Introduction

A distinguishing feature of perfect information games is that each player knows exactly how the game has progressed up to his move. As we explained in the previous chapter, the backwards induction principle derived for this class of games cannot be extended to games of imperfect information. On one dimension this chapter broadens our enquiry to include games with all types of information sets. But the range applications is narrowed on another dimension, for we focus on payoff structures with special properties. Moreover since the extensive form is ill suited for analyzing imperfect information games, this chapter uses the strategic form to make headway.

In some games the payoffs are aligned so that rational players would choose certain strategies regardless of what the other players do. Other games lack this property, but their strategic forms lead one or more players to avoid rule out some strategies without considering what the other players might choose. Strategies which are optimal for a player regardless of whether the other players play rationally or not are called dominant. Strategies that yield payoffs less than could be obtained from playing an alternative strategy irrespective of what the other players pick, are said to be dominated. This chapter analyzes how the notion of dominance helps solves certain classes of games.

Strictly Dominant Strategies

The choice that yields the highest payoffs for one player might not depend on what the other players choose. When this situation arises we say the player has a dominant strategy. When all players have a dominant strategy, then strategic considerations do not affect the choices of anybody. Although a player’s payoff might depend on the choices of the other players, and although he might wish at least one of others would make a different choice, when a dominant strategy exists, the player has no reason to introspect about the objectives of the other players in order to make his own decision.

Mealtime

A game in which everybody has a dominant strategy that coincides with each player receiving their maximal payout is sometimes referred to as a win-win situation. Suppose a family of four visit a restaurant together, and the children are seated side by side on a bench facing their parents. Each child must decide in which side of the bench to sit. The left side is preferable to the ride because it has easier access to the corridor, whereas the right side of the bench butts up against the wall. All this is immediately obvious to everyone in the family, but there are other relevant considerations too. One sibling is right handed and the other eats with her left. If both children want to sit on the same side of the bench, there is a tussle that creates pandemonium, pleasurable fro the children but not their parents, bringing disciplinary action, so that ultimately the strategy profile of picking which is costly to both siblings. Utility in this case is measure by the number of jabs received in the shoulder by one’s sibling. These jabs are not punished providing no commotion is created. The strategic form is depicted below in Figure 7.1.

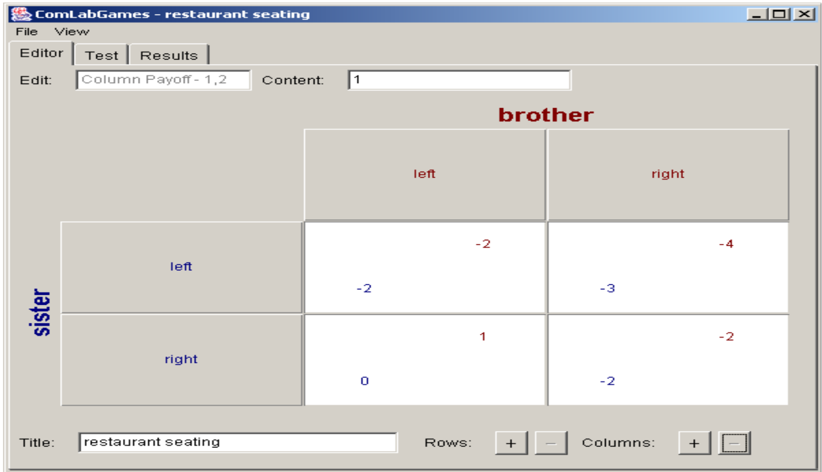

Figure 7.1
Mealtime

Notice that although the seating arrangement profile (right,left) could be considered a win-win situation, the sister might not regard it as fair. After all the brother benefits the equivalent of one shoulder jab. Nevertheless this example shows that the use of dominant strategies can yield Pareto optimality, situations where pursuing a different strategy profile that improves the utility of can only be achieved at the expense of somebody else. It would, however, be mistaken to conclude this necessarily the case.

Exercise *Consider*

Prisoner's dilemma

The prisoner's dilemma is another example of a game with dominant strategies for both players. Lacking evidence that might solve a heinous crime and facing public outcry, the police force in a country with a poor civil rights record randomly arrests two strangers, and charges them with committing the crime together. They are questioned separately, and given the opportunity to plea bargain. If neither of the arrested parties, or prisoners, confess to the heinous crime both are detained for a stiffer interrogation, but ultimately allowed to go free after a human rights organization has brought attention to their plight. If one of them coop confesses that they pared up and committed the crime together, he is forgiven for the remorse he has shown and the other prisoner is executed in public. If both parties confess a long gaol sentence meted out to both prisoners. The payoffs for the strategic form of the Prisoner's dilemma game are depicted in Figure 7.2 reflect the ineuqalities the description

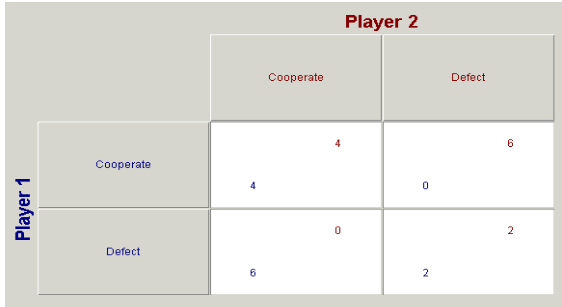

Figure 7.2
The Strategic Form of the Prisoners Dilemma

low reflect and the is. The best move for each prisoner does not depend on what the other prisoner does. The Prisoners' dilemma is perhaps the most famous example of a simultaneous move game in which both players have dominant strategies, because if both players did not play the dominant strategy profile, then they would both be better off. The strategy space for both prisoners are the two choices, Cooperate and Defect. The strategic form, represented by the matrix in Figure 7.2, illuminates the only reasonable outcome for this game. Regardless of which strategy the second prisoner selects, the first prisoner is better off choosing Defect. In the event of receives 6 versus 4 and 2 versus 0. A symmetric argument applies to the second prisoner. Thefore the solution to this game is for both prisoners to select Defect. In the prisoners' dilemma, defecting is a dominant strategy for both players.

The prisoner's dilemma is celebrated in game theory because there are many situations which resemble it, and because the prediction it yields is quite powerful, and collectively dissasterous for both parties.

When output is the result of teamwork that rely on unmonitored individual effort, individuals typically contribute less than they would if their marginal product could be observed.

Common property such as freeways, rivers and oceans are congested, overfished, pollute and exploited beyond renewal rates, because the private benefits of making an extra trip, are not offset against the costs an extra vehicle.

The scientific profession depends on researchers advancing science by revealing their knowledge to each other, as they discover it but individuals are rewarded for their personal contributions. Thus researchers release their new information strategically in order to maximize the personal credit accorded to them.

Airlines and the broadband are notable because of an overcapacity problem. Each firm other industries strive to take a larger share of the market pushing the price into an elastic region where the total revenue to the industry falls as the price drops

Arms race countries strive for detente but individually build up greater influence and strategic play.

Experimental results

Students from two programs played two rounds of prisoner's dilemma game and one group played only one round. Summary of the data in Table 7.2. shows that players 1 played the dominant strategy, defect more often 73.24 percent and players 2 defected 67.61% of the time. However the result is significantly different from the theoretical prediction that they should play "defect" all the time.

	Cooperate	Defect	Total
Cooperate	5.63% (4)	21.13% (15)	26.76% (19)
Defect	26.76% (19)	46.48% (33)	73.24% (52)
Total	32.39% (23)	67.61% (48)	100% (71)

Table 7.2: Experimental results for Prisoner's dilemma game

One would expect that deviations might have occurred in the first round and that subjects changed their strategy in the next round. Figure 7.2a shows that subjects who cooperated in round one are more likely to cooperate in the next round. Similarly subjects who defected in the first round defected in the second round 71.43 percent of the time if they were Player 2 and 80 percent of the time if they were Player 1. Because they were enough people who deviated we checked if the expected payoffs from deviating exceed the predicted payoffs. If Player 1 deviated from defect his expected payoff was $0.84 and if he played the dominant strategy he can expect the payoff of 3.46. If Player 2 deviated from the dominant strategy his expected payoff was

