $p = 2/3$ and $q = 1/2$ form the unique Nash equilibrium of A, so neither player has any incentive to change from this point.

Definition 9.19 (Nash Flow)**.** The Nash flow for a 2×2 matrix game A is a continuous function $f_A : S_{2,2} \to S_{2,2}$. As usual, let R and C denote Rose's and Colin's payoff matrices and let (p, q) be a point in the strategy space. The precise definition of the function will describe how to compute the new point $f_A(p, q) = (p', q')$.

Point (p, q) in the strategy space is associated with the strategies $\mathbf{p} = \begin{bmatrix} p & 1-p \end{bmatrix}$ for Rose and $\mathbf{q} = \begin{bmatrix} q \\ 1-q \end{bmatrix}$ for Colin. Rose's (expected) payoffs for playing row 1 or 2 against Colin's strategy $\mathbf{q}$ are given by

$$\begin{bmatrix} a_1 \\ a_2 \end{bmatrix} = R\mathbf{q}.$$

So a_1 is Rose's payoff if she plays row 1 against Colin's $\mathbf{q}$ strategy, and a_2 is her payoff if she plays row 2. Similarly, if Rose plays $\mathbf{p}$, Colin's payoffs for playing either column 1 or column 2 are

$$\begin{bmatrix} b_1 & b_2 \end{bmatrix} = \mathbf{p}C.$$

To articulate the players' updated strategies, first define the *plus function* as follows:

$$(x)^+ = \begin{cases} x & \text{if } x \geq 0, \\ 0 & \text{if } x < 0. \end{cases}$$

Now, at last, we can define Nash flow for the point (p, q). Let $f_A(p, q) = (p', q')$ where

$$p' = \frac{p + (a_1 - a_2)^+}{1 + |a_1 - a_2|}, \qquad q' = \frac{q + (b_1 - b_2)^+}{1 + |b_1 - b_2|}.$$

Let's investigate the situation. The behavior of p' and q' are very similar, so this analysis focuses on p'. First note that the denominator in the expression for p' is positive and the numerator is nonnegative, so $p' \geq 0$. Since this denominator is always at least as large as the numerator, $p' \leq 1$. Therefore $0 \leq p' \leq 1$, so p' does correspond to a strategy for Rose.

If $a_1 = a_2$, then $p' = p$ and Rose will not alter her strategy. She gets the same payoff playing row 1 or row 2 so she has no incentive to change. Next suppose Rose does better playing row 2 than row 1, so $a_2 > a_1$. In this case the formula says $p' = \frac{p}{1+t}$ where $t = |a_1 - a_2| > 0$. If $p = 0$, then Rose is already playing the pure strategy of row 2 and $p' = 0$, so she will not change. On the other hand, if $p > 0$, then $p' < p$ so Rose's new strategy will have her playing the first row with lower probability and the second with higher probability. In the $a_2 < a_1$ case, Rose does better playing the first row than the second. Here, the formula simplifies to $p' = \frac{p+t}{1+t}$ where $t = a_1 - a_2 > 0$. If $p = 1$, then Rose is already playing the pure strategy of choosing row 1 and $p' = p$ so she will not change. Otherwise, $p < 1$ and $p' > p$, so Rose's new strategy will have her play the first row with lower probability and the second with higher, just as desired.

In sum, this modification scheme gives Rose a sensible response. If row 1 and row 2 give her equal payoffs, she does not change strategy. If row 1 gives Rose a better payoff than row 2, then she modifies her strategy to play row 1 more frequently (if possible). Similarly, if row 2 gives a better payoff than row 1, then Rose alters her strategy to play row 2 more frequently (if possible). Rose's new strategy p' will be exactly the same as her original p if and only if p is a best response to q. A similar analysis for Colin results in the following key property.

Lemma 9.20. *For every 2×2 matrix A, the fixed points (p, q) of Nash flow f_A are precisely those points for which $\mathbf{p} = \begin{bmatrix} p & 1-p \end{bmatrix}$ and $\mathbf{q} = \begin{bmatrix} q \\ 1-q \end{bmatrix}$ form a Nash equilibrium of A.*

Nash flow can be very helpful for getting a sense of the strategy space for a particular game. Let's reconsider this function for another one of the dilemmas from Chapter 7.

Example 9.21. Figure 9.16 depicts the Nash flow associated with the Dating Dilemma game given by the matrix $\begin{bmatrix} (2,1) & (0,0) \\ (0,0) & (1,2) \end{bmatrix}$. This game has two pure Nash equilibria corresponding to the case when the players go to the same venue. These outcomes correspond to the fixed points

(0, 0) and (1, 1) in the Nash flow. There is an additional Nash equilibrium when Rose plays $\begin{bmatrix} 2/3 & 1/3 \end{bmatrix}$ and Colin plays $\begin{bmatrix} 1/3 \\ 2/3 \end{bmatrix}$.

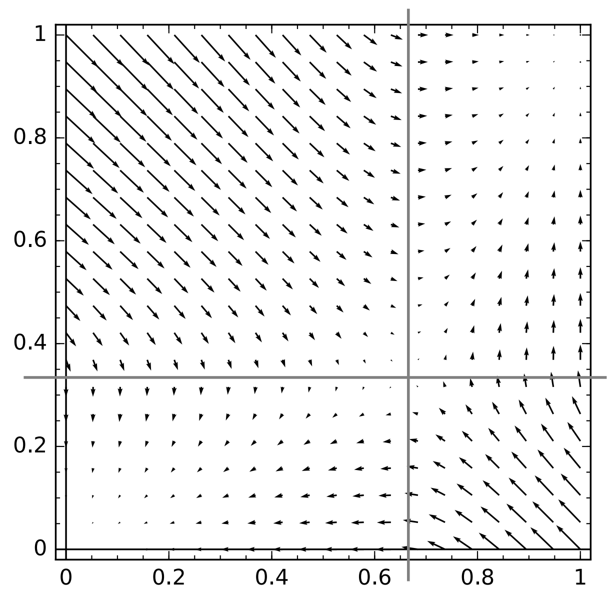

Figure 9.16. Nash Flow for the Dating Dilemma

In closing, we note the considerable value of Nash flow for studying stability properties of Nash equilibria (as in the discussion of evolution in Chapter 8). It's evident in the Dating Dilemma that the two pure strategy Nash equilibria are very stable—all points in the strategy space near either of these points are directed toward it. The mixed strategy equilibrium is less stable since some nearby points are directed away from it.

The Proof. We are finally ready to prove Nash's Equilibrium Theorem. This proof is, in fact, a fairly straightforward consequence of Corollary 9.16 and Lemma 9.20.

Theorem. *Every* 2×2 *matrix game* A *has a Nash equilibrium.*

Proof. The existence of a Nash equilibrium for the 2×2 game A follows from the statements below:

(1) Nash flow is a continuous function $f_A : S_{2,2} \to S_{2,2}$ with the property that every fixed point corresponds to a Nash equilibrium of A (Lemma 9.20).

(2) The strategy space $S_{2,2}$ has the fixed point property (Corollary 9.16), and thus f_A has a fixed point. □

Higher Dimensions. Just as we defined for every 2×2 matrix game A the Nash flow $f_A : S_{2,2} \to S_{2,2}$, it is possible to define a (continuous) Nash flow function for an arbitrary $m \times n$ matrix game. The details appear explicitly in Appendix C, but the key is the following generalization of Lemma 9.20.

Lemma 9.22. *For every $m \times n$ matrix A, the fixed points of Nash flow $f_A : S_{m,n} \to S_{m,n}$ are precisely those points that correspond to a Nash equilibrium of A.*

With this in place, the full proof of the Nash Equilibrium Theorem for an arbitrary $m \times n$ matrix game follows from the same reasoning as in the 2×2 case.

Theorem (Nash)**.** *Every matrix game has a Nash equilibrium.*

Proof. Let A be an $m \times n$ matrix game with strategy space $S_{m,n}$. The proof that A has a Nash equilibrium follows from the following two properties:

(1) Nash flow is a continuous function $f_A : S_{m,n} \to S_{m,n}$ with the property that every fixed point corresponds to a Nash equilibrium of A (Lemma 9.22).

(2) The strategy space $S_{m,n}$ has the fixed point property (Lemma 9.17), so f_A has a fixed point. □

Exercises

(1) Find all fixed points for the functions below:
  (a) $f : [0,2] \to [0,2]$ defined by $f(x) = 1 + x^2/4$,
  (b) $g : [-1,2] \to [-1,2]$ defined by $g(x) = x^2/3 + 2x/3 - 1/4$,
  (c) $h : [0,4] \to [0,4]$ defined by $h(x) = x^2 - 3x + 3$.

(2) Find all fixed points for the functions below:
  (a) $f : \mathbb{R}^2 \to \mathbb{R}^2$ defined by $f(x,y) = (y-1, x+y)$,
  (b) $g : \mathbb{R}^2 \to \mathbb{R}^2$ defined by $g(x,y) = (2-y, -2x+1)$,
  (c) $h : \mathbb{R}^2 \to \mathbb{R}^2$ defined by $h(x,y) = (xy-1, 3x+2)$.

(3) Define the square $S = \{(x,y) \in \mathbb{R}^2 \mid 0 \le x \le 1 \text{ and } 0 \le y \le 1\}$. Sketch (as in Figure 9.8) each function from S to itself:
  (a) $f : S \to S$ given by $f(x,y) = (x^2, y^2)$,
  (b) $g : S \to S$ given by $g(x,y) = (x/2, 1-y/2)$,
  (c) $h : S \to S$ given by $h(x,y) = (1-y, x)$.

(4) For $A = \begin{bmatrix} (0,5) & (3,3) \\ (6,2) & (1,4) \end{bmatrix}$, evaluate the Nash flow function f_A at each point:
  (a) $f_A(1/3, 1/3)$,
  (b) $f_A(1/2, 1/4)$,
  (c) $f_A(0,t)$ for $0 \le t \le 1$,
  (d) $f_A(t,t)$ for $0 \le t \le 1$.

(5) This problem concerns functions that may be discontinuous.
  (a) Construct a function $f : [0,1] \to [0,1]$ with no fixed point.
  (b) For every positive integer k, construct a function $f : [0,1] \to [0,1]$ which has exactly k fixed points.

(6) Find a continuous function $f : \mathbb{R}^2 \to \mathbb{R}^2$ with the following given property:
  (a) f has no fixed points.
  (b) f has exactly one fixed point.
  (c) f is not the identity but has infinitely many fixed points.

(7) Consider the matrix $A = \begin{bmatrix} (1,2) & (2,3) \\ (0,3) & (3,1) \end{bmatrix}$ from Example 9.18. When Rose plays the mixed strategy $\begin{bmatrix} p & 1-p \end{bmatrix}$, Colin gets a higher payoff playing the first column when $p < 2/3$ and a higher payoff playing

the second column when $p > 2/3$. This corresponds to the fact that the Nash flow function goes upward to the left of the line $p = 2/3$ and downward to the right. By considering Rose's payoffs we can determine that the Nash flow function will go rightward above the line $q = 1/2$ and leftward below the line $q = 1/2$. Based on this we obtain a rough plot of Nash flow shown in the figure below (compare with Figure 9.15).

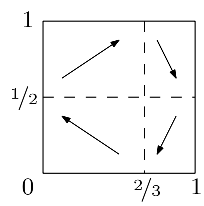

Find rough plots of Nash flow for the following matrix games:

(a) $\begin{bmatrix} (3,1) & (0,6) \\ (1,4) & (5,3) \end{bmatrix}$, (b) $\begin{bmatrix} (1,3) & (3,8) \\ (9,3) & (2,2) \end{bmatrix}$, (c) $\begin{bmatrix} (1,6) & (3,4) \\ (5,2) & (4,3) \end{bmatrix}$.

(8) Prove that the given set does not have the fixed point property.
   (a) $\{x \in \mathbb{R} \mid 0 < x < 1\}$,
   (b) $\{(x,y) \in \mathbb{R}^2 \mid x \geq 0 \text{ and } y \geq 0\}$,
   (c) $\{(x,y) \in \mathbb{R}^2 \mid 0 < x < 1 \text{ and } 0 < y < 1\}$,
   (d) $\{(x,y) \in \mathbb{R}^2 \mid 0 < x^2 + y^2 < 1\}$.

(9) Construct continuous functions with the following properties:
   (a) $f : [0,\infty) \to [0,\infty)$ has 0 as its unique fixed point.
   (b) $g : [0,1] \to [0,1]$ had 0, 1 as its only fixed points.
   (c) $h : \mathbb{R} \to \mathbb{R}$ has exactly k fixed points.

(10) Let $a < b$ and consider the closed interval $[a,b]$.
   (a) Prove that $[a,b]$ has the fixed point property by generalizing the proof of Theorem 9.9.
   (b) Prove that $[a,b]$ is topologically equivalent to $[0,1]$ by generalizing Example 9.14.

(11) This exercise uses Sperner's Lemma to prove that Hex cannot end in a draw. Suppose that Rose is marking cells with $*$ and trying to connect the lower left side with the upper right side, while Colin is marking cells with $\circ$ and trying to connect the lower right side with the upper left. Assume that they continue play until all cells

are occupied. Now construct a subdivided triangle on top of the Hex board as shown in the figure below:

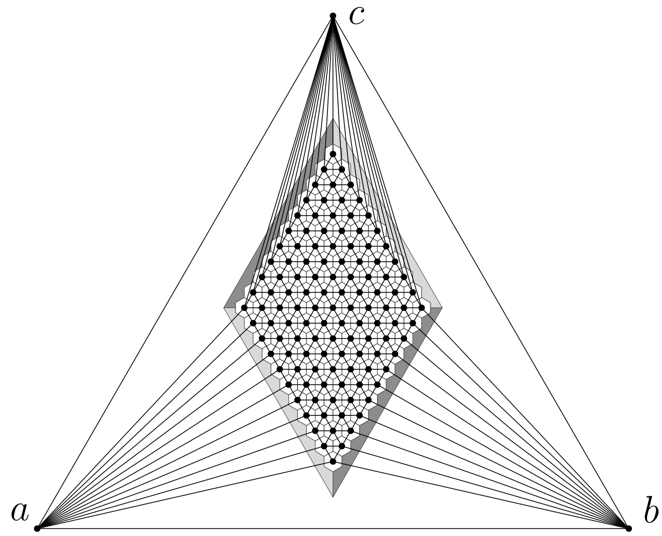

Next we label vertices of the subdivided triangle. Give a label 1, b label 2, c label 3, and label every other vertex v according to the following:

- v gets label 1 if its cell contains $*$ and this cell can be joined to the lower left side by a connected path of cells labeled $*$.
- v gets label 2 if its cell contains $\circ$ and this cell can be joined to the lower right side by a connected path of cells labeled $\circ$.
- v gets label 3 if neither of the above conditions apply.

Prove that either Rose or Colin won the game.

(12) If A is an $n \times n$ matrix and $\mathbf{b} \in \mathbb{R}^n$, the function $f : \mathbb{R}^n \to \mathbb{R}^n$ given by $f(\mathbf{x}) = A\mathbf{x} + \mathbf{b}$ is an *affine transformation*.

(a) Show that the composition of two affine transformations is an affine transformation.

(b) Determine when the affine transformation $f(\mathbf{x}) = A\mathbf{x} + \mathbf{b}$ is invertible, and find a formula for the inverse when it exists.

(c) Let $\mathbf{e_i} \in \mathbb{R}^n$ be the vector with a 1 in coordinate i and 0 elsewhere. If $\mathbf{x_1}, \dots, \mathbf{x_n}$ are linearly independent, find an invertible matrix A so that $A\mathbf{x_i} = \mathbf{e_i}$ for every $1 \le i \le n$.

(d) If $\mathbf{x_0}, \dots, \mathbf{x_n} \in \mathbb{R}^n$ do not lie on a common hyperplane, show that there exists an invertible affine transformation $f : \mathbb{R}^n \to \mathbb{R}^n$ so that $f(\mathbf{x_0}) = \mathbf{0}$ and $f(\mathbf{x_i}) = \mathbf{e_i}$ for $1 \leq i \leq n$.
(e) Prove that every simplex in $\mathbb{R}^n$ is topologically equivalent to Δ_n.
(f) Prove that any two simplices in $\mathbb{R}^n$ are topologically equivalent.

(13) Let A be a symmetric 2×2 matrix game with payoff matrix $R = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ for Rose (so Colin's payoff matrix is $C = \begin{bmatrix} a & c \\ b & d \end{bmatrix}$). In this exercise we will follow the paradigm from this chapter to prove that A has a symmetric Nash equilibrium (i.e. we will prove Theorem 8.11 for the special case of 2×2 matrices).
(a) Show that for every $0 \leq t \leq 1$ the Nash flow function satisfies $f_A(t,t) = (u,u)$ for some u.
(b) Define a function $g : [0,1] \to [0,1]$ by the rule $g(t) = u$ if $f_A(t,t) = (u,u)$. Use a theorem from this chapter to show that g has a fixed point, and then deduce from this that A has a symmetric Nash equilibrium.

(14*) Let T be the boundary of the triangle Δ_2. Prove that there does not exist a continuous function $f : \Delta_2 \to T$ with the property that every $\mathbf{x} \in T$ is a fixed point of f. Hint: Use Sperner's Lemma.