Chapter Three: Static Games*

Game-theoretic modeling often begins with the simplest of structures, either in extensive or in normal form. Such simple structures are meant to define the players, their available actions and their priorities. These basic structures can then be extended in two main directions: incomplete information or their repetition with discounting of the future. However, even basic game structures can provide meaningful and relevant models of real-world issues.

3.1 The Meaning of Mixed Strategies

Since its very beginnings with von Neumann's findings, Game Theory has faced a serious interpretation problem: the overwhelming majority of equilibria involves some probabilities that a strategy or another will be used. But what does it mean to use a strategy with probability? There are several arguments that fit various circumstances.

3.1.1 Unpredictability

A mixed-strategy equilibrium typically arises when the best-reply to best-reply to best-reply process described earlier fails to stabilize. So, the probability mixings of each player is the only way to prevent any other from finding a *single pure* best reply. Indeed, any deviation from the mixing that would yield a single pure best reply would usually make that new mixing sub-optimal in response since it would re-initiate the best-reply to best-reply sequence. So, using the right probabilities of play is often a way to avoid being exploited by another player. The unpredictability removes any incentive to each of the players to stick to any single best reply. This argument is best illustrated by the well known game of "Matching Pennies:" Two players simultaneously display a penny's side (Heads or Tails.) If the displayed sides match one player, say Blue, wins. Otherwise, the other player Red wins. This simultaneous-play zero-sum game is illustrated in Figure 3.1 together with its obvious Nash equilibrium.

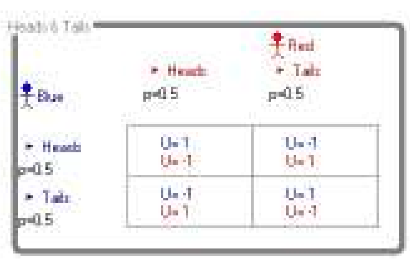

Figure 3.1

If Red can be rationally expected to have even the smallest bias for choosing Heads, with probability $p > \frac{1}{2}$, the expected value to Blue of choosing heads becomes

$$\mathbb{E}(\text{Heads}) = p - (1-p) = 2p - 1 > 0$$
$$> \mathbb{E}(\text{Tails}) = -p + (1-p) = 1 - 2p$$

It follows that Blue will choose Heads with certainty and, in doing so, will gain an advantage. Since the game is zero-sum, this advantage is at the expense of Red. So, deviating from the $p = \frac{1}{2}$ expected play opens the door to exploitation.

Of course, the argument does not explain where this expectation comes from. One possible way to form an expectation is to observe the frequency of prior play if the game is indefinitely repeated against the same opponent. This suggests the next interpretation.

3.1.2 Learning

One early argument in favor of probabilistic play was inspired by the assumptions of perfect knowledge and rationality. If a player can think in terms of what the others are thinking about what they are thinking, ad infinitum, then they may engage into an exercize of "fictitious play." Blue could think in the Matching Pennies game: if I start with Head the best reply for Red would be Tails to which I would best respond with Tails, to which Red would best respond with Heads... As this tought process continues, Blue can form averages of his and his opponent's best choices and fictitiously pick his own and his opponent's best replies to these averages. Continuing in this fashion the averages clearly converge to $p = \frac{1}{2}$ in the game of Matching Pennies. Indeed, they can be proven to converge to a Nash equilibrium in all zero-sum games. Unfortunately, this is not always the case when the game becomes non-zero sum. There are counter-examples where convergence never takes place.

3.1.3 Threat or Pledges

In many extensive form games, an expected move is interpretable as a threat or a pledge of making a specific choice when one's turn comes. Indeed, the equilibrium exhibited in Figure 1.17, in Chapter One, can be interpreted in precisely such terms: Red may threaten to play Safe with just enough probability $p = \frac{2}{3}$ to discourage, at least partly, the Cheat-Blue to choose Continue. And this is a better plan than the alternative equilibrium that would have her play Safe with certainty, a plan that discourages even the Honest-Blue to pick Continue. Indeed, her expected payoff of playing the game with this probabilistic threat is $\mathbb{E}(\text{@Start}) = \frac{-5}{8}$, whereas it is $\mathbb{E}(\text{@Start}) = -1$, when always playing Safe. So, the threat of Safe with probability $p = \frac{2}{3}$ partly deters, with probability $p = \frac{5}{6}$, the Cheat-Blue from choosing Continue. Of course, Red's thinking is based on that precise expectation which yields her belief $b = \frac{1}{3}$ of facing the Cheat-Blue, a prerequisite to making the threat credible.

3.1.4 Uncertainty

Perhaps the most sophisticated attempt at explaining mixed strategy equilibria is due to Nobel laureate John Harsanyi who proposed a process known as "purification." Harsanyi introduces a small interval of uncertainty about each sides true payoffs. That uncertainty is pictured as a uniform probability that the true player's payoff fall at any point of the interval. In turn, this translates into a belief by all other players that they are facing any particular type. When making expected utility calculations with such beliefs, all players end up with probabilities that each of their moves is a best reply to the pure best choices of each of the other sides' types. This multilateral uncertainty then yields probabilistic best replies.

But if the uncertainty is very small, one may investigate the limit as the interval of uncertainty shrinks to a single point of these probabilities. It turns out that, under some not unreasonable conditions, these probabilities converge to a Nash equilibrium of the game. This explanation is dependent upon some technical assumptions. For instance, a Nash equilibrium that involves "locally dominated" strategies can fail to arise from that process.[1] But the concept is still appealing: "almost negligible" uncertainty over the players' payoffs can justify the choice of mixed strategy equilibria.

3.1.5 Perfect Recall

Except for the simplest cases, a strategy involves a number of decisions at several of the same players' turns. The probabilities involved in a behavioral strategy should therefore translate into probabilities on the very moves it involves. And one would expect that behavioral strategies would be equivalent to the mixed strategies of the normal form.

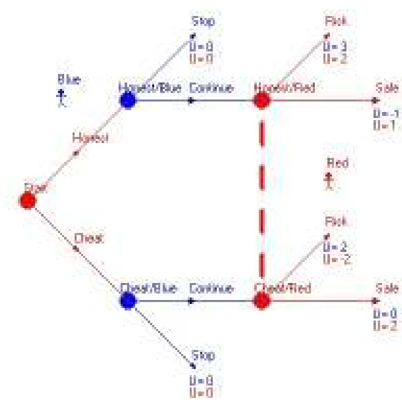

Figure 3.1: A Game without Perfect Recall

This is unfortunately not the case as examplified by the game of Figure 3.1. When writing the normal of that game one obtains:

[1]A locally dominated strategy is one that fails to be a best reply to a specific mix of other players' strategies. By contrast, a "globally" dominated strategy is one that is never a best reply and can be ignored in all strategic thinking (and even the very specification of a normal form game.)

Table

Blue \ Red	Hon&Risk	Hon&Safe	Cheat&Risk	Cheat&Safe
Stop&Stop	U= 0 U= 0	U= 0 U= 0	U= 0 U= 0	U= 0 U= 0
Cont&Stop	U= 3 U= 2	U= -1 U= 1	U= 0 U= 0	U= 0 U= 0
Stop&Cont	U= 0 U= 0	U= 0 U= 0	U= 2 U= -2	U= 0 U= 2
Cont&Cont	U= 3 U= 2	U= -1 U= 1	U= 2 U= -2	U= 0 U= 2

Figure 3.2: A non-Matching Normal Form

Solving this normal form yields a number of Nash equilibria, for instance, the pure strategy Cont&Cont for Blue against the mixed strategy $(\frac{1}{2}, 0, 0, \frac{1}{2})$ for Red. But then, either Honest or Safe must have probability zero. But that contradicts the non-zero probabilities of Hon&Risk and Cheat&Safe. So, this Nash equilibrium of the normal form has no corresponding behavioral strategy equilibrium in the extensive form.

The condition known as "perfect recall" guarantees that this mismatch between behavioral and mixed strategies never happens. In a nutshell, perfect recall means that no player ever forgets previously held information. Here, if and when Red reaches her second turn, the information set means that she essentially forgot what she did at her first turn.

3.2 Normal Form Games

As was seen in Chapter One, the normal form of a game does not adequately capture issues of timing or information. But when such issues are not central the normal form provides an attractive modeling option. This is often the case when two or more individuals make quasi-simultaneous strategic decisions that cannot be revisited in view of the unfolding of the game and that involve no expectation of consequences in future game playing. The Prisoner's dilemma and the Battle of the Sexes, at least in the terms discussed previously, are typical examples: the two prisoners make their decisions independently while possible future retaliation is neglected. Indeed, introducing the very possibility of retaliation profoundly affects the strategic picture, a subject that will be discussed at length in Chapter Five. Similarly, the Battle of the Sexes is viewed as a one-shot issue: perhaps the couple failed to communicate beforehand and must make their respective decisions on the spot without the chance to consult each other. Better communication could solve the problem by providing an expectation of what the other will do. However, the strategic issue remains whole when they attempt to reach an expectation: indeed, there are three Nash equilibria in that game. Which one will prevail?

3.2.1 Some Famous Static Games

The Battle of the Sexes (see Chapter One exercises) belongs to the class of "coordination games." The game has two pure equilibria: He chooses the fight and She goes along. Or She chooses the ballet and He goes along. There is a third Nash equilibrium where they each make probabilistic choices that depend on how much they value their respective outcomes. That solution is picture in Figure 3.1.

It is easy to verify that the probabilities involved change with the respective magnitudes of the payoffs each side assigns to the various outcomes: one simply edits the game in *GamePlan* and solves it. Coordination games can be far more complex than that. They can involve more choices and more players. Another famous such game is the "Stag Hunt" (see homework.)

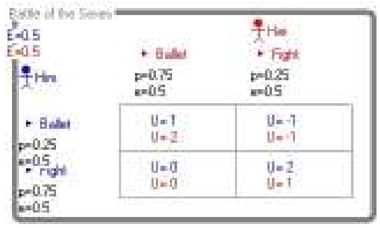

Figure 3.3: A Solution of the Battle of the Sexes

Another prominent normal form game is Chicken. This is supposedly a formal model of the "game of Chicken" described by Herman Kahn as a metaphor for nuclear crises. The game, supposedly played by California teenagers in the 1950s, involves two cars speeding toward each other in the middle of the road. The first one to swerve is "chicken." The game is really one of timing and uncertainty about the other side's resolve. But game theory was not that sophisticated at the time and the "game model" pictured in Figure 3.4 was proposed.

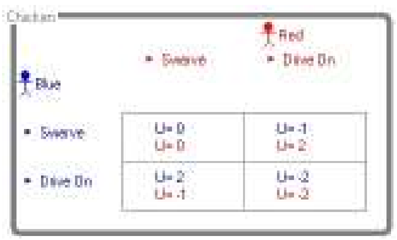

Figure 3.4: Chicken

If each side commits in advance to a strategy it has to make assumptions about the other. The standard search for a stable cell, in the best reply to best reply process, can, for instance, begin in the up-right one: under that assumption Blue would not choose to Drive On and Red would not choose to Swerve. So, that cell provides a Nash equilibrium. But so does the down-left cell where Blue drives on and Red swerves. Worse, there is also a mixed-strategy equilibrium that begs for an interpretation.

Volumes have been written about the above game of Chicken, mostly in the International Relations literature where it is still widely viewed as a game model of nuclear crises. Needless to say, volumes have also been devoted to criticizing it as a model and improving it in various ways.

3.2.3 More Strategies and More Players

There are several interesting structures that arise from the so-called 2 by 2 games, meaning normal form games with just two players and two strategies per player (see homework.) But the possibilities explode geometrically when adding strategies and players. A two-person game with three strategies per player (3 by 3) is pictured in Figure 3.5.

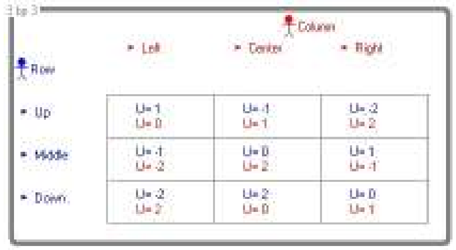

Figure 3.5: A "3 by 3" Game

The standard search for a stable cell can begin, for instance, in the up-left corner. Row would not want to switch strategy from there but Column would switch to Right, followed by Row's switch to Middle, Column's switch to Center, Row's switch to Down, Column's switch to Left and Row's switch to Up, back to the beginning. There is, indeed, no pure strategy Nash equilibrium in that game. There is, however, a single mixed-strategy equilibrium.

A three-person game with two strategies per player is pictured in Figure 3.4 using the *GamePlan* format. The left-hand side table shows the payoffs when Green chooses Coop while the right-hand side table shows them when Green picks Dfct. This is in fact a possible version of a 3-player Prisoner's Dilemma with a twist: if all three prisoners confess (Dfct), the evidence becomes overwhelming and the prosecutor sends them all off to prison for life! As a result, the unanimous choice of Dfct does not form an equilibrium and one side should instead hold off confession (Coop) expecting the other two to tell on each other. But this yields three possible pure-strategy equilibria, one for each of the three possible sides that chooses to not confess. There are in fact four further mixed-strategy equilibria (see homework.)

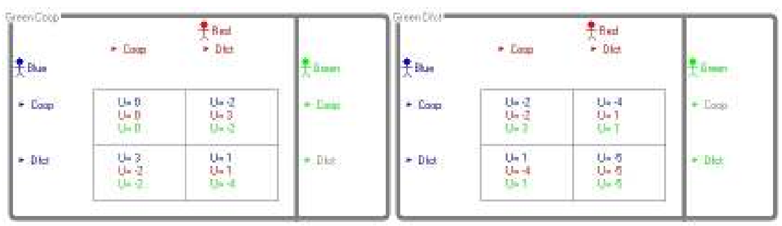

Figure 3.6: A "2 by 2 by 2" Game

Three or more player games can offer interesting modeling opportunities, especially to understand how some subset of players can coordinate their play successfully to the expense of the others.

3.2.4 Zero-Sum Games

John von Neumann was first to establish that normal form two-person zero-sum games always admit an equilibrium. Such a game is easy to spot: the payoffs in each cell must add up to zero (or to a given constant.) One easily sees that the game of Figure 3.3 does not have zero-sum. One game that has zero-sum is pictured in Figure 3.7.

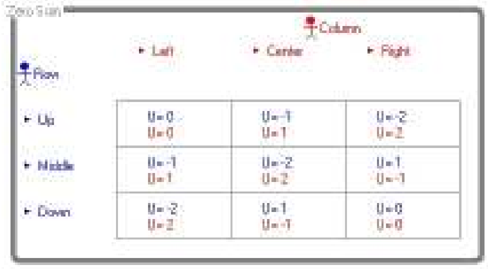

Figure 3.7: A 3-by-3 Zero-Sum Game

The same kind of search for equilibrium, a stable cell, can be conducted, sometime successfully if a pure-strategy equilibrium does exist. Here, it is easy to see that this is not the case. The only difference between zero and non-zero sum normal form games really lies in the characterization of their solution set: in zero sum games, any convex combination of solutions is still a solution whereas this doesn't hold in non-zero sum games. The solution of a zero sum game is always a convex set.

An interesting class of zero-sum games is given by issues of optimal allocation of resources between several theaters of conflict. Say there are three theaters of battle and each side has a fixed number of assets to allocate between the theaters. When one side allocates more than the other on one theater it wins that battle. An if one side wins two battles it wins the game otherwise it is a draw. This class of games is known as "Colonel Blotto" games.

3.2.5 Solving Normal Form Games

Aside from identifying dominant strategies, the simplest approach to solving a normal form game is the search of a stable cell in the best reply to best reply process that has been successfully applied to a few games so far. Unfortunately, that search will succeed only with a very small subset of games, those that have a pure-strategy equilibrium. For the vast majority of other games there is a need for more systematic and efficient techniques.

It is not difficult to show that a two-person zero-sum game is equivalent to a linear programming problem that has a solution. Indeed, John von Neumann's result is often called the "Minimax Theorem" because it shows the existence of a value V for the game, the optimal payoff received by one side that attempts to maximize it while the other (which receives -V since the game is zero-sum) attempts to minimize it. Linear programming problems can be solved systematically with various techniques of optimization such as the "Simplex Algorithm."

A two-person non-zero sum game in normal form can be much harder to solve. One main technique is called complementary pivot programming, more specifically the Lemke-Howson Algorithm. Contrary to zero-sum games, that method is not guaranteed to identify all equilibria and its extension to more than two players is just as unreliable. Various studies have established that solving non-zero sum games in normal form, especially those with more than two players, are "hard" problems that involve high mathematical complexity.

3.3 Extensive Form Games

The extensive form is the best modeling option when timing or information are important aspects of the players' strategic thinking. In particular, if players can adapt their choices to what is observed in the unfolding of the game, then chances are that the normal form could erase some critical strategic aspects.

3.3.1 Open versus Secret Ballots

A panel of three judges must vote on an appeal. Bob, the Chief Judge is a hard liner. He wants to uphold the law and tends to resent all appeals for leniency based of flimsy mitigating factors, as in the case under review. Gina is a liberal but also a consensus builder. In her heart of hearts she would like to grant the appeal, but she values highly her colleagues' opinions. She values a unanimous vote to reject the appeal as highly as seeing the appeal granted by a vote of two to one. Ronald is all for toughness when it comes to crime, but he is a born dissenter. And dissenting with the chief judge is what motivates Ron the most. His favorite outcome is a vote of two against one, with Bob casting the dissenting vote. His next best is to be in the minority.

Bob's as chief judge must decide on a voting procedure. One possibility is to ask each of his colleagues in turn, according to seniority, to express their vote to the panel. The meeting would begin with Bob announcing his vote rejecting the appeal. He would then ask Gina and Ronald to announce their preference in turn. In order to know what to expect, Bob sets up the Sequential Voting game model of Figure 3.8.

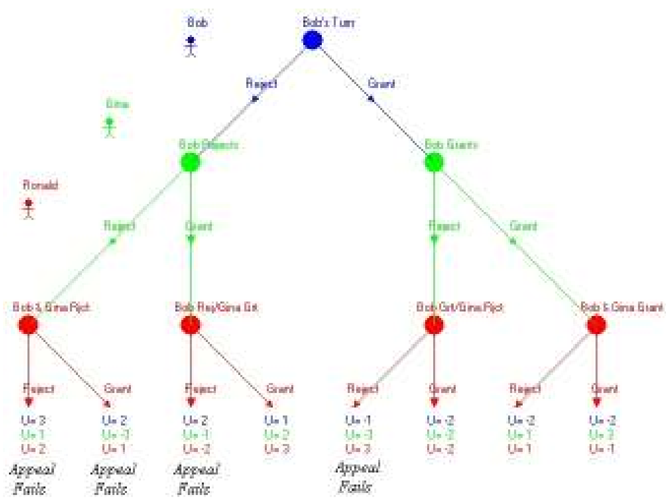

Figure 3.8: Sequential Voting

Bob having perfect knowledge of the game chose each judge's payoffs carefully to reflect their known priorities. He also knows how to apply the backward induction of rational behavior from the end moves. But he has an even more effective solving tool: a copy of *GamePlan* that he received as a gift from an admiring clerk. The solution of the game comes as a reality check: Ronald and Gina vote against him to grant the appeal!

Perhaps changing the voting order would produce a different result? But this could create protocol issues: Gina could object on being last to vote, after the two "boys." So, Bob has an idea: he can hold a secret ballot vote! He edits his *GamePlan* file and solves the game to obtain, to his delight, a completely rational unanimous vote to reject the appeal.

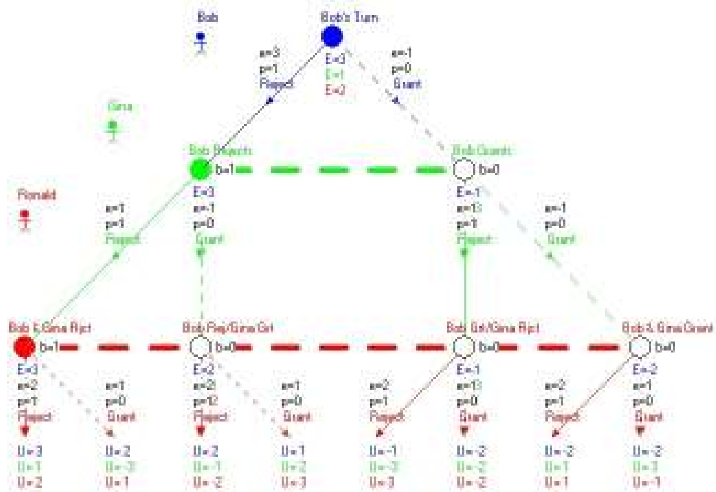

Figure 3.9: A Solution of the Secret Ballot Vote

Unfortunately, the same solution as in the sequential vote also appears as well as a probabilistic one according to *GamePlan*. But, even that last one has a minuscule probability $\mathbb{P}(\text{Appeal succeeds}) = \frac{1}{28}$. At least, Bob now has a serious shot at his favorite outcome by adopting a perfectly legal voting procedure!

3.4 The Subgame Perfect Equilibrium

As discussed with the edited Game 1 of Chapter 1, the Nash equilibrium can specify totally non-credible plans. Indeed, the implicit threat of choosing Left in the normal form of Figure 1.10 results in the Nash equilibrium {Stop,Left}, a better outcome than the other Nash equilibrium {Continue,Right} of the same game. But when seen in the extensive form of Figure 1.4, the threat appears rather flimsy: Blue can easily reason that it would never be carried out rationally by Red if tested by the choice Continue. However, the Nash equlibrium {Continue,Right} does not suffer from the same flaw.

In the extensive form game {Stop,Left} is still a Nash equilibrium while {Continue,Right} has a stronger property described as "subgame perfect." The concept was introduced by Selten in the following terms: a subgame of a game is any self-contained part of a game that constitutes a whole game in its own right. In particular, any move starting at a node of a subgame must still be in the subgame, and so should the node it leads to and any information set that contains such a node. In Figure 1.4, the subgame beginning at node Next is a subgame.[2] For Selten, a Nash equilibrium is "subgame perfect" (SPE) if it is a Nash equilibrium in all subgames. Just like their weaker Nash equilibrium cousin SPEs always exist.

However, the concept is irrelevant to normal form games. Indeed, a normal form game can be viewed as an extensive form game such as that of Figures 1.5 or 3.9 where all players essentially move at the same time. So, there is only one subgame in any such game: indeed, any part of it would have to separate an information set, something not allowed in the very definition of a subgame. In normal form games, any Nash equilibrium is an SPE.[3]

3.5 Continuous Games

A game is called "continuous" when the choice set of at least one player is the continuum to begin with, not just an extension of a discrete set to the probability distributions over that set. It may merely be a matter of interpretation when payoffs are linear in the choices. But it is unambiguous when payoffs are non-linear. Such non-linearities arise quite naturally in Economics. The Cournot Duopoly and Oligopoly models of Chapter One are good examples. Another good example is the Bertrand Duopoly.

3.5.1 The Bertrand duopoly

Joseph Bertrand, in a critique of Augustin Cournot's duopoly model in the late 19^{th} Century, pointed out that most businesses choose price rather than quantity. The idea that businesses compete on price can be approached in several ways. The following is a relatively simple model.

Assume that the consumers for a homogeneous product (such as gasoline) are distributed uniformly along Main Street, pictured as the segment $[0, 1]$ of the x-axis. Two suppliers A and B are located at points a and b on that segment. We assume that $0 \leq a < b \leq 1$. Business A charges price p_A and business B charges price p_B for the exact same product (say, a gallon of gas.)

Typical consumer x (located at position x) derives utility:

$$\mathrm{U}_A = \mathrm{U} - p_A - c(x - a)^2$$

by shopping at A and similarly, with p_A replaced by p_B and a replaced by b for U_B, the utility of shopping at B. $\mathrm{U} > 0$ is his satisfaction of obtaining the product. The quadratic term is a cost of traveling to either place. Clearly, consumer x will prefer shopping at A when $\mathrm{U}_A > \mathrm{U}_B$. Solving for equality $\mathrm{U}_A = \mathrm{U}_B$ yields the critical consumer z such that all $x < z$ shop at A and all $x > z$ shop at B. One has:

$$z = h + (p_B - p_A) \div 2cd \qquad (1)$$

---

[2]Although Blue never makes any move in that subgame, he is still present and has payoffs. So, this still counts as a game in Selten's definition.

[3]Selten also defined a "trembling-hand perfect" equilibrium for normal games that has fallen out of use.

with $d = b - a$ and $h = (a + b) \div 2$. So, with the given prices, business A will capture z customer and business B will get the remainder $(1 - z)$. If we further assume the same linear costs of supply kz for quantity z for both, we can write the profit (objective) functions for each business:

$$\mathrm{V}_A = z(p_A - k) \qquad \text{and} \qquad \mathrm{V}_B = (1 - z)(p_B - k) \tag{2}$$

In order to optimize profits, one differentiates V_A with respect to p_A to obtain the optimum p_A that is a function of p_B:

$$p_A = cdh + \tfrac{1}{2}(k + p_B) \tag{3}$$

and, similarly (verifying the down-concavity of profits in own decision variable):

$$p_B = cd(1 - h) + \tfrac{1}{2}(k + p_A) \tag{4}$$

As with Cournot, an equilibrium is achieved by solving (3) and (4) simultaneously, yielding:

$$p_A = k + \tfrac{2}{3}cd(1 + h) \text{ and } \quad p_B = k + \tfrac{2}{3}cd(2 - h) \tag{5}$$

There are many variations that are Bertrand-type duopolies or oligopolies that make price rather than quantity appear as the decision variable. Once equilibrium prices are determined as above, one can turn to the positioning game of choosing locations a and b (see homework.)

3.5.2 The Cournot Oligopoly

The Cournot Duopoly can be generalized to $n \geq 2$ players. Consider the demand defined by f as follows:[4]

$$p = f(Q) = \begin{cases} a - bQ & \text{for } 0 \leq Q \leq \frac{a}{b} \\ 0 & \text{for } Q > \frac{a}{b} \end{cases} \tag{6}$$

with $a, b > 0$. Total demand $Q = \sum_{j=1}^{n} q_j$ is therefore assumed to clear at price p given by (6). Again, neglecting costs, each "oligopolist" $i \in \{1, 2, .., n\}$ wishes to maximize:

$$\mathrm{U}_i(q_i, Q_{-i}) = q_i f(\sum_{j=1}^{n} q_j) \tag{7}$$

where Q_{-i} denotes $\{q_j | j \neq i\}$. One easily obtains from (7) the first order condition of optimization:[5]

$$\tfrac{\partial}{\partial q_i}\mathrm{U}_i(q_i, Q_{-i}) = a - b\sum_{j=1}^{n} q_j - bq_i = 0 \tag{8}$$

or $$q_i = \varphi_i(Q_{-i}) = \tfrac{a}{2b} - \tfrac{1}{2}\sum_{j \neq i} q_j \tag{9}$$

---

[4]Technically, f is called the "inverse" demand function since, as Bertrand pointed out, it is price that drives demand, not conversely.

[5]The second order condition reads $\frac{\partial^2}{\partial q_i^2}\mathrm{U}_i(q_i, Q_{-i}) = -2b < 0$.

Again, φ_i is called i's "reaction function." A Nash-Cournot equilibrium $\{q_i^*\}_{i=1,..,n}$ is reached when for all i: $q_i^* = \varphi_i(Q_{-i}^*)$. One easily obtains (for all i):

$$q_i^* = q^* = \frac{a}{b(n+1)} \tag{10}$$

with utility for each oligopolist i:

$$\mathrm{U}_i(q_i^*, Q_{-i}^*) = \frac{a^2}{b(n+1)^2} \tag{11}$$

Consider now the possibility that all n oligopolists get together in secret an try to engineer supply in order to raise prices. They proceed as follows: let us *all* commit to produce the same quantity $q_j \equiv r$ that would maximize all our identical utilities:

$$\mathrm{U}_i(r) = r(a - nbr)$$

One easily obtain the optimal such $r = \frac{a}{2nb}$ by writing the first order condition. But then, all i would enjoy utility:

$$\mathrm{U}_i(r) = \frac{a^2}{4nb} > \frac{a^2}{b(n+1)^2} \qquad \text{for any } n \geq 2 \tag{12}$$

This process known as "collusion" is forbidden by law in most countries but is sometimes hard to detect. Worse, there is the possibility that will be discussed when studying dynamic games of "tacit collusion" that would be undetectable.

3.5.3 A General Theorem

A Nash equilibrium in a continuous static game is simply a set of strategies, one per player, that cannot be individually improved upon by any player given the other players' own strategies in the set. There are several generalizations of Nash's Theorem to various types of games. The standard result for continuous games involves three main assumptions:

a) Continuity of payoffs in all decision variables;
b) Convexity and compactness of each player's decision space;[6]
c) Quasi-concavity of the players' own payoffs in their own decision variables.[7]

A standard theorem is:

Theorem: Assume that a continuous game has finitely many players each with decision variables that can only be chosen within a convex and compact subset of Euclidean space. Assume further that each player's payoff is continuous in all players' decision variables and is quasi-concave in each player's own decisions. Then the game admits a Nash equilibrium.

3.6 Homework

3.6.1 Rock-Paper-Scissors.

[6]Typically, a player's action space is a convex, closed and bounded subset of Euclidean space.
[7]A function $f(X)$ is quasi-concave if the set $Q(y) = \{X | f(X) \geq y\}$ is convex. In particular, a concave function is quasi-concave, but the converse is not true.

Write the payoff matrix of this well known game and obtain the single Nash equilibrium.

3.6.2 The Stag Hunt

The original story is attributed to Jean-Jacques Rousseau: three men can go hunting for rabbits individually, or they can join forces and hunt for a stag. Write a simple normal form representation of this typical social dilemma between individualism and cooperation and comment on your solutions. Generalize to the case where only two need to join forces to succeed in the stag hunt.

3.6.3 The Tragedy of the Commons

When more than two players are involved in a Prisoner's Dilemma like situation, one often refers to the so-called "Tragedy of the Commons," the idea that individualistic behavior can deplete common resources to the detriment of all, if left unchecked. Generalize the Prisoner's Dilemma of Chapter One to a 3-player symmetric game. Discuss the following two cases on how you attribute payoffs according to the number of defections: (1) when it is best to defect against two defectors; and (2) when it is worst to defect against two defectors. Solve for Nash equilibria in each case and comment.

3.6.4 Changing Order in Voting

a) Show that the probability that the appeal is granted is $\mathbb{P}(\text{Appeal Granted}) = \frac{1}{28}$ in the the mixed equilibrium of the simultaneous-vote appeal game of section 3.3.1.

b) Exchange the order of play for Gina and Ronald in the sequential game, solve and comment.

3.6.5 The Tenure Game

Stella Starr is an assistant professor at Big City State University (BCSU) where she is on the tenure track. She has been recently approached by Posh College located in Middle of Nowhere, a small college town, that has made her an offer with tenure. Her dilemma is that she would rather live in Big City but can't take the uncertainty about getting tenure at BCSU. So, she decides to go see her dean to ask for an early tenure review. The catch is that Posh College is not willing to wait forever and is considering other, perhaps less stellar candidates. So, Stella is facing a serious timing issue: if she waits too long to accept Posh's offer they might hire someone else before her early tenure review comes through. Meanwhile, BCSU would rather stick to the standard tenure clock but would not like to lose her. Posh officials are a bit worried that they are being played and that Stella is only using their offer as leverage at BCSU.

You are a game theorist colleague of Stella. What advice would you offer on the basis of a well developped game model?

3.6.6 The Bertrand Duopoly

Assuming that both sides will announce their optimal prices given by (3) and (4) they can now choose their respective locations a and b. Prove that that the optimal location for A is $a = 0$ and the optimal location for B is $b = 1$.

3.6.7 The Bertrand Oligopoly

Extend the Bertrand Duopoly to the case where a third players C located at c and charging price p_C is added, assuming $0 \le a < b < c \le 1$. Obtain formulas for the equilibrium prices.

3.6.8 The Dating Game

Gwen is a senior at her high school and she is admired by many of her classmates. Indeed, Gwen is smart and pretty, and she is one of the best athletes in the school. Yet, Gwen's unusual mix of qualities may be a serious handicap in the dating game for the forthcoming Prom night. Indeed, she so impresses her peers that none of her classmates has managed to foster the courage to ask her to be their date at the prom.

Gwen's best friend Linda is well aware of this paradoxical situation. So she decides to take action. She leaks to the class that Gwen once confided in her that she really likes a tall guy with green eyes and a talent for chess. Unfortunately, there are two boys in the class who fit that profile very well: Bill and Rob. Needless to say, both are crazy about Gwen. Unfortunately, they are also very sensitive guys and thus extremely afraid of being rejected.

Gwen in fact likes Bill and Rob equally. But she is a bit vain and would be most flattered if both boys decided to ask her, so that she could have her pick. In order to help sort things out, Linda who has heard that you are an expert in Game Theory asks *you* for your help. In Linda's thinking, Nature is going to choose randomly (and fairly) one the two boys to have the first shot at inviting Gwen. But if and when he makes his decision, each boy will be uncertain about whether:

a) He is the first mover. In other words, even if he is only second to move, he will not necessarily know it;
b) The other boy might have just chickened out; or
c) The other boy asked Gwen but she turned him down.

And at her turn Gwen will only know whom of the two boys is asking her. She won't know whether he is the first mover or if the other boy has been scared away.

You will design a game model using *GamePlan* to represent the above strategic issues and you will offer your best advice to Linda.