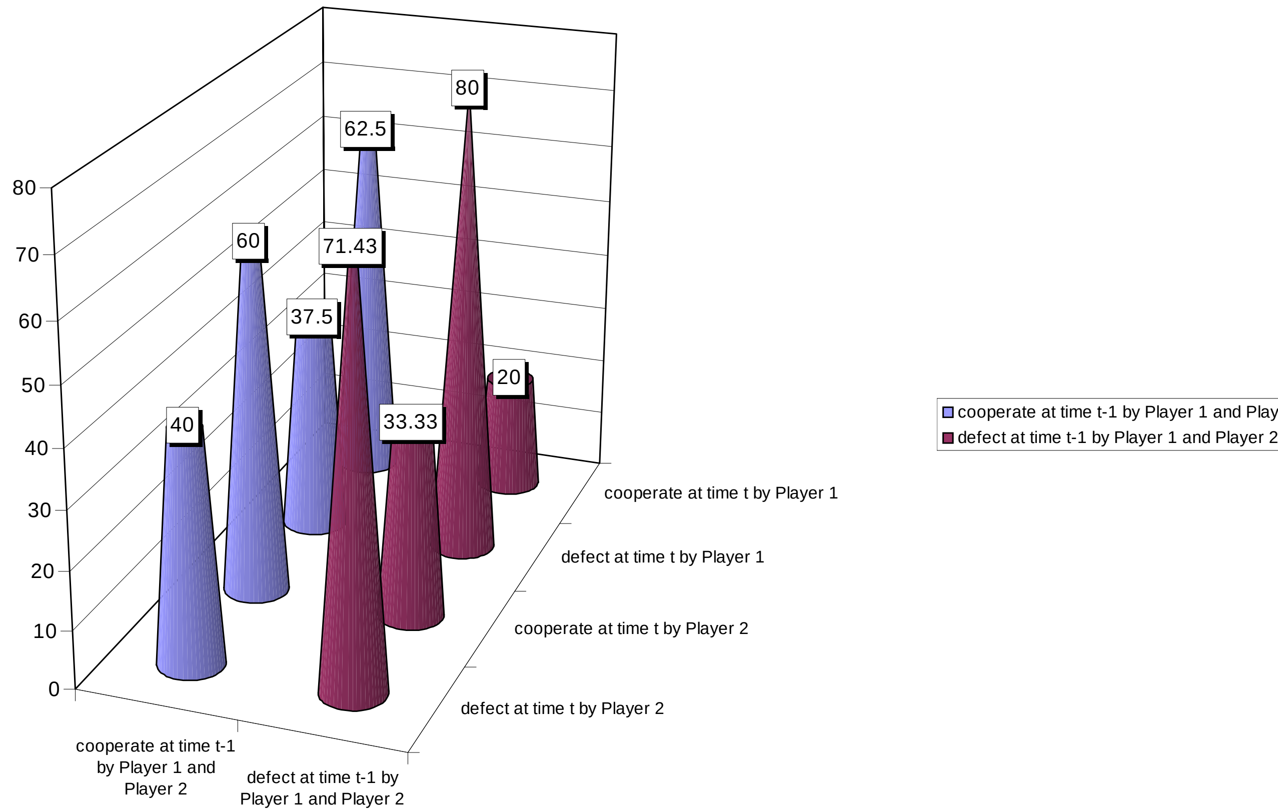

Figure 7.2a: Likelihood of staying with the same choice in the next round

Exercise *Two telecommunications firms simultaneously decide whether to pursue local or global networking strategies. If one firm chooses a local strategy, its net revenues are $3 billion per year if its rival chooses the same strategy in a different part of the world, but loses $1 billion per year if the other firm goes global. In this case the rival would earn $5 billion per year. If both firms choose a global strategy, then competition between the firms in the single global market reduce the total net revenue to the industry to $2 billion per year, both firms earning half the industry profits.*

1. *Write down the strategic form as a bimatrix.*
2. *Conduct an experiment of the game and tabulate the results.*
3. *What is there a dominant strategy?*
4. *Test the hypothesis that the subjects play the dominant strategy.*

A formal definition of dominant strategies

A strictly dominant strategy for a player yields a strictly higher expected payoff than any other strategy available to the player, regardless of the strategies chosen by everyone else. In terms of the notation defined in the previous chapter let:

$$S_{-n} \equiv S_1 \times S_2 \times \ldots \times S_{n-1} \times S_{n+1} \times \ldots S_N$$

denote the strategy space for all the other players in the game aside from Player n, with typical element s_{-n}. The strategy $\bar{s}_n$ for the n^{th} player is strictly dominant if it satisfies the inequalities

$$u_n(\bar{s}_n, s_{-n}) > u_n(s_n, s_{-n})$$

for all $s_{-n} \in S_{-n}$ and $s_n \in S_n$ if $s_n \neq \bar{s}_n$. Note that if a dominant strategy exists, it is unique by definition.

This leads us to the second rule for playing strategically.

Rule 2 Play a dominant strategy if it exists.

Strictly Dominated Strategies

At the opposite extreme of dominant strategies are dominated strategies, which yield lower payoffs than could be obtained by playing another strategy, regardless of what the other players choose. Thus a strictly dominated strategy for a player yields a lower expected payoff than at least one other strategy available to the player, regardless of the strategies chosen by everyone else. It follows that if one player has a dominant strategy in a game, then all the other strategies are dominated. If a player has only two strategies to pick from, and one of them is dominant, then the other one must be dominated. For example in the prisoners' dilemma game, cooperate is a dominated strategy, because defect strictly dominates cooperate. However if the player can choose between more than two strategies, the existence of a dominated strategy does not guarantee a dominant strategy. In such cases discarding dominated strategies narrows the search for a solution strategy, but not does identify a unique solution.

Essay

Games with dominated strategies do not necessarily have dominant strategies. While there is at most one dominant strategy, there many be many dominated strategies. For example in the essay game, whose strategic form is depicted in Figure 7.5, the strategy of submitting gibberish is dominated by submitting an original contribution. Furthermore if the teacher grades the essay for exposition and content, the student does best by submitting an original contribution. However if the teacher only checks spelling and grammar, the payoff from plagiarizing exceeds the payoff from submitting an original contribution.

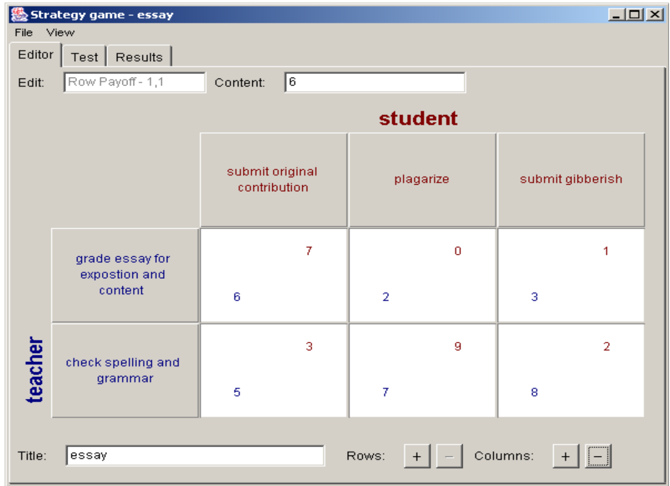

Figure 7.3
Essay

Marketing Groceries

The essay game highlights a strategy that is dominated by a pure strategy, but the following example illustrates a strategy which is dominated by a mixed strategy. Supermarkets compete with local grocery stores for business. We model the nature of their competition as a choice over three attributes shoppers value, the prices of their products, the range of products, the length of the checkout, the air-conditioning system within the store, ease of close parking and other customer services, and the hours of operations each day and proximity to demanders. Supermarkets typically occupy a very different location in the attribute spectrum to a corner store franchise, and on a much larger scale, so this is reflected in the bimatrix payoff entries.

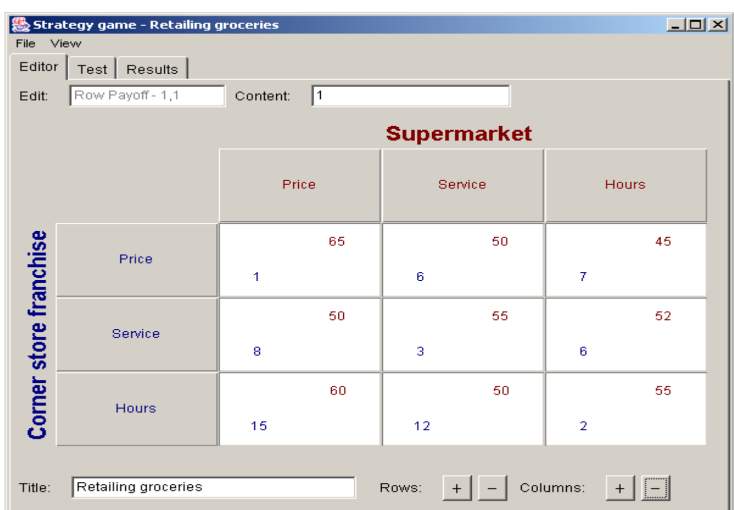

Figure 7.4

Retailing groceries

We can show that the supermarket's hours strategy is dominated by some mixtures of the price and service strategies. Let π denote the probability that the supermarket chooses a price strategy, and $(1-\pi)$ denote the probability that the supermarket chooses a service strategy. this mixture dominates the hours strategy if

$$\pi 65 + (1-\pi)50 > 45$$

$$\pi 50 + (1-\pi)55 > 52$$

$$\pi 60 + (1-\pi)50 > 55$$

Rearranging we obtain

$$\pi > -1/3$$

$$53/55 > \pi$$

$$\pi > 1/2$$

Hence all mixtures of π satisfying the inequalities

$$1/2 < \pi < 53/55$$

dominate the hours strategy.

Notation for dominated strategies

In terms of the notation defined in the previous chapter let:

$$S_{-n} \equiv S_1 \times S_2 \times \ldots \times S_{n-1} \times S_{n+1} \times \ldots \times S_N$$

denote the pure strategy space for all the other players in the game aside from Player $n,$ with typical element s_{-n}, let d_n denote the number of pure strategies at the disposal of player $n,$ and let Π_n denote the d_n dimensional simplex that defines the set of mixed strategies for player $n.$ The strategy $\underline{s}_n$ for the n^{th} player is strictly dominated by a mixed strategy $\pi_n^* \equiv (\pi_{1d}^*, \pi_{2d}^*, \ldots, \pi_{nd_n}^*) \in \Pi_n$ if the inequalities

$$u_n(\underline{s}_n, s_{-n}) < \sum_{d=1}^{d_n} \pi_{nd}^* u_n(s_{nd}, s_{-n})$$

are satisfied for all $s_{-n} \in S_{-n}$.

This leads us to our third rule.

Rule 3 Each player should discard their own dominated strategies.

Experiment

Subjects played this game three rounds, each time with a different player. Subjects played the game before the discussion about the dominance was introduced. Subjects who were assigned a supermarket role never played hours strategy that was dominated by the mixture of price and services. However they did occasionally, five times out of 45, selected services (see Table 7.4). Subjects who were assigned the role of the corner store franchise always played hours strategy.

		Supermarket	
		Price	***Services***
Corner store franchise	***Hours***	88.89 % n=40	11.11 % n=5

Table 7.4

Exercise *Dominated strategies*

1. *In a two player game the row player can choose three strategies, respectively named bottom, middle and top, while the column player can choose two strategies, called left and right. The row player receives a payoff of zero if he plays the bottom strategy, and a payoff of two if he plays the middle strategy, regardless of whether the column player plays left or right. The two player does not have a dominant strategy. What inequalities must the payoffs from the top strategy satisfy?*

Iterative Dominance

Rules 2 and 3 rely on a player recognizing strategies to play or avoid independently of how others behave. If all players recognized situations in which these two rules applied and abided by them, and one of the players realized that, then this particular player should exploit this knowledge to his own advantage by refining the set of strategies the other players will use. Knowing which strategies the other players have eliminated reduced the dimension of his problem, ruling out possible courses of action that might otherwise look reasonable.

A Second Restaurant

The easiest situation to recognize where iterative dominance applies is when all the players except one have a dominant strategy. If the player who does not have a dominant strategy is rational and believes that everyone else recognizes and plays dominant strategies, she will choose her own strategy based on the assumption that everyone else is choosing their dominant strategy.

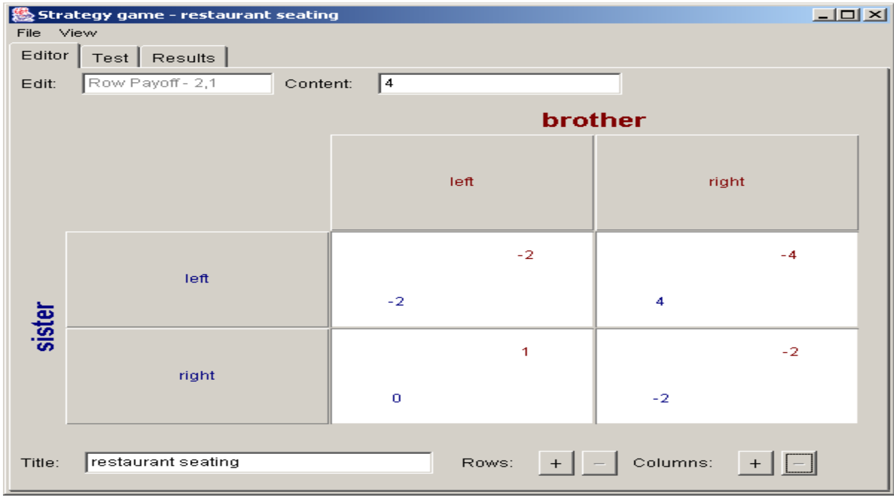

Figure 7.2
Another restaurant

Exercise *At a third restaurant, the payoffs to the siblings are*

1. *Solve the game using the dominant strategy principle*

2. *Conduct an experiment based on the strategic form of this third restaurant*

3. *Test the hypothesis that the outcomes from the experiment play the outcome*

4. *Split the sample into two groups, and test the hypothesis that the distribution of the outcomes change between the first time a subject plays an experiment and later times.*

5. *Test the hypothesis that players assigned to a role in which there is a dominant strategy converge to the solution more quickly than those who solve the games using the principle of iterated dominance.*

Groceries Revisited

The next easiest situation to recognize is one when one player eliminates a dominated strategy by one player the remaining players have dominant strategies.

For example, upon inspecting the payoff matrix in Figure 7.4 facing the rival supermarket, the corner store franchise might realize that its rival will not choose the hours strategy because it is dominated by a mixed strategy of playing price and service. If the franchise owner knows that the supermarket will not choose the hours strategy, then player the franchise owner can simplify his own choice problem by considering a game in which the supermarket is not allowed to play the dominated strategy. Eliminating the supermarket's hours strategy from consideration yields a simultaneous move game with a smaller strategy set for the supermarket, with a strategic form displayed in Figure 7.5.

Figure 7.5

Reducing the strategy space of the supermarket

The corner store franchise has a dominant strategy in the reduced game. Regardless of whether the supermarket chooses a price or a service strategy, the corner store should pick hours. If the supermarket chooses service, the corner store's profits from choosing hours are more than twice the amount it could earn by choosing service or price, while if the supermarket chooses price to compete on, the corner store's profits are more than 50 percent greater.

We could add another link to this chain of reasoning. Suppose the marketing manager at the supermarket is just as capable as the corner store franchise owner of make such logical inferences about strategic play. In this case the supermarket's marketing manager will recognize that the corner store franchise owner knows that the supermarket will not choose an hours' strategy and will consequently choose that strategy itself. Referring again to Figure 7.5, if the corner store franchise picks hours, the supermarket's profits from choosing price is 60, but its profits from choosing service are only 50. The supermarket thus settles on a price strategy.

Market games

Part 7 of this book is devoted to market games, studying how strategic play between traders affects market outcomes, supplemented by experiments that are conducted in real (continuous) time, rather than mediated through (discrete time) moves. We can, however, probe some the strategic aspects of play by solving the strategic form of price and quantity setting games between producers, for the moment treating buyers as a demand curve rather than as rational agents.

We suppose that consumer demand for a product is a linear function of price taking the form

$$q = \alpha - \beta p$$

where q is the industry quantity demanded, p is the price consumers pay for the product, while α and β are two positive coefficients that respectively determine its overall level and price sensitivity. Higher values of α raise the level of demand for any given price, while higher values of β cause the decline in quantity demanded to be more pronounced as the price increases. The latter point is reflected in the formula for the price elasticity of this linear demand schedule, denoted ε_d and defined as

$$\varepsilon_d \equiv \frac{p}{q}\beta$$

In this first application we assume that the cost to each supplier of producing a unit of the commodity, denoted $c,$ is a parameter that does not depend on the scale of production. In other words the marginal cost of production is constant.

In a Bertrand pricing game between two producers, both firms set the price of their product and fulfill all orders on demand. We denote by p_1 the price set by the first firm, and p_2 the price set by the second firm. The vector (p_1,p_2) are the choice variables in this simultaneous move game. If the prices charged by the firms differ, consumers only buy from the lowest price supplier, and if they charge the same price demand is split equally between the two suppliers. As a function of $(p_1,p_2),$ the net profit to the first firm, denoted by $\pi_1(p_1,p_2),$ is therefore:

$$\pi_1(p_1,p_2) = \begin{cases} (\alpha - \beta p_1)(p_1 - c) \text{ if } p_1 < p_2 \\ \frac{1}{2}(\alpha - \beta p_1)(p_1 - c) \text{ if } p_1 = p_2 \\ 0 \text{ if } p_2 < p_1 \end{cases}$$

Similarly, the net profits of the second firm are:

$$\pi_2(p_1,p_2) = \begin{cases} (\alpha - \beta p_2)(p_2 - c) \text{ if } p_2 < p_1 \\ \frac{1}{2}(\alpha - \beta p_2)(p_2 - c) \text{ if } p_1 = p_2 \\ 0 \text{ if } p_1 < p_2 \end{cases}$$

This game can be solved by iteratively removing dominated strategies. Charging a very high price

For experimental purposes it is useful to imagine that prices charged from a discrete set (perhaps because the units prices are denominated in whole dollars or cents) a discrete version of the pricing game. Consider the following Bertrand pricing game:

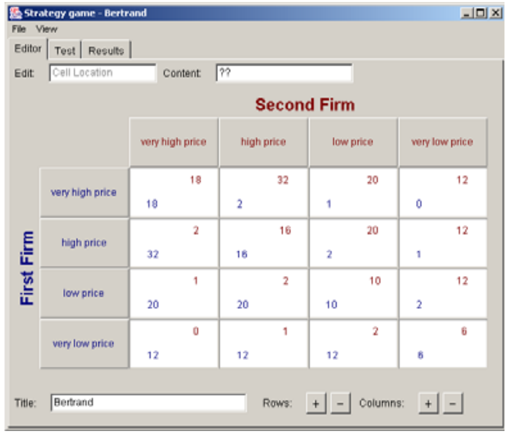

Figure 7.6
Price competition

Note that charging the highest price is a dominated strategy.

Exercise *Suppose $\alpha = 120$ and $\beta = 3$. Also assume that prices are in 50 cent intervals, and that the first firm charges in round dollar units only, and the second always is always sells at the 50 cent remainder. Suppose $c = 8$.*

1. *Write down a matrix representation of the game*
2. *Conduct an experiment of the market based on that representation.*
3. *Solve the strategic game implied by the*
4. *Test whether the distribution of outcomes differs significantly from the solution*

Duopoly quantity games

An alternative to competing on price is to compete on quantity, by choosing production levels and letting the market determine the price that keeps inventory levels roughly constant. Denote by q_1 and q_2 the quantities chosen by the firms. Then the industry price is derived from the demand curve as

$$p = (\alpha - q_1 - q_2)/\beta$$

Net profits to the i^{th} firm are now denoted as the mapping from

$$\lambda_i(q_1, q_2) = \left[\frac{(\alpha - q_1 - q_2)}{\beta} - c \right] q_i$$

The marginal profit from increasing quantity is found by differentiating $\lambda_1(q_1, q_2)$ with respect to q_1 to obtain

$$\frac{\partial \lambda_1(q_1, q_2)}{\partial q_1} = \frac{(\alpha - 2q_1 - q_2)}{\beta} - c$$

Notice that the marginal profit function is declining in the quantity of both firms. This means that profits fall off uniformly as the player moves away from the optimum. Thus iteratively removing dominated strategies would work here as well.

In this chapter we focus on a discrete version of the problem

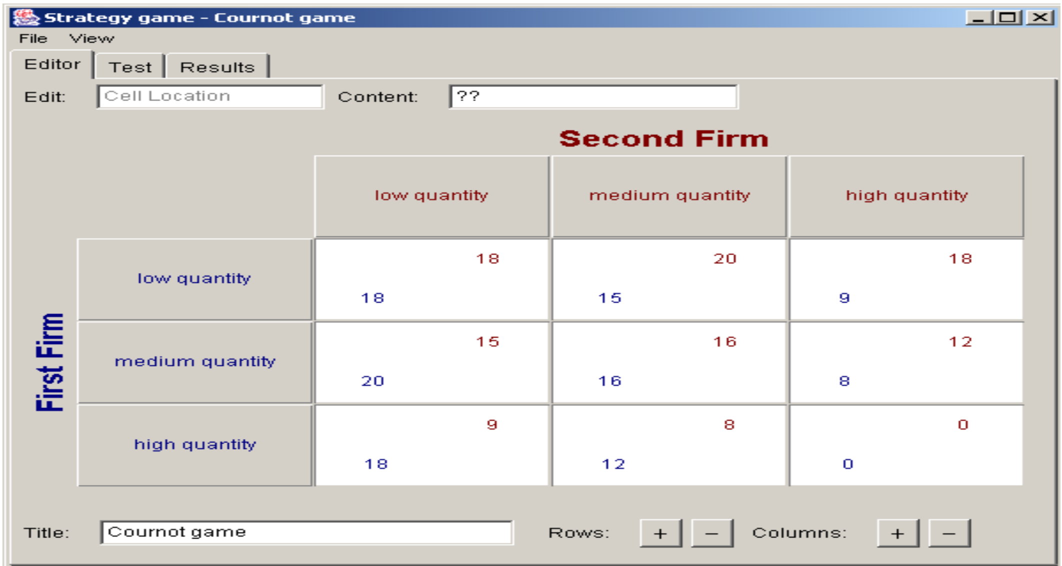

Figure 7.7
Quantity competition

These examples show that the dimensions on which companies compete can affect quite significant affects on the market outcome.

The Cournot Limit

Notice that increasing the number of firms in the pricing game has very little effect on the prices that are charged. Now let us investigate what happens as the number of firms increase when firms compete on quantity. Denote by q_n the quantity chosen by the n^{th} firm, where $n \in \{1, \ldots, N\}$ and for convenience define the industry quantity as

$$q = \sum_{n=1}^{N} q_n$$

Then the industry price is derived from the demand curve as

$$p = (\alpha - q)/\beta$$

Net profits to the n^{th} firm are now denoted as the mapping from

$$\lambda_n(q_1, \ldots, q_N) = \left[\frac{(\alpha - q)}{\beta} - c \right] q_n$$

The marginal profit from increasing quantity is found by differentiating $\lambda_1(q_1, q_2)$ with respect to q_n to obtain

$$\frac{\partial \lambda_n(q_1, \ldots, q_N)}{\partial q_n} = \frac{(\alpha - q_n - q)}{\beta} - c$$

We use the principle of removing iteratively dominated strategies to solve the game and show convergence to the price equals marginal cost condition.

Rules for Rational Players

This example illustrates the technique of solving a game by iteratively removing strictly dominated strategies. It not only assumes that both players are rational, but also assumes that the row player expects the column player to be rational, and expects the column player to expect the row player to be rational, and so on. Service is not a strictly dominated strategy for the supermarket in the original game. The supermarket marketing manager eliminates service because he expects the franchise owner to play hours, because the marketing manager expects the franchise owner to believe that the marketing manager will not to play hours. In this way, the concept of a dominant strategy can be extended to the notion of iterated dominance. Recognizing that one player has a dominated strategy, the remaining players automatically rule out the possibility that it will be chosen. Ruling out this strategy might then reveal a dominated strategy for another player in the reduced game. The successive elimination of dominated strategies is called iterative dominance. Thus the game of retailing groceries is an example of a game which can be solved using the principle of iterative dominance.

The principle of iterative dominance makes more demands on the sophistication of players than the principle of dominance by itself. In simultaneous move games, applying the principle of dominance makes sense regardless of how other players behave. In contrast, a player who chooses according to the principle of iterative dominance does so because he believes the other players choose according to that principle too. This renders the principle of iterative dominance less robust than the principle of dominance, and that is the price of a tighter prediction. Applying this principle to solve games requires more sophistication than each player simply being able to recognize a dominant strategy fro themselves. Each player must be able to recognize all the dominated strategies that everyone faces, reduce the strategy space of each player as called for, and then repeat the process.

Having acknowledged that limitation, however, it is important to recognize that something akin to iterative dominance has already been applied when the designer of the game designs the game in the first place. In the restaurant game, for example, the strategies of throwing a tantrum or bribing a sibling are not modeled as part of the game. these strategies might well be relevant for some families, and if they are they should obviously So the fact that this game in Figure 7. does not allow these strategies is an indicator that they are not relevant, presumably because it is not reasonable to play them under any circumstances that might arise in the game.

This remark underscores another motivation for iteratively eliminating dominated strategies: it is possible to iteratively expand many games by enlarging the strategy space with dominated strategies. The exercises will work on inventing plausible examples of that phenomenon, and also recognizing dominated and iteratively dominated strategies in both strategic form and extensive form games. There are many games that contain dominated strategies, and it is not always clear that all the dominated strategies are written down. Iterated dominance thus plays an important latent role in game design.

This is not to minimize the significance of the solicitation required to follow the rule of iterated dominance, nor even to argue that games which are dominant solvable should not be written down as games at all. Rather it should be interpreted as a statement about the ability to model strategic play as a rational behavior by designing models in which players make choices recognizing that the outcomes and their individual payoffs depends on what everyone else does.

Is the algorithm of iteratively removing strictly dominated strategies unique?

Question: Can we have different solutions if we use different sequence of truncations?

Answer: No

Fact: Different sequences of complete truncation of STRICTLY dominated strategies lead to the same set of solutions. The key to proving this point is that if a strategy is revealed to be dominated it will remain dominated if another strategy is removed first

Suppose that only by removing this strategy A a strategy fro another player is revealed dominated. then note that the later the A strategy is removed, then the other will surly be revealed dominated.

Theorem *The order in which strictly dominated strategies is revealed does not determine which strategies will be removed.*

We conclude this section by

Rule 4 Successively eliminate any dominated strategies from consideration in an iterative fashion

Exercise *Consider Figure 7.6 in which the payoff structure has been modified. There is a Juke box on the wall to the right of the dining booth, and the child nearest it controls which songs will be played. The siblings have diametrically opposed tastes in musical appreciation.*

1. *Run the game and tabulate the results.*
2. *Prove that player has a dominant strategy, and the other does not.*
3. *Then use the principle of iterated dominance to solve the game.*
4. *Compare the tabulated results with the outcome predicted by game theory.*

Weak Dominance

Many implausible strategies survive the application of the iterated dominance rule, leading one to question whether the strict dominance is too stringent a requirement. Rather than require that a strategy be strictly dominant after eliminating strictly dominated strategies, we might require something less demanding. For example suppose the selected strategy is at least as good as any other strategy regardless of what the other players do, and in all pairwise comparisons with the other strategies, the selected strategy is strictly better than the alternative for at least one strategic profile played by the other players. The selected strategy is then called a weakly dominant strategy.

Teamwork

For example, consider the following team game. A plumber and his assistant work with each other on a contract basis, and are employed to undertake some repairs. The plumber is paid considerably more than his assistant, \$500 versus \$200, but they value leisure time equally at \$100 for the time spent on the job. If only one worker shows up, he cannot undertake the repairs by himself and must reschedule the work, which costs the equivalent of \$100 in terms of lost time and inconvenience.

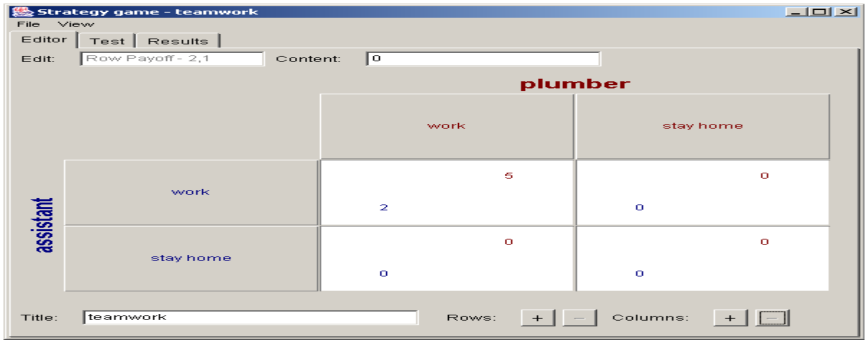

Figure 7.9
Teamwork

Individually, both players are rewarded more from working than by any other outcome. In this vital respect their interests fully coincide. Yet this game cannot be solved by strict dominance because the benefits from staying home are independent of what the other player chooses. Likewise, there are no strictly dominated strategies. However working is a weakly dominant strategy for both players, because it yields at least as much utility as the alternative, and depending on what the team member does, maybe more. Choosing work guarantees the plumber at least 0, the same level of utility as choosing stay home, and possibly 5.

First Price Auctions

Auctions are becoming an increasingly important mechanism for disposing of goods and awarding contracts to buy components and services. For this reason Part 6 of this textbook is devoted to studying how the number of objects up for sale (or purchase), different bidding rules, incomplete information about the valuations of bidders, the numbers of bidders and their opportunities for collusion, affect their bidding behavior and the prices the winners pay. Nevertheless some basic principles about optimal bidding in auctions can be explained without recourse to many of the complications that characterize auctions in reality.

In this introduction we focus on first and second price sealed bid auctions where just two bidders compete with each other for a single auctioned item, each of whom knows the valuations of both bidders. In a sealed bid auction each bidder submits a price, and the auctioneer awards the item to the bidder submitting the highest price. In a first price auction, the highest bidder pays his bid, while in a second price auction the highest bidder only pays the price of the second highest bid.

Suppose one bidder values the item at $4 million, and the other at $2 million. Figure 7.11 depicts the strategic form of the game. To calculate the expected payoffs when both players make the same bid, we assume each bidder stands an equal chance of winning the auction. Accordingly consider the top left matrix entries, which give the expected payout to both players ($0.5 million, $1.5 million) if both players bid $1 million. Since there is a 50 percent chance the high valuation player will win in this case, and his net gain would be $3 million, his expected gain is $1.5 million. Similarly the net gain to the low valuation player conditional on winning is only $1 million, so the expected value of both bidders submitting a price of $1 million is $0.5 million to him. The other matrix entries are completed in a similar fashion.

Figure 7.10

A First Price Sealed Bid Auction

The solution to this auction can be found by applying the rule of iterative dominance, and then selecting the weakly dominant strategy. The column player bidding $5 is (strictly) dominated by bidding $4. Upon deleting the last column of the matrix, it is now evident that the row player bidding $4 or more is dominated by bidding $3. The strategic form of the reduced game appears as Figure 7.13

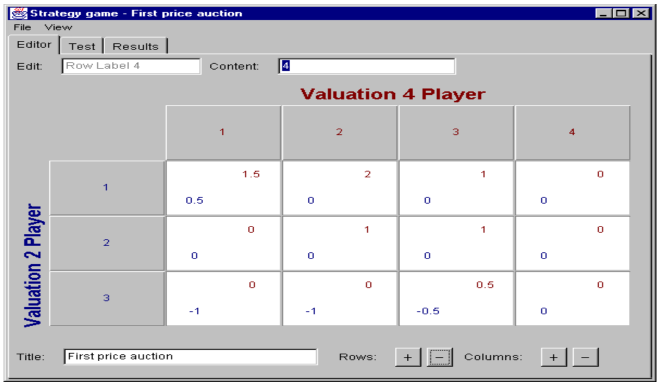

Figure 7.11

A reduced game

In the reduced auction the column player bidding $4 is dominated by bidding $3 in the reduced game. Continuing the algorithm of iteratively deleting dominated strategies, we remove the last column. Now the row player bidding $3 is dominated by bidding $2. so the bottom row of the reduced strategic form can also be deleted. Eliminating the iteratively dominated strategies shrinks the strategic form still further, as Figure 7.12 shows.

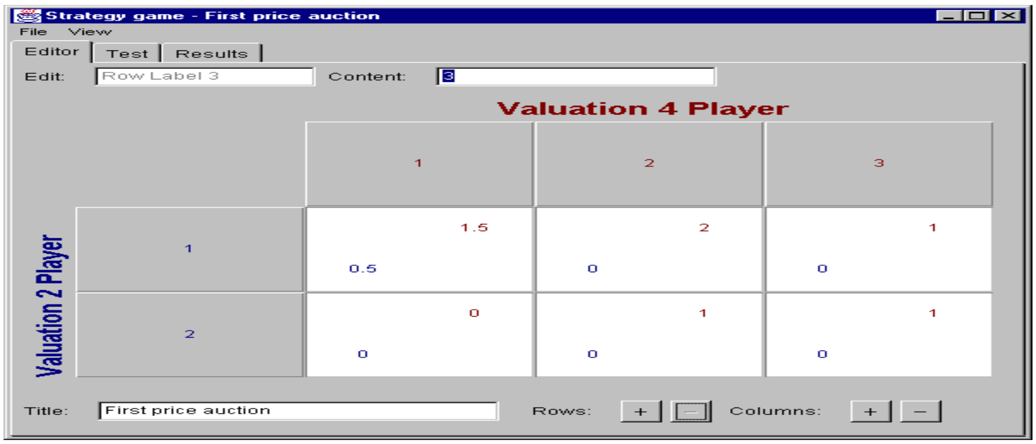

Figure 7.12

A further iteration

In the reduced game depicted in Figure 7.12 both players have weakly dominant strategies. Indeed we shall argue that having the row player bid $1 and the column player bid $2 uniquely solves the game.

Experimental results

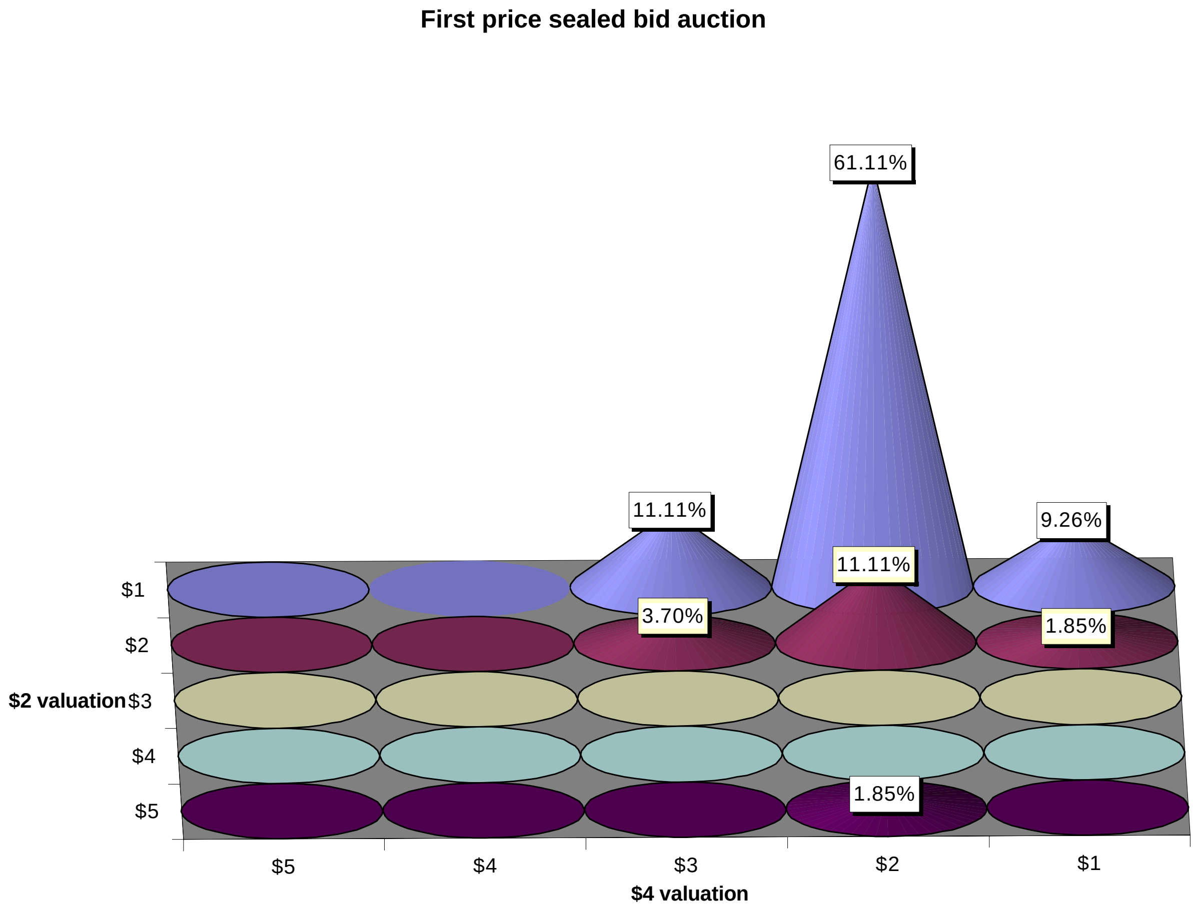

Exercise *In the example above we could define the valuations of the two players by the vector (2,4). This exercise investigates how the solution and the distribution of experimental outcomes vary with the valuations of the players. Answer the following questions when the payoffs are defined by (2,5), (3,5), (3,7) and (6,7).*

1. *Depict the matrix representation of the strategic form for a second price sealed bid auction in which only these two players submit bids.*

2. *Conduct an auction experiment and tabulate the results.*

3. *Compare the proportion of times the winner is the high valuation player, and the mean price of the winning bid conditional on the type of player winning the auction, and the standard deviation of the winning bid.*

4. *Repeat the exercise for the losing bid.*

Second Price Auctions

As we mentioned above, the only difference between first price and second price sealed bid auctions is the amount the highest bidder pays. In a first price auction the winner pays his bid, but in a second price auction the highest bidder is obligated to pay only the second highest bid in exchange for the auctioned object. If bidders adopt the same bidding strategy in both types of auctions, then except in the case of ties, first price auctions yield more revenue to the seller than second price auctions. However it hardly seems reasonable that bidders would choose the same bidding strategy when the rules of the auction change. Since first price auctions require the winner to pay a higher price, all players presumably adjust their bids downwards. Which factor is more important in theory and in practice?

We can investigate these questions by appealing the example above. As before suppose there are only two bidders, one of whom has a valuation of $4 and the other a valuation of $2. A strategic form for the second price sealed bid auction can be depicted in a similar way to the first price auction, by modifying the previous example to take account of the new bidding rule. Thus if the high valuation player bids $5 million and the low valuation player bids less than that, then the high valuation player wins the auction and pays the bid of the low valuation player. Similarly if the low valuation player bids $3 million then the item is sold to the high valuation player for $3 million for a gain of $1 million. Inspecting the cell corresponding to the (bid $4 million, bid $5 million) strategy profile, we see the payoffs are ($0, $1 million) as required. Figure 7.13 illustrates.

Figure 7.12
A Second Price Sealed Bid Auction

Reviewing the rows corresponding to payoffs for the Valuation 2 player, observe that bidding $2 is a weakly dominant strategy. Similarly an inspection of the column payoffs to the Valuation 4 player reveals that bidding $4 is weakly dominant. Therefore the unique solution to this game is for each player to bid his/her own valuation.

This example illustrates two results in auction theory that have much broader application. First, bidders should bid their true valuations in second price sealed bid auctions, a result we establish in Chapter 15. Second, although the solution strategies for the first and second price auctions are quite different, the revenues generated by the two auctions are identical. This is an example two auctions which are revenue equivalent but not strategically equivalent. We explore these concepts much more broadly in Chapters 15 through 18.

Experimental results

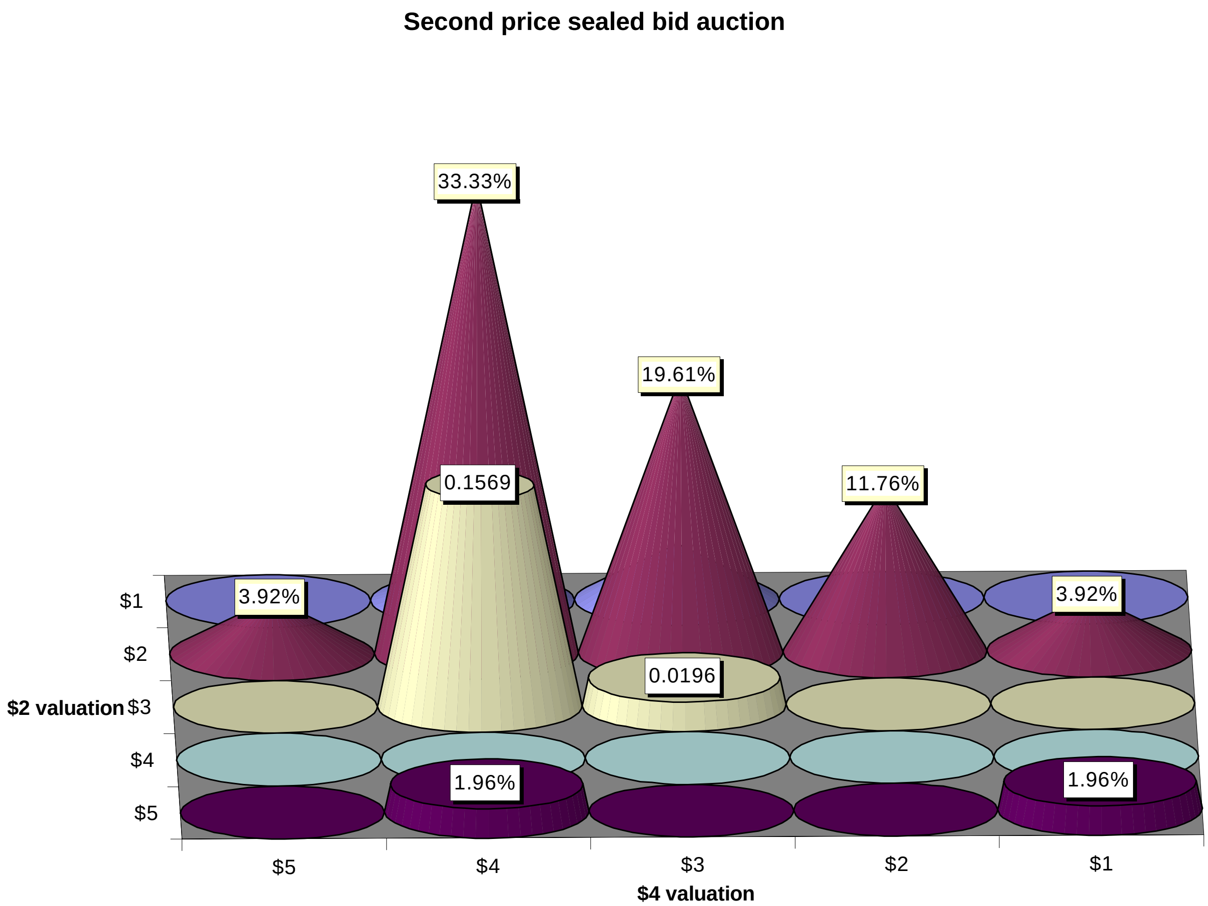

Exercise *This exercise investigates how the solution to a second price sealed bid auction and the distribution of experimental outcomes vary with the valuations of the players. Answer the following questions when the payoffs are defined by (2,5), (3,5), (3,7) and (6,7).*

1. *Depict the matrix representation of the strategic form for a second price sealed bid auction in which only these two players submit bids.*

2. *Run both auctions and tabulate the results.*

3. *Compare the proportion of times the winner is the high valuation player, and the mean price of the winning bid conditional on the type of player winning the auction, and the standard deviation of the winning bid.*

4. *Repeat the exercise for the losing bid.*

5. *Compare the distribution of the wining bid in the first price auction with the distribution of the losing bid in the second price auction.*

Principle of Weak Dominance

The formal definitions of a weak dominance parallel those for strict dominance. As before, let: S_{-n} denote the strategy space for all the other players in the game aside from Player n, with typical element s_{-n}. The strategy $\bar{s}_n$ for the n^{th} player is weakly dominant if it satisfies the inequalities

$$u_n(\bar{s}_n, s_{-n}) \geq u_n(s_n, s_{-n})$$

for all $s_{-n} \in S_{-n}$ and $s_n \in S_n$ if $s_n \neq \bar{s}_n$ and if there exists some strategy profile for all the other players, called $\widehat{s}_{-n}$ such that

$$u_n(\bar{s}_n, \widehat{s}_{-n}) \geq u_n(s_n, \widehat{s}_{-n})$$

for all $s_n \in S_n$ if $s_n \neq \bar{s}_n$.

By definition in every game there is at most one a weakly dominant strategy for each player

Rule 5 If a weakly dominant strategy exists, play it.

Weakly Dominated Strategies

One could further relax the dominance principle. Up till now we have only been removing strictly dominated strategies. What about removing a strategy that yields payoffs no higher than can be obtained from choosing another strategy regardless of what the other players do, where a strict inequality applies if the other players follow a certain strategy profile? This is the hallmark of a weakly dominated strategy, and at first sight the idea of iteratively eliminating the weakly dominated strategies seems almost as appealing as iteratively removing strictly dominated strategies. Indeed one of the most compelling cases for iteratively removing weakly dominated strategies is based on the fact that in perfect information games, the backwards induction algorithm corresponds to the iteratively eliminating weakly dominated strategies.

Formally a weakly dominated strategy can be defined as follows.

Suspect

A policeman has the opportunity to attempt apprehending an unarmed suspect and his alternative use of time is to eat a doughnut. The officer will be rewarded for bringing this person to the station for questioning but only if this can be accomplished without shooting the suspect. The suspect chooses between surrendering himself to the policeman when ordered to, or fleeing the scene. If the suspect tries to escape, the only way to prevent is for the policeman to shoot him. Figure 7.13 displays the extensive form of this game with payoffs that accurately rank the priorities of each player. As we have mentioned on previous occasions, the actual magnitudes do not affect the solution to this game.

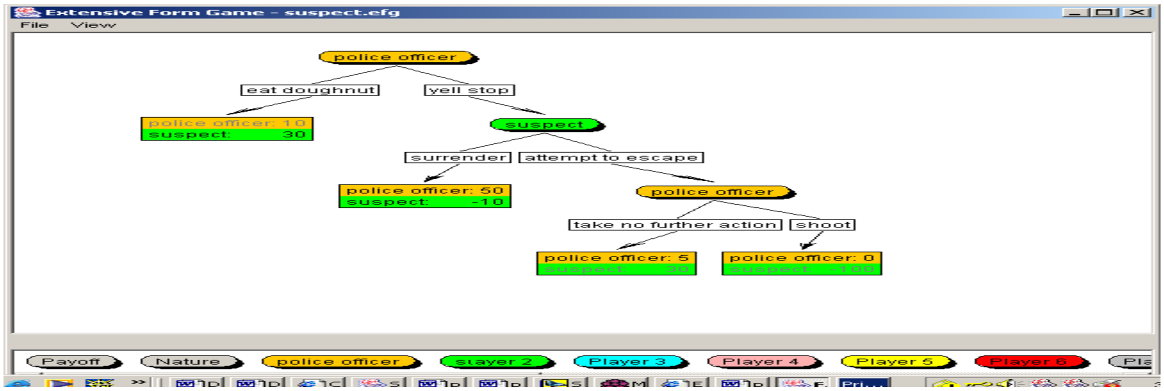

Figure 7.13
Suspect

If the suspect's decision node is reached and he flees, the policeman holds his fire, because he would rather suffer the indignity of the suspect ignoring his command to stop, than suffer the consequences of shooting an unarmed suspect. The reduced game is now displayed in its extensive form in Figure 7.14.

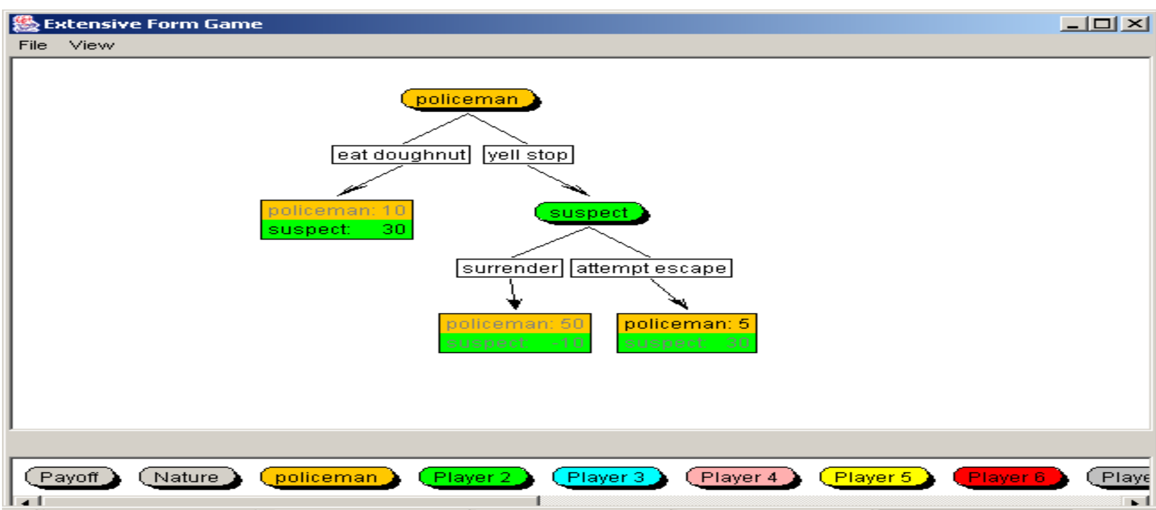

Figure 7.14
Reduced game of Suspect

From Figure 7.14, the suspect will flee, anticipating that the policeman will hold his fire. We now show how the backwards induction derivation of this solution corresponds to the elimination of weakly dominated strategies

The strategic form of suspect is given in Figure 7.15. The policeman has three strategies, to eat the doughnut, to command the suspect to stop but not to shoot if the suspect flees, and to shoot the suspect if he does not the policeman's command. The suspect has two strategies, to surrender when so commanded, and to flee when commanded to surrender.

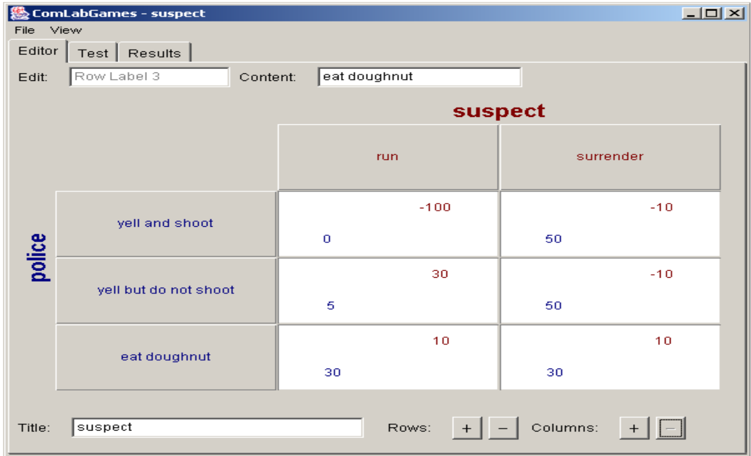

Figure 7.15
Suspect

Note that yelling and shooting is weakly dominated by yelling but not shooting. Eliminating yelling and shooting from consideration we are thus left with the strategic form of the reduced game displayed in Figure 7.16, whose extensive form is depicted in Figure 7.14.

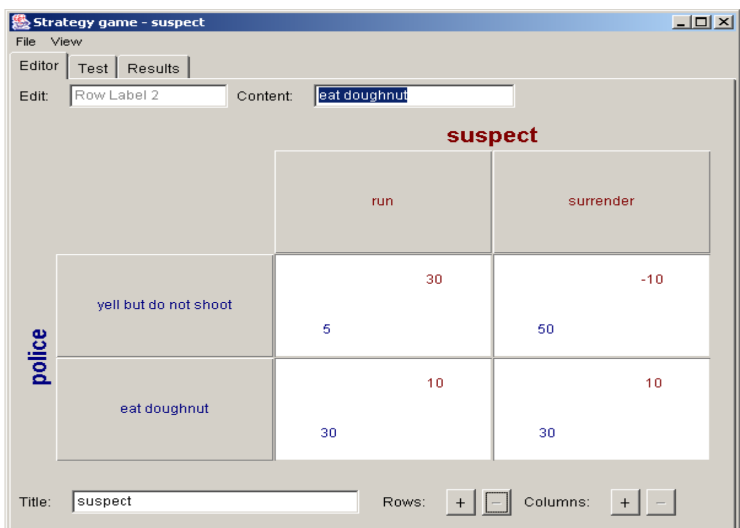

Figure 7.16
Reduced Strategic form of Reduced Suspect game

At this point we can see the suspect has a weakly dominant strategy to run, so the best response of the policeman is to eat a doughnut, as we found from using the backwards induction algorithm. Of course the fact that both algorithms produce the same solution in games of perfect information is comforting, but is not sufficient reason for using iteratively removing weakly dominated strategies in imperfect information games.

Iterative removal of weakly dominated strategies

One limitation of using this algorithm carelessly is that multiple solutions might be overlooked. The following simultaneous move game displayed in Figure 7.17 shows why adopting this suggestion might lead us to ignore some of the solutions to the game.

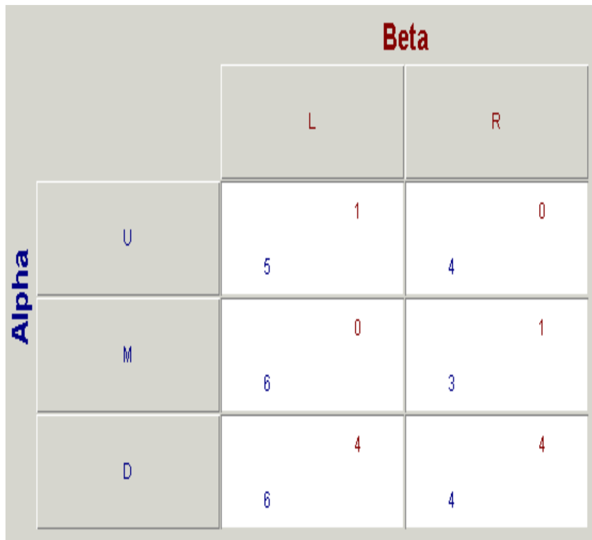

Figure 7.17
Multiple solutions

Note that U and M are weakly dominated by D for Alpha. Removing U first the reduced game becomes:

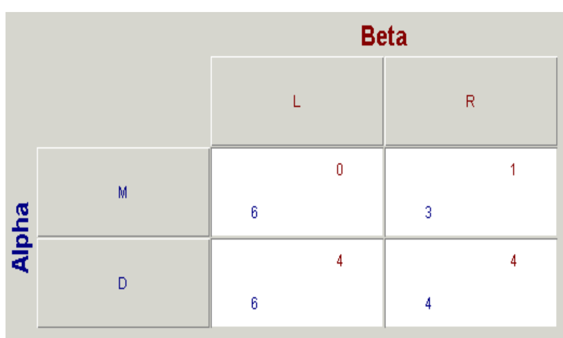

Figure 7.18
Removing U first

Noting that in Figure 7.18 that D and R are weakly dominant strategies for Alpha and Beta respectively, the unique solution of this reduced game is (D, R).

Now consider the consequences of removing M is removed first. We obtain the strategic form in Figure 7.19 for this reduced game.

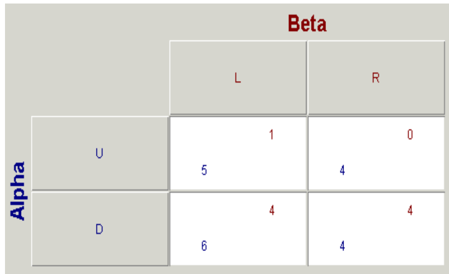

Figure 7.19
Removing M first

Notice D and L are weakly dominant strategies for Alpha and Beta, the unique solution for this reduced game is (D, L).

This example demonstrates that, in contrast to the iterative removal of strictly dominated strategies, the solution found by iteratively removing weakly dominated strategies is path dependent. Therefore if one follows the principle of iteratively removing weakly dominated strategies, then all the possible paths of reduced form games should be considered. However as we shall see in the chapter on sequential games with imperfect information, even this precaution does not prevent certain plausible candidates for solutions from being eliminated. In summary, we advocate using the principle of backwards induction, which corresponds to iteratively removing weakly dominated strategies in perfect information games, but not its wholesale use.

Applying the dominance principle in strategic play

The advantage of appealing to the dominance principle is that a strict interpretation of this principle is not at all controversial. We have argued that dominance is simple yet powerful principle that applies to many strategic situations. The only argument against picking a dominant strategy is that it might be hard to find in a complex environment.

We distinguished between strict dominance and weak dominance, and argued that choosing a weakly dominant strategy where possible is sound strategy. As we demonstrated the principle of backwards induction is in fact an application of the more general principle of weak dominance. This shows that if one accepts the principle of choosing a weakly dominant strategy, then the argument for applying backwards induction cannot be denied.

Applying iterative dominance further refines the set of choices that players are likely to pick and outcomes they are likely to reach. The successive elimination of dominated strategies is called iterative dominance. The principle of iterative dominance makes more demands on the sophistication of players than the principle of dominance by itself. In simultaneous move games, applying the principle of dominance makes sense regardless of how other players behave. In contrast, a player who chooses according to the principle of iterative dominance does so because he believes the other players choose according to that principle too. This renders the principle of iterative dominance less robust than the principle of dominance, and that is the price of a tighter prediction. For example suppose you believe somebody is stupid, and will pick a particular dominated strategy. This might reveal a new strategy for you that was previously dominated, and now would pay a higher reward. However the application of iterated dominance is intuitively appealing, especially when experienced players participate in games with which they are familiar.

Iteratively removing weakly dominated strategies is much less convincing. In an example we showed how different solutions might .At a minimum this suggest that all possible paths of iterative removal should be taken. But in fact matters are worse than this. In the chapter on complete information games we will argue that playing a strategy that is not weakly dominated but can be removed in this fashion might make good sense. Thus we do not advocate iteratively removing weakly dominated strategies, only strictly dominated ones.

The main disadvantage of the dominance principle is that does not always provide guidance about how to make a good strategic decision. In the next chapter we explain that to give more specific advice to strategic players, less plausible criterion must be employed. The solution concepts developed in the next chapter are less convincing, because they eliminate more plausible strategic behavior. Our approach of starting with principles that require little rationality, and progressing to those that make more demands of players, is a conscious expositional choice, because we wish to emphasize the trade off between the benefits of a more precise prediction at the cost of having less confidence in it.

Rivalry

This section investigates an industry where there are a fixed number of sellers competing against each other to supply a homogeneous product to a consumer population. We start by analyzing the case of two producers, duopoly. Then we extend our analysis to cover any finite number of suppliers, and show how the industry is affected as that number diverges. This leads us into a discussion of entry and exit by firms in response, and in anticipation of, changing product demand and input supply conditions. The results we derive in this section are based on dominance explained earlier in this chapter. There we showed that if firms set prices, the solution for the industry is for each firm to charge marginal cost, but if firms set quantities, then the solution would entail each firm making strictly positive profits. We now explore extenstions of these results.

Price competition between duopolists

We now explore the outcome of a game, where consumers can only submit market orders to one of two firms, each of which submits limit orders, typical of retail establishments for example. The monopolist first consider a form of price competition. Upon seeing a lower price from an alternative supplier buyers are free to exchange the good for a full refund and obtain it at the alternative source,

If producers have the same unit cost c, then price is driven down to marginal cost, and all consumers with valuations above marginal cost purchase the good at that price. The result can be established as follows. First we show that both firms charging c is a Nash equilibrium. Then we demonstrate that it is unique.

We remark that in this game a weakly dominant strategy is each consumer is to buy if the price falls below her valuation.

Give good away is dominated by charging a positive price. Now by iterative dominance suppose there is a positive price less than c which is not iteratively dominated. Repeat the argument to show that it is. Do the same from the top. Get equilibrium by iterative dominance.

Notice that neither firm makes any profits by charging c Neither firm will offer units of the good at below cost because then the lowest priced firm will sell units to consumers with valuations above the price, and hence make a loss. this proves that the solution price at least covers costs. Next, suppose one firm offers the good at some price above c. The best response of the other firm is to offer the good at a price between c and its rival's price, and in that way capture all the demand. By symmetry this implies that the first firm is not best responding to the second.

Figure 7.20 illustrates some of these points. In this diagram we assume that demand is drawn from a uniform distribution. the best response of each agent is buy a unit of the common

Figure 7.20

A minor variation of this argument can be used to establish that if firms have different cost, c_1 and c_2 respectively, with $c_1 > v$ say, then all consumers with valuations of c_1 or more will buy the good from the second firm at a price just below c_1. By the arguments we have made above, it is not optimal for the second firm to be undercut by the first first firm if the first firm charges above c_2. Yet the first firm charging is less than c_1 is an iteratively dominated strategy.

Figure 21.1 illustrates this case. Again we have mapped in the demand curve for the consumers since there is a weakly dominant strategy

We now explore how the price and aggregate quantity changes as the number of sellers. If at least one seller chooses a pricing policy, the solution price is driven to marginal cost with any number of sellers. Indeed the payoff from undercutting the other sellers increases with the number of firms in the market. For example if all N firms except one charge some price $p_0 > c$ then the profits of the remaining firm jump from $N^{-1}(p_0 - c)q$ to $(p_0 - dp - c)(q + q'dp)$ by charging a marginally lower price. In this sense, larger numbers of firms reinforce, but do not change, the prediction that price is equated with industry marginal cost.

Exercise *Here are some suggestions for experiments that vary the game specifications about product demand, supply costs, how much information producers have about consumer demand, and how information consumers have about the cost structure of the firms. In each of the following experiments, compare the average price of a sale, the sample variance of the prices, the total quanties sold, the how much of the gains from trade are exhausted.*

Matching Prices

Although by no means universal, some stores have a policy of matching prices on identical products available at rival stores. This policy effectively prevents one store from undercutting another to scoop out demanders when both have the same cost structure. Consequently each store considers the anticipated price reduction when offering discounts its own products. Thus in contrast to games where price competition dirves producers back to marginal costs the collective behavior of consumers unwittingly help producers to restrict ouptut and realize some monopoly rent by keeping each producer aware of his competitor’s prices. From the perspective of consumers this is another application of the prisoner’s dilemma we discussed in Chapter 6 on dominant strategies. If consumers could conspire to prevent stores from knowing the amount their competitors discounted, then their rivalry would generate marginal cost pricing.

To establish the proposition about matching prices, we consider a game in which both firms post limit order prices, consumers then choose whether to purchase the product or not, and if so, select a retailer. The last step is for consumers who have purchased from the most expensive retailer to seek the discount afforded by teh other supplier. The discount by a contractual obligation, which we assume is always honored. Alternatively the store refunds the customer for the full amount who is then free to purchase the good at its competitor for the advertised price.

As before, let $Q(p)$ denote the quantity demanded by consumers from a monopolist at price $p,$ and initially suppose the wholesale purchasing and marketing costs are c per unit for both retailers. Recognizing that they will receive the lowest price regardless of where they shop, consumers select the retailer at random, and we assume that each store attracts half the total demand. It follows that each store chooses a limit price recognizing that it either

$$\pi_j(p-c)Q(p)$$

Comparing this criterion function with the optimization problem a unifrom price monopoplist solves, the two functions only differ by teh factor of proportionality $\pi,$ and therefore generate teh same price solution. Therefore each store choosses the value maximizing price for the industry and shares the full monopoly rent in proportion to the share of custom they attract through random selection.

Similar to a trigger strategy in which everyone shares the rent.

Exercise *There are several ways of playing games where producers are forced to match prices. One mechanism is for producers to submit limit orders, customers to submit market orders, and for firms to automatically submit refunds, or retrospective discounts, on all items purchased at prices above the minimum limit sell order submitted to customers by either firm. Compare the outcome with the monopoly solution.*

Capacity

Rather than choosing price firms pick the quantity the intend to produceAnother way firms compete with each other isHow could we stop such debilitating competition?

Capacity commitments would be one way.

An alternative mechanism for selling the good is to have each customer approach the store of their choice and make an offer. After all the customers have sorted themselves into one or the other store, the firms respectively decide how many to serve and at what uniform price. Sealed bid pricing

Considering the monopoly problem of the previous lecture, let us now introduce a second seller with same marginal cost schedule, