16

Game Theory and Cost Allocation

16.1 Introduction

In most decision-making situations, our profits (and losses) are determined not only by our decisions, but by the decisions taken by outside forces (e.g., our competitors, the weather, etc.). A useful classification is whether the outside force is indifferent or mischievous. We, for example, classify the weather as indifferent because its decision is indifferent to our actions, in spite of how we might feel during a rainstorm after washing the car and forgetting the umbrella. A competitor, however, generally takes into account the likelihood of our taking various decisions and as a result tends to make decisions that are mischievous relative to our welfare. In this chapter, we analyze situations involving a mischievous outside force. The standard terminology applied to the problem to be considered is *game theory*. Situations in which these problems might arise are in the choice of a marketing or price strategy, international affairs, military combat, and many negotiation situations. For example, the probability a competitor executes an oil embargo against us probably depends upon whether we have elected a strategy of building up a strategic petroleum reserve. Frequently, the essential part of the problem is deciding how two or more cooperating parties "split the pie". That is, allocate costs or profits of a joint project. For a thorough introduction to game theory, see Fudenberg and Tirole (1993).

16.2 Two-Person Games

In so-called two-person game theory, the key feature is *each of the two players must make a crucial decision ignorant of the other player's decision*. Only after both players have committed to their respective decisions does each player learn of the other player's decision and each player receives a payoff that depends solely on the two decisions. Two-person game theory is further classified according to whether the payoffs are constant sum or variable sum. In a constant sum game, the total payoff summed over both players is constant. Usually this constant is assumed to be zero, so one player's gain is exactly the other player's loss. The following example illustrates a constant sum game.

A game is to be played between two players called *Blue* and *Gold*. It is a single simultaneous move game. Each player must make her single move in ignorance of the other player's move. Both moves are then revealed and then one player pays the other an amount specified by the payoff table below:

Payoff from Blue to Gold

		Blue's Move	
		a	**b**
	a	4	–6
Gold's Move	**b**	–5	8
	c	3	–4

Blue must choose one of two moves, (*a*) or (*b*), while Gold has a choice among three moves, (*a*), (*b*), or (*c*). For example, if Gold chooses move (*b*) and Blue chooses move (*a*), then Gold pays Blue 5 million dollars. If Gold chooses (*c*) and Blue chooses (*a*), then Blue pays Gold 3 million dollars.

16.2.1 The Minimax Strategy

This game does not have an obvious strategy for either player. If Gold is tempted to make move (*b*) in the hopes of winning the 8 million dollar prize, then Blue will be equally tempted to make move (*a*), so as to win 5 million from Gold. For this example, it is clear each player will want to consider a random strategy. Any player who follows a pure strategy of always making the same move is easily beaten. Therefore, define:

BM_i = probability Blue makes move i, $i = a$ or b,
GM_i = probability Gold makes move i, $i = a$, b, or c.

How should Blue choose the probabilities BM_i? Blue might observe that:

If Gold chooses move (*a*), my expected loss is:
$4\ BMA - 6\ BMB$.

If Gold chooses move (*b*), my expected loss is:
$-5\ BMA + 8\ BMB$.

If Gold chooses move (c), my expected loss is:
$3\ BMA - 4\ BMB$.

So, there are three possible expected losses depending upon which decision is made by Gold. If Blue is conservative, a reasonable criterion is to choose the BM_i, so as to minimize the maximum expected loss. This philosophy is called the *minimax strategy*. Stated another way, Blue wants to choose the probabilities BM_i, so, no matter what Gold does, Blue's maximum expected loss is minimized. If *LB* is the maximum expected loss to Blue, the problem can be stated as the LP:

```
MIN = LB;
! Probabilities must sum to 1;
                BMA      + BMB   = 1;
! Expected loss if Gold chooses (a);
     -LB + 4 * BMA   - 6 * BMB <= 0;
! Expected loss if Gold chooses (b);
     -LB - 5 * BMA   + 8 * BMB <= 0;
! Expected loss if Gold chooses (c);
     -LB + 3 * BMA   - 4 * BMB <= 0;
```

The solution is:

```
Optimal solution found at step:         2
Objective value:                0.2000000

Variable           Value        Reduced Cost
      LB       0.2000000           0.0000000
     BMA       0.6000000           0.0000000
     BMB       0.4000000           0.0000000

     Row    Slack or Surplus      Dual Price
       1       0.2000000            1.000000
       2       0.0000000          -0.2000000
       3       0.2000000           0.0000000
       4       0.0000000           0.3500000
       5       0.0000000           0.6500000
```

The interpretation is, if Blue chooses move (*a*) with probability 0.6 and move (*b*) with probability 0.4, then Blue's expected loss is never greater than 0.2, regardless of Gold's move.

If Gold follows a similar argument, but phrases the argument in terms of maximizing the minimum expected profit, *PG*, instead of minimizing maximum loss, then Gold's problem is:

```
MAX = PG;
! Probabilities sum to 1;
          GMA     + GMB     + GMC  = 1;
! Expected profit if Blue chooses (a);
-PG + 4 * GMA - 5 * GMB + 3 * GMC >= 0;
! Expected profit if Blue chooses (b);
-PG - 6 * GMA + 8 * GMB - 4 * GMC >= 0;
```

The solution to Gold's problem is:

```
Optimal solution found at step:         1
Objective value:                0.2000000

Variable           Value        Reduced Cost
      PG       0.2000000           0.0000000
     GMA       0.0000000           0.1999999
     GMB       0.3500000           0.0000000
     GMC       0.6500000           0.0000000

     Row    Slack or Surplus      Dual Price
       1       0.2000000            1.000000
       2       0.0000000           0.2000000
       3       0.0000000          -0.6000000
       4       0.0000000          -0.4000000
```

The interpretation is, if Gold chooses move (*b*) with probability 0.35, move (*c*) with probability 0.65 and never move (*a*), then Gold's expected profit is never less than 0.2. Notice Gold's lowest expected profit equals Blue's highest expected loss. From Blue's point of view the expected transfer to Gold is at least 0.2. The only possible expected transfer is then 0.2. This means if both players follow the random strategies just derived, then on every play of the game there is an expected transfer of 0.2 units from Blue to Gold. The game is biased in Gold's favor at the rate of 0.2 million dollars per play. The strategy of randomly choosing among alternatives to keep the opponent guessing, is sometimes also known as a mixed strategy.

If you look closely at the solutions to Blue's LP and to Gold's LP, you will note a surprising similarity. The dual prices of Blue's LP equal the probabilities in Gold's LP and the negatives of Gold's dual prices equal the probabilities of Blue's LP. Looking more closely, you can note each LP is really the dual of the other one. This is always true for a two-person game of the type just considered and mathematicians have long been excited by this fact.

16.3 Two-Person Non-Constant Sum Games

There are many situations where the welfare, utility, or profit of one person depends not only on his decisions, but also on the decisions of others. A two-person game is a special case of the above in which:

1. there are exactly two players/decision makers,
2. each must make one decision,
3. in ignorance of the other's decision, and
4. the loss incurred by each is a function of both decisions.

A two-person constant sum game (frequently more narrowly called a zero sum game) is the special case of the above where:

(4a) the losses to both are in the same commodity (e.g., dollars) and
(4b) the total loss is a constant independent of players' decisions.

Thus, in a constant sum game the sole effect of the decisions is to determine how a "constant sized pie" is allocated. Ordinary linear programming can be used to solve two-person constant sum games.

When (1), (2) and (3) apply, but (4b) does not, then we have a two-person non-constant sum game. Ordinary linear programming cannot be used to solve these games. However, closely related algorithms, known as linear complementarity algorithms, are commonly applied. Sometimes a two-person non-constant sum game is also called a bimatrix game.

As an example, consider two firms, each of which is about to introduce an improved version of an already popular consumer product. The versions are very similar, so one firm's profit is very much affected by its own advertising decision as well as the decision of its competitor. The major decision for each firm is presumed to be simply the level of advertising. Suppose the losses (in millions of dollars) as a function of decision are given by Figure 16.1. The example illustrates that each player need not have exactly the same kinds of alternatives.

Figure 16.1 Two Person, Non-constant Sum Game

		Firm A		
		No Advertise	Advertise Medium	Advertise High
Firm B	No Advertise	-4 -4	-3 -2	-5 1
	Advertise	-1 -5	-2 -1	-1 0

Negative losses correspond to profits.

16.3.1 Prisoner's Dilemma

This cost matrix has the so-called prisoner's dilemma cost structure. This name arises from a setting in which two accomplices in crime find themselves in separate jail cells. If neither prisoner cooperates with the authorities (thus the two cooperate), both will receive a medium punishment. If one of them provides evidence against the other, the other will get severe punishment while the one who provides evidence will get light punishment, *if the other does not provide evidence against the first.* If each provides evidence against the other, they both receive severe punishment. Clearly, the best thing for the two as a group is for the two to cooperate with each other. However, individually there is a strong temptation to defect.

The prisoner's dilemma is common in practice, especially in advertising. The only way of getting to Mackinac Island in northern Michigan is via ferry from Mackinaw City. Three different companies, Sheplers, the Arnold Line, and the Star Line operate such ferries. As you approach Mackinaw City by car, you may notice up to a mile before the ferry landing, that each company has one or more small roadside stands offering to sell ferry tickets for their line. Frequent users of the ferry service proceed directly to the well-marked dock area and buy a ticket after parking the car and just before boarding the ferry (no cars are allowed on Mackinac Island). No reserved seats are sold, so there is no advantage to buying the tickets in advance at the stands. First time visitors, however, are tempted to buy a ticket at a company specific stand because the signs suggest that this is the safe thing to do. The "socially" most efficient arrangement would be to have no advanced ticket booths. If a company does not have a stand, however, while its competitors do, then this company will lose a significant fraction of the first time visitor market.

The same situation exists with the two firms in our little numerical example. For example, if *A* does not advertise, but *B* does, then *A* makes 1 million and *B* makes 5 million of profit. Total profit would be maximized if neither advertised. However, if either knew the other would not advertise, then the one who thought he had such clairvoyance would have a temptation to advertise.

Later, it will be useful to have a loss table with all entries strictly positive. The relative attractiveness of an alternative is not affected if the same constant is added to all entries. Figure 16.2 was obtained by adding +6 to every entry in Figure 16.1:

Figure 16.2 Two Person, Non-constant Sum Game

		Firm A		
		No Advertise	Advertise Medium	Advertise High
Firm B	No Advertise	2 / 2	3 / 4	1 / 7
	Advertise	5 / 1	4 / 5	5 / 6

We will henceforth work with the data in Figure 16.2.

16.3.2 Choosing a Strategy

Our example illustrates that we might wish our own choice to be:

i. somewhat unpredictable by our competitor, and
ii. robust in the sense that, regardless of how unpredictable our competitor is, our expected profit is high.

Thus, we are lead (again) to the idea of a random or mixed strategy. By making our decision random (e.g., by flipping a coin) we tend to satisfy (i). By biasing the coin appropriately, we tend to satisfy (ii).

For our example, define a_1, a_2, a_3 as the probability *A* chooses the alternative "No advertise", "Advertise Medium", and "Advertise High", respectively. Similarly, b_1 and b_2 are the probabilities that *B* applies to alternatives "No Advertise" and "Advertise", respectively. How should firm *A* choose a_1, a_2, and a_3? How should firm *B* choose b_1 and b_2?

For a bimatrix game, it is difficult to define a solution that is simultaneously optimum for both. We can, however, define an equilibrium stable set of strategies. A stable solution has the feature that, given *B*'s choice for b_1 and b_2, *A* is not motivated to change his probabilities a_1, a_2, and a_3. Likewise, given a_1, a_2, and a_3, *B* is not motivated to change b_1 and b_2. Such a solution, where no player is motivated to unilaterally change his or her strategy, is sometimes also known as a Nash equilibrium. There may be bimatrix games with several stable solutions.

What can we say beforehand about a strategy of A's that is stable? Some of the a_i's may be zero while for others we may have $a_i > 0$. An important observation which is not immediately obvious is the following: the expected loss to A of choosing alternative i is the same over all i for which $a_i > 0$. If this were not true, then A could reduce his overall expected loss by increasing the probability associated with the lower loss alternative. Denote the expected loss to A by v_A. Also, the fact that $a_i = 0$ must imply the expected loss from choosing i is > v_A. These observations imply that, with regard to A's behavior, we must have:

$2b_1 + 5b_2 \geq v_A$ (with equality if $a_1 > 0$),
$3b_1 + 4b_2 \geq v_A$ (with equality if $a_2 > 0$),
$b_1 + 5b_2 \geq v_A$ (with equality if $a_3 > 0$).

Symmetric arguments for B imply:

$2a_1 + 4a_2 + 7a_3 \geq v_B$ (with equality if $b_1 > 0$),
$a_1 + 5a_2 + 6a_3 \geq v_B$ (with equality if $b_2 > 0$).

We also have the nonnegativity constraints:

$a_i \geq 0$ and $b_i \geq 0$, for all alternatives i.

Because the a_i and b_i are probabilities, we wish to add the constraints $a_1 + a_2 + a_3 = 1$ and $b_1 + b_2 = 1$.

If we explicitly add slack (or surplus if you wish) variables, we can write:

$2b_1 + 5b_2 - slka1 = v_A$
$3b_1 + 4b_2 - slka2 = v_A$
$b_1 + 5b_2 - slka3 = v_A$
$2a_1 + 4a_2 + 7a_3 - slkb1 = v_B$
$a_1 + 5a_2 + 6a_3 - slkb2 = v_B$
$a_1 + a_2 + a_3 = 1$
$b_1 + b_2 = 1$
$a_i \geq 0$, $b_i \geq 0$, $slka_i \geq 0$, and $slkb_i \geq 0$, for all alternatives i.
$slka1 * a1 = 0$
$slka2 * a2 = 0$
$slka3 * a3 = 0$
$slkb1 * b1 = 0$
$slkb2 * b2 = 0$

The last five constraints are known as the complementarity conditions. The entire model is known as a linear complementarity problem.

Rather than use a specialized linear complementarity algorithm, we will simply use the integer programming capabilities for LINGO to model the problem as follows:

```
MODEL: ! Two person nonconstant sum game.(BIMATRX);
 SETS:
   OPTA: PA, SLKA, NOTUA, COSA;
   OPTB: PB, SLKB, NOTUB, COSB;
   BXA( OPTB, OPTA): C2A, C2B;
 ENDSETS
 DATA:
 OPTB = BNAD BYAD;
 OPTA = ANAD AMAD AHAD;
  C2A =  2    3    1  ! C2A( I, J) = cost to A if B;
         5    4    5; ! chooses row I, A chooses col J;
  C2B =  2    4    7  ! C2B( I, J) = cost to B if B;
         1    5    6; ! chooses row I, A chooses col J;
 ENDDATA
!--------------------------------------------------;
! Conditions for A, for each option J;
 @FOR( OPTA( J):
! Set CBSTA= cost of strategy J, if J is used by A;
   CBSTA = COSA( J) - SLKA( J);
   COSA( J) = @SUM( OPTB( I): C2A( I, J) * PB( I));
! Force SLKA( J) = 0 if strategy J is used;
   SLKA( J) <= NOTUA( J) * @MAX( OPTB( I):
     C2A( I, J));
! NOTUA( J) = 1 if strategy J is not used;
   PA( J) <= 1 - NOTUA( J);
! Either strategy J is used or it is not used;
   @BIN( NOTUA( J));
       );
! A must make a decision;
 @SUM( OPTA( J): PA( J)) = 1;
! Conditions for B;
 @FOR( OPTB( I):
! Set CBSTB = cost of strategy I, if I is used by
      B;
   CBSTB =  COSB( I)  - SLKB( I);
   COSB( I) = @SUM( OPTA( J): C2B( I, J) * PA( J));
! Force SLKB( I) = 0 if strategy I is used;
   SLKB( I) <= NOTUB( I) * @MAX( OPTA( J):
      C2B( I, J));
! NOTUB( I) = 1 if strategy I is not used;
   PB( I) <= 1 - NOTUB( I);
! Either strategy I is used or it is not used;
   @BIN( NOTUB( I));
       );
!  B must make a decision;
 @SUM( OPTB( I): PB( I)) = 1;
  END
```

A solution is:

```
             Variable           Value
                CBSTA        3.666667
                CBSTB        5.500000
            PA( AMAD)       0.5000000
            PA( AHAD)       0.5000000
          SLKA( ANAD)       0.3333333
         NOTUA( ANAD)        1.000000
          COSA( ANAD)        4.000000
          COSA( AMAD)        3.666667
          COSA( AHAD)        3.666667
            PB( BNAD)       0.3333333
            PB( BYAD)       0.6666667
          COSB( BNAD)        5.500000
          COSB( BYAD)        5.500000
```

The solution indicates that firm *A* should not use option 1(No ads) and should randomly choose with equal probability between options 2 and 3. Firm *B* should choose its option 2(Advertise) twice as frequently as it chooses its option 1(Do not advertise).

The objective function value, reduced costs and dual prices can be disregarded. Using our original loss table, we can calculate the following:

Situation			Weighted Contribution To Total Loss of	
A	**B**	**Probability**	**A**	**B**
No Ads	No Ads	0 × 1/3	0	0
No Ads	Ads	0 × 2/3	0	0
Advertise Medium	No Ads	1/2 × 1/3	(1/6) × (−3)	(1/6) × (−2)
Advertise Medium	Ads	1/2 × 2/3	(1/3) × (−2)	(1/3) × (−1)
Advertise High	No Ads	1/2 × 1/3	(1/6) × (−5)	(1/6) × (1)
Advertise High	Ads	1/2 × 2/3	(1/3) × (−1)	(1/3) × (0)
			−2.3333	−0.5

Thus, following the randomized strategy suggested, *A* would have an expected profit of 2.33 million; whereas, *B* would have an expected profit of 0.5 million. Contrast this with the fact that, if *A* and *B* cooperated, they could each have an expected profit of 4 million.

16.3.3 Bimatrix Games with Several Solutions

When a nonconstant sum game has multiple or alternative stable solutions, life gets more complicated. The essential observation is we must look outside our narrow definition of “stable solution” to decide which of the stable solutions, if any, would be selected in reality.

Consider the following nonconstant sum two-person game:

Figure 16.3 Bimatrix Games

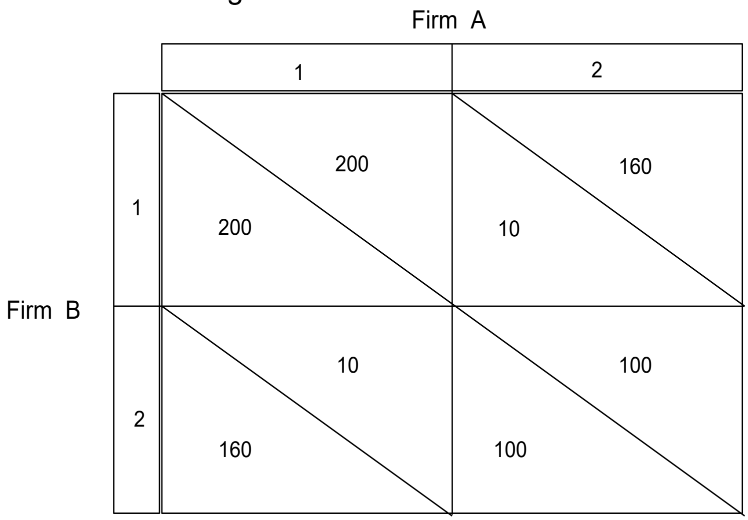

As before, the numbers represent losses.

First, observe the one solution that is stable according to our definition: (*I*) Firm *A* always chooses option 1 and Firm *B* always chooses option 2. Firm *A* is not motivated to switch to 2 because its losses would increase to 100 from 10. Similarly, *B* would not switch to 1 from 2 because its losses would increase to 200 from 160. The game is symmetric in the players, so similar arguments apply to the solution (*II*): *B* always chooses 1 and *A* always chooses 2.

Which solution would result in reality? It probably depends upon such things as the relative wealth of the two firms. Suppose:

i. *A* is the wealthier firm,
ii. the game is repeated week after week, and
iii. currently the firms are using solution *II*.

After some very elementary analysis, *A* concludes it much prefers solution *I*. To move things in this direction, *A* switches to option 1. Now, it becomes what applied mathematicians call a game of "chicken". Both players are taking punishment at the rate of 200 per week. Either player could improve its lot by $200 - 160 = 40$ by unilaterally switching to its option 2. However, its lot would be improved a lot more (i.e., $200 - 10 = 190$) if its opponent unilaterally switched. At this point, a rational *B* would probably take a glance at *A*'s balance sheet and decide *B* switching to option 2 is not such a bad decision. When a game theory problem has multiple solutions, any given player would like to choose that stable solution which is best for it. If the player has the wherewithal to force such a solution (e.g., because of its financial size), then this solution is sometimes called a Stackelberg equilibrium.

If it is not clear which firm is wealthier, then the two firms may decide a cooperative solution is best (e.g., alternate between solutions *I* and *II* in alternate weeks). At this point, however, federal antitrust authorities might express a keen interest in this bimatrix game.

We conclude a "stable" solution is stable only in a local sense. When there are multiple stable solutions, we should really look at all of them and take into account other considerations in addition to the loss matrix.

The above two-player non-cooperative game analysis involved only two players. It can be extended to three or more players, however, the number of variables and constraints increases multiplicatively. For three players you will need three cubes rather than two matrices in order to describe the payoffs to a given player *X*, given that *X* chose alternative *i*, and player *Y* chose alternative *j*, and player *Z* chose alternative *k*.

16.4 Nonconstant-Sum Cooperative Games with > 2 Players

The most unrealistic assumption underlying classical two-person constant-sum game theory is the sum of the payoffs to all players must sum to zero (actually a constant, without loss of generality). In reality, the total benefits are almost never constant. Usually, total benefits increase if the players cooperate, so these situations are sometimes called cooperative games. In these nonconstant-sum games, the difficulty then becomes one of deciding how these additional benefits due to cooperation should be distributed among the players.

There are two styles for analyzing nonconstant sum games. If we restrict ourselves to two persons, then so-called bimatrix game theory extends the methods for two-person constant sum games to nonconstant sum games. If there are three or more players, then *n-person game theory* can be used in selecting a decision strategy. The following example illustrates the essential concepts of *n*-person game theory.

Three property owners, *A*, *B*, and *C*, own adjacent lakefront property on a large lake. A piece of property on a large lake has higher value if it is protected from wave action by a seawall. *A*, *B*, and *C* are each considering building a seawall on their properties. A seawall is cheaper to build on a given piece of property if either or both of the neighbors have seawalls. For our example, *A* and *C* already have expensive buildings on their properties. *B* does not have buildings and separates *A* from *C* (i.e., *B* is between *A* and *C)*. The net benefits of a seawall for the three owners are summarized below:

Owners Who Cooperate, i.e., Build While Others Do Not	**Net Benefit to Cooperating Owners**
A alone	1.2
B alone	0
C alone	1
A and *B*	4
A and *C*	3
B and *C*	4
A, *B*, and *C*	7

Obviously, all three owners should cooperate and build a unified seawall because then their total benefits will be maximized. It appears *B* should be compensated in some manner because he has no motivation to build a seawall by himself. Linear programming can provide some help in selecting an acceptable allocation of benefits.

Denote by v_A, v_B, and v_C the net benefits, which are to be allocated to owners A, B, and C. No owner or set of owners will accept an allocation that is less than that, which they would enjoy if they acted alone. Thus, we can conclude:

$$v_A \geq 1.2$$
$$v_B \geq 0$$
$$v_C \geq 1$$
$$v_A + v_B \geq 4$$
$$v_A + v_C \geq 3$$
$$v_B + v_C \geq 4$$
$$v_A + v_B + v_C \leq 7$$

That is, any allocation satisfying the above constraints should be self-enforcing. No owner would be motivated to not cooperate. He cannot do better by himself. The above constraints describe what is called the "core" of the game. Any solution (e.g., $v_A = 3$, $v_B = 1$, $v_C = 3$) satisfying these constraints is said to be in the core.

Various objective functions might be appended to this set of constraints to give an LP. The objective could take into account secondary considerations. For example, we might choose to maximize the minimum benefit. The LP in this case is:

Maximize z

subject to $z \leq v_A$; $z \leq v_B$; $z \leq v_C$

$$v_A \geq 1.2$$
$$v_C \geq 1$$
$$v_A + v_B \geq 4$$
$$v_A + v_C \geq 3$$
$$v_A + v_B + v_C \leq 7$$

A solution is $v_A = v_B = v_C = 2.3333$.

Note the core can be empty. That is, there is no feasible solution. This would be true, for example, if the value of the coalition A, B, C was 5.4 rather than 7. This situation is rather interesting. Total benefits are maximized by everyone cooperating. However, total cooperation is inherently unstable when benefits are 5.4. There will always be a pair of players who find it advantageous to form a subcoalition and improve their benefits (at the considerable expense of the player left out). As an example, suppose the allocations to A, B, and C under full cooperation are 1.2, 2.1, and 2.1, respectively. At this point, A would suggest to B the two of them exclude C and cooperate between the two of them. A would suggest to B the allocation of 1.8, 2.2, and 1. This is consistent with the fact that A and B can achieve a total of 4 when cooperating. At this point, C might suggest to A that the two of them cooperate and thereby select an allocation of 1.9, 0, 1.1. This is inconsistent with the fact that A and C can achieve a total of 3 when cooperating. At this point, B suggests to C etc. Thus, when the core is empty, it may be everyone agrees that full cooperation can be better for everyone. There nevertheless must be an enforcement mechanism to prevent "greedy" members from pulling out of the coalition.

16.4.1 Shapley Value

Another popular allocation method for cooperative games is the Shapley Value. The rule for the Shapley Value allocation is that each player should be awarded his average marginal contribution to the coalition if one considers all possible sequences for forming the full coalition. The following table illustrates for the previous example:

	Marginal value of player		
Sequence	A	B	C
A B C	1.2	2.8	3
A C B	1.2	4	1.8
B A C	4	0	3
B C A	3	0	4
C A B	2	4	1
C B A	3	3	1
Total:	14.4	13.8	13.8
Average:	2.4	2.3	2.3

Thus, the Shapley value allocates slightly more, 2.4, to Player *A* in our example. For this example, as with most typical cooperative games, the Shapley Value allocation is in the core if the core is non-empty.

16.5 The Stable Marriage/Assignment Problem

The stable marriage problem is the multi-person interpretation of the assignment problem. Although the major application of the stable marriage model is in college admissions and large scale labor markets, the problem historically has been explained as the "marriage" problem of assigning each of *n* men to exactly one of *n* women, and vice versa. Instead of there being a single objective function, each man provides a preference ranking of each of the women, and each woman provides a ranking of each of the men. An assignment is said to be stable if for every man *i*, and woman *j*, either: a) man *i* prefers the woman he currently is assigned to over woman *j*, or b) woman *j* prefers the man she is currently assigned to over man *i*. Otherwise, man *i* and woman *j* would be motivated to abandon their current partners and "elope". The stable marriage assignment method has been used for assigning medical residents and interns to hospitals in the U.S. since 1952. Each year, thousands of prospective interns rank each of the hospitals in which they are interested, and each hospital ranks each of the interns in which they are interested. Then a neutral agency, the National Resident Matching Program (NRMP) assigns interns to hospitals using methods described below. A similar system is used in Canada and Scotland. Norway and Singapore use a similar approach to assign students to schools and universities. Roth (1984) gives a very interesting history of how the U.S. medical profession came to use the stable marriage assignment method embodied in NRMP. Roth, Sonmez, and Unver(2005) describe the establishment of a system, based on the marriage assignment method, for matching kidney donors with people needing kidneys.

In any multi-player problem, the following questions should always be asked: a) Is there always a stable assignment or more generally an equilibrium solution? b) Can there be multiple stable solutions?
c) If yes, what criterion should we use for choosing among the multiple solutions? d) Is the solution Pareto optimal, i.e., undominated? e) Is our method for solving the problem, in particular how we answer (c), incentive compatible? That is, does the method motivate the players to provide accurate input information, e.g., rankings, for our method?

We illustrate ideas with the following 3-man, 3-woman example from Gale and Shapley(1962). A 1 means most attractive. A 3 means least attractive.

```
 MAN = ADAM  BOB  CHUCK;
WOMAN= ALICE BARB CARMEN;
! Men(row) preference for women(col);
MPREF =
         1    2    3  !ADAM;
         3    1    2  !BOB;
         2    3    1; !CHUCK;
! Women(col) preference for men(row);
WPREF =
         3    2    1 !ADAM;
         1    3    2 !BOB;
         2    1    3;!CHUCK;
! Thus, Adam's first choice is Alice.
     Alice's first choice is Bob;
```

We shall see from this example that the answer to question (b) is that, yes, there can be multiple stable solutions. In this example, giving each man his first choice (and incidentally, each woman her third choice) is feasible, giving the assignment: Adam with Alice, Bob with Barb, and Chuck with Carmen. It is stable because no man is motivated to switch. A second stable solution is possible by giving each woman her first choice (and incidentally, each man his third choice), namely: Adam with Carmen, Bob with Alice, and Chuck with Barb. It is stable because no woman is motivated to switch. A third, less obvious stable solution is to give everyone their second choice: Adam with Barb, Bob with Carmen, and Chuck with Alice. All other assignments are unstable.

How to solve the problem?

Gale and Shapley (1962) show that an intuitive iterative courtship type of method can be made into a rigorous algorithm for finding a stable assignment. The algorithm proceeds as follows:

1) Each man proposes, or is tentatively assigned, to his first choice woman.
2) If every woman has exactly one man assigned to her, then stop. We have a stable assignment.
3) Else, each woman who has two or more men assigned to her rejects all but one of the men assigned to her, tentatively keeping the one most attractive to her of the men that just proposed to her.
4) Each man just rejected in (3) proposes/is assigned to the next most attractive woman on his list.
5) Go to (2).

This version of the algorithm will produce the first solution mentioned above in which all men get there first choice. Obviously, there is the female version of this algorithm in which the roles of men and woman or exchanged. That version gives the second solution above. Gale and Shapley(1962) make the following observations: i) Regarding our question (a) above, this algorithm will always find a stable solution; ii) If both the male and the female versions of the algorithm give the same assignment, then that is the unique stable solution; iii) When men propose first, the solution is optimal for the men in the sense that there is no other stable solution in which any man does better. Similarly, the version in which women propose first, results in a solution that is optimal for women.

The two Gale/Shapley algorithms can only give a solution in which men are treated very well, or a solution in which women are treated very well. What about a solution in which everyone is treated "moderately well"? Vande Vate(1989) showed that it is possible to formulate the stable marriage assignment problem as a linear program. The key observation is: if we consider any man i and woman j in a stable solution, then one of the following must hold: a) i and j are assigned to each other, or b) man i is assigned

to some other woman k whom he prefers to j, or c) woman j is assigned to some man k whom she prefers to i. If none of (a), (b), and (c) hold, then i and j both prefer each other to their current mates and they are tempted to elope.
Define the parameters and sets:

$mpref_{ij}$ = the preference position of woman j for man i,
e.g., if man 2's first choice is woman 3, then $mpref_{23} = 1$,

$wpref_{ij}$ = the preference position of man i for woman j,
e.g., if woman 3's second choice is man 1, then $wpref_{13} = 2$,

$SM(i,j)$ = the set of women that man i prefers to woman j,

$= \{ k : mpref_{ik} < mpref_{ij} \}$,

$SW(i,j)$ = the set of men that woman j prefers to man i,

$= \{ k : mpref_{kj} < mpref_{ij} \}$,

Define the variables:

y_{ij} = 1 if man i and woman j are assigned to each other.

The "no eloping" stability conditions (a), (b), and (c) above correspond to the linear constraints:

For all men i and women j:

$y_{ij} + \Sigma_{k \text{ in } SM(i,j)}\ y_{ik} + \Sigma_{k \text{ in } SW(k,j)}\ y_{kj}$.

A remaining question is, what objective function should we use? We already saw a solution above in which men were treated well but women were treated poorly, and a solution in which women were treated well but men were treated poorly. How about a solution in which minimizes the worst that anyone gets treated? The following LINGO model illustrates.

```
! Stable Marriage Assignment(stable_marriage3);
SETS:
   MAN: AM;
 WOMAN: AW;
 MXW(MAN,WOMAN): MPREF, WPREF, Y, RM, RW;
ENDSETS
DATA:
! Example from Gale and Shapley(1962);
  MAN = ADAM  BOB  CHUCK;
 WOMAN= ALICE BARB CARMEN;
 ! Men(row) preference for women(col);
 MPREF =
          1    2    3 !ADAM;
          3    1    2 !BOB;
          2    3    1;!CHUCK;
 ! Women(col) preference for men(row);
 WPREF =
          3    2    1 !ADAM;
          1    3    2 !BOB;
          2    1    3;!CHUCK;
```

```
! Thus, Adam's first choice is Alice.
     Alice's first choice is Bob;
! This data set has 3 stable assignments;
ENDDATA
! Y(i,j) = 1 if man i is assigned to woman j;

! Each man must be assigned;
  @FOR(MAN(i):
    @SUM(WOMAN(j): Y(i,j)) = 1;
      );
! Each woman must be assigned;
  @FOR(WOMAN(j):
    @SUM(MAN(i): Y(i,j)) = 1;
      );

! Stability conditions: Either man i and woman are
   assigned to each other, or
   man i gets a woman k he prefers to j, or
   woman j, gets a man k she prefers to i;
  @FOR( MXW(i,j):
     Y(i,j)
   + @SUM(WOMAN(k)| MPREF(i,k) #LT# MPREF(i,j): Y(i,k))
   + @SUM(  MAN(k)| WPREF(k,j) #LT# WPREF(i,j): Y(k,j)) >= 1
        );

! Compute actual assigned rank for each man and woman;
@FOR( MAN(i):
  AM(i) = @SUM( WOMAN(k): MPREF(i,k)*Y(i,k));
  PWORST >= AM(i);
    );
@FOR(WOMAN(j):
  AW(j) = @SUM(   MAN(k): WPREF(k,j)*Y(k,j));
  PWORST >= AW(j);
   );

! Minimize the worst given to anyone;
  MIN = PWORST;
```

When solved, we get the solution:

```
          Variable           Value
     Y( ADAM, BARB)        1.000000
    Y( BOB, CARMEN)        1.000000
   Y( CHUCK, ALICE)        1.000000
```

In the “Men first” solution, every woman got her third choice. In the “Woman first” solution, every man got his third choice. In this solution, the worst anyone gets is their second choice. In fact, everyone gets their second choice. McVitie and Wilson(1971) present an algorithm for efficiently enumerating all stable solutions.

For this example, we have an answer to question (d) above. It is easy to see that each of the three solutions is Pareto optimal. In the “Women first” solution, clearly the women cannot do any better, and the men cannot do any better without hurting one of the women. Similar comments apply to the other two solutions.

With regard to the incentive compatibility question, (e) above, Roth, Rothblum, and Vande Vate provide a partial answer, namely, if the "Men first" algorithm is used then there is nothing to be gained by a man misrepresenting his preferences. This is somewhat intuitive in that if the "Men first" rule is used, then the resulting solution gives each man the best solution possible among all stable solutions. We may reasonably restrict ourselves to stable solutions. Thus, if some man misrepresents his preferences, this might cause a different stable solution to result in which this man might be treated worse, but definitely no better. Abdulkadiroglu, Pathak, and Roth(2005), mention that New York City, when assigning students to highschools, uses a "Students first" variant of the marriage assignment algorithm so as to motivate students to state their true preferences among highschools they are considering attending.

16.5.1 The Stable Room-mate Matching Problem

The stable room-mate problem is the multi-person interpretation of the 2-matching optimization problem. A college wants to match incoming freshman, two to a room in a freshman dormitory. Each student provides a ranking of all other potential room-mates. A matching is stable if there are no two students, *i* and *j*, who are not room-mates such that *i* prefers *j* to his current room-mate, and *j* prefers *i* to his current room-mate. The stable marriage problem can be interpreted as a special case of the room-mate matching problem in which people give very unattractive rankings to people of the same sex.

In contrast to the stable marriage problem, there need not be a stable solution to a stable room-mate problem. The following 4-person example due to Gale and Shapley(1962) illustrates a situation with no stable matching.

```
! Example from Gale and Shapley;
  PERSON = AL   BOB  CAL DON;
! Row preference for col;
  PREF =
      99  1  2  3
       2 99  1  3
       1  2 99  3
       1  2  3 99;
 ! E.g., AL
 ! The 99 is to indicate that a person cannot be
   matched to himself.
```

Consider, for example, the solution: AL with BOB, and CAL with DON. It is not stable because BOB is matched with his second choice and CAL is matched with his third choice, whereas if BOB and CAL got together, BOB would get his first choice and CAL would get his second choice. That would give us the solution BOB with CAL, and AL with DON. This is solution is not stable, however, because then AL and CAL would discover that they could improve their lot by getting together to give: AL with CAL, and BOB with DON. This solution is not stable, etc. In the terminology of game theory, the marriage assignment problem always has a core. The room-mate matching problem may not have a core.

Irving(1985) gives an efficient algorithm for detecting whether a room-mates problem has a stable matching, and if yes, finding a stable matching. The room-mates problem can also be solved by formulating it as a mathematical program as illustrated by the following LINGO model for finding a stable room-mate matching among 8 potential room-mates. This example from Irving(1985) has three stable matchings.

```
! Stable Roommate Matching(stable_roommate8);
! Each of 2n people specify a rank, 1, 2,..., 2n-1, for
  each other person.  We want to pair up the people into
  a stable set of pairs, i.e., there are no two people
  i and j who are not paired up, but would prefer to be
  paired up rather than be paired with their current partner.
  It may be that there is no such a stable pairing.  This
  LINGO model will find such a pairing if one exists, and
  will minimize the worst that any person gets treated under
  this pairing.
SETS:
 PERSON: AP;
 PXP(PERSON,PERSON): PREF, Y, R, NOSTAB;
ENDSETS
DATA:
! Example from Irving(1985);
  PERSON = 1..8;
! Row preference for col;
 PREF=!1  2  3  4  5  6  7  8;
     99  1  7  3  2  4  5  6
      3 99  1  7  6  2  4  5
      7  3 99  1  5  6  2  4
      1  7  3 99  4  5  6  2
      2  4  5  6 99  1  7  3
      6  2  4  5  3 99  1  7
      5  6  2  4  7  3 99  1
      4  5  6  2  1  6  3 99;
 ! E.g., the first choice of 1 is 2.  The first choice
   of 8 is 5.
 ! The 99 is to indicate that a person cannot be
   matched to himself.
 ! This data set has 3 stable matchings;
ENDDATA

! Y(i,j) = 1 if PERSON i and j are matched, for  i < j;

  NP = @SIZE(PERSON);
! Each person must be assigned;
  @FOR(PERSON(i):
    @SUM(PERSON(k)| k #GT# i: Y(i,k))
  + @SUM(PERSON(k)| k #LT# i: Y(k,i)) = 1;
      );

! Turn off the lower diagonal part of Y;
   @SUM( PXP(i,j)| i #GT# j: Y(i,j)) = 0;

! Enforce monogamy by making the Y(i,j) = 0 or 1;
   @FOR( PXP(i,j):
       @BIN(Y(i,j))
        );

! Stability conditions: Either person i and person j
   are assigned to each other, or
```

```
     person i gets a person k he prefers to j, or
     person j gets a person k he prefers to i, or
     there is no stable solution;
    @FOR( PXP(i,j)| i #LT# j:
       Y(i,j)
      +@SUM(PERSON(k)| k #LT# i #AND# PREF(i,k) #LT# PREF(i,j): Y(k,i))
      +@SUM(PERSON(k)| k #GT# i #AND# PREF(i,k) #LT# PREF(i,j): Y(i,k))
      +@SUM(PERSON(k)| k #LT# j #AND# PREF(j,k) #LT# PREF(j,i): Y(k,j))
      +@SUM(PERSON(k)| k #GT# j #AND# PREF(j,k) #LT# PREF(j,i): Y(j,k))
      + NOSTAB(i,j) >= 1
          );

  ! Compute actual assigned rank for each person;
  @FOR( PERSON(i):
    AP(i) = @SUM( PERSON(k)| i #LT# k: PREF(i,k)*Y(i,k))
          + @SUM( PERSON(k)| k #LT# i: PREF(i,k)*Y(k,i));
    PWORST >= AP(i);
      );

  ! Compute number of instabilities;
    NUMUSTAB = @SUM(PXP(i,j): NOSTAB(i,j));
  ! Apply most weight to getting a stable solution;
    MIN = NP*NP*NUMUSTAB + PWORST;
```

Notice in the resulting solution below, there is a stable matching, i.e. NUMUSTAB = 0, and, no participant received worse than his second choice.

```
Variable           Value
NUMUSTAB        0.000000
Y( 1, 5)        1.000000
Y( 2, 6)        1.000000
Y( 3, 7)        1.000000
Y( 4, 8)        1.000000
```

16.6 Should We Behave Non-Optimally to Obtain Information?

One of the arts of modeling is knowing which details to leave out of the model. Unfortunately, the most likely details left out of a model are the things that are difficult to quantify. One kind of difficult-to-quantify feature is the value of information. There are a number of situations where, if value of information is considered, then one may wish to behave non-optimally, at least in the short run. Three situations to consider are: 1) We would like to gain information about a customer or supplier, e.g., a more precise description of the customer's demand curve or credit-worthiness, 2) We do not want to communicate too much information to a competitor, or 3) We want to communicate information to a business partner, e.g., a supplier.

As an example of (1) suppose we extend credit to some customers. If our initial credit optimization model says "never extend credit to customers with profile X", then we may nevertheless wish to occasionally extend credit to such customers in order to have up-to-date information of the credit worthiness of customers with profile X. In the inventory setting where unsatisfied demand is lost and not observed, Ding and Puterman(2002) suggest that it may be worthwhile to stock a little more than "optimal" so as to get a better estimate of customer demand.

Regarding (2), we may wish to behave non-optimally so as to not reveal too much information. Any good poker player knows that one should occasionally bluff by placing a large bet, even though the odds associated with the current hand do not justify a large bet. If other players know you never bluff, then they will drop out early and not give you the chance of winning large bets, any time you make a large bet. Similarly, there was a rumor at the end of World War II that Britain allowed a bombing attack on Coventry on one occasion even though Britain knew in advance of the attack, thanks to its code-breaking. The argument was that if Britain had sent up a large fleet of fighter in advance to meet the incoming German bombers, the Germans would have known, earlier than Britain desired that Britain had broken the German communications code.

An example of (3) comes from inventory control. An optimal inventory model may recommend using a very large order size. If we use a smaller order size, however, we will be giving more timely information to our supplier about retail demand for his product. In between orders, the supplier has no additional information about how his product is selling. In the extreme, if we used an order size of 1, then the supplier would have very up-to-date information about retail demand and could do better planning.

In probability theory there is a problem class known as the multi-armed bandit problem that is similar to case (1). A decision maker (DM) must decide which one of several slot machines (one armed bandits) should be selected for the next bet. The DM strongly suspects that the expected payoff is different for different machines. From a simple pure optimization perspective, the DM would bet only on the machine with the highest expected payoff. From an information perspective, however, the DM wants to scatter the bets a little bit in order to better estimate the expected payoff of each machine. This trade-off between optimization vs. experimentation is sometimes called the explore vs. exploit decision.

16.7 Problems

1. Both Big Blue, Inc. and Golden Apple, Inc. are "market oriented" companies and feel market share is everything. The two of them have 100% of the market for a certain industrial product. Blue and Gold are now planning the marketing campaigns for the upcoming selling season. Each company has three alternative marketing strategies available for the season. Gold's market share as a function of both the Blue and Gold decisions are tabulated below:

Payment To Blue by Gold as a Function of Both Decisions

		Blue Decision		
		A	**B**	**C**
	X	.4	.8	.6
Gold Decision	**Y**	.3	.7	.4
	Z	.5	.9	.5

Both Blue and Gold know the above matrix applies. Each must make their decision before learning the decision of the other. There are no other considerations.

a) What decision do you recommend for Gold?

b) What decision do you recommend for Blue?

2. Formulate an LP for finding the optimal policies for Blue and Gold when confronted with the following game:

Payment To Blue By Gold as a Function of Both Decisions

		Blue Decision			
		A	**B**	**C**	**D**
Gold Decision	**X**	2	–2	1	6
	Y	–1	4	5	–1

3. Two competing manufacturing firms are contemplating their advertising options for the upcoming season. The profits for each firm as a function of the actions of both firms are shown below. Both firms know this table:

Profit Contributions

		Fulcher Fasteners		
		Option A	**Option B**	**Option C**
Repicky Rivets	**Option Y**	4 10	8 4	6 6
	Option X	8 8	12 2	10 4

a) Which pair of actions is most profitable for the pair?

b) Which pairs of actions are stable?

c) Presuming side payments are legal, how much would which firm have to pay the other firm in order to convince them to stick with the most profitable pair of actions?

4. The three neighboring communities of Parched, Cactus and Tombstone are located in the desert and are analyzing their options for improving their water supplies. An aqueduct to the mountains would satisfy all their needs and cost in total $730,000. Alternatively, Parched and Cactus could dig and share an artesian well of sufficient capacity, which would cost $580,000. A similar option for Cactus and Tombstone would cost $500,000. Parched, Cactus and Tombstone could each individually distribute shallow wells over their respective surface areas to satisfy their needs for respective costs of $300,000, $350,000 and $250,000.

   Formulate and solve a simple LP for finding a plausible way of allocating the $730,000 cost of an aqueduct among the three communities.

5. Sportcasters say Team *I* is out of the running if the number of games already won by *I* plus the number of remaining games for Team *I* is less than the games already won by the league leader. It is frequently the case that a team is mathematically out of the running even before that point is reached. By Team *I* being mathematically out of the running, we mean there is no combination of wins and losses for the remaining games in the season such that Team *I* could end the season having won more games than any other team. A third-place team might find itself mathematically though not obviously out of the running if the first and second place teams have all their remaining games against each other.

   Formulate a linear program that will not have a feasible solution if Team *I* is no longer in the running.

The following variables may be of interest:

x_{jk} = number of times Team j may beat Team k in the season's remaining games and Team I still win more games than anyone else.

The following constants should be used:

R_{jk} = number of remaining games between Team j and Team k. Note the number of times j beats k plus the number of times k beats j must equal R_{jk}.

T_k = total number of games won by Team k to date. Thus, the number of games won at season's end by Team k is T_k plus the number of times it beat other teams.

6. In the 1983 NBA basketball draft, two teams were tied for having the first draft pick, the reason being that they had equally dismal records the previous year. The tie was resolved by two flips of a coin. Houston was given the opportunity to call the first flip. Houston called it correctly and therefore was eligible to call the second flip. Houston also called the second flip correctly and thereby won the right to negotiate with the top-ranked college star, Ralph Sampson. Suppose you are in a similar two-flip situation. You suspect the special coin used may be biased, but you have no idea which way. If you are given the opportunity to call the first flip, should you definitely accept, be indifferent, or definitely reject the opportunity (and let the other team call the first flip). State your assumptions explicitly.

7. A recent auction for a farm described it as consisting of two tracts as follows:

   Tract 1: 40 acres, all tillable, good drainage.
   Tract 2: 35 acres, of which 30 acres are tillable, 5 acres containing pasture, drainage ditch and small pond.

   The format of the auction was described as follows. First Tract 1 and Tract 2 will each be auctioned individually. Upon completion of bidding on Tract 1 and Tract 2, there will be a 15 minute intermission. After that period of time, this property will be put together as one tract of farmland. There will be a premium added to the total dollar price of Tract 1 and Tract 2. This total dollar amount will be the starting price of the 75 acres. If, at that time, no one bids, then the property will go to the highest bidders on Tracts 1 and 2. Otherwise, if the bid increases, then it will be sold as one.

   Can you think of some modest changes in the auction procedure that might increase the total amount received for the seller? What are some of the game theory issues facing the individual bidders in this case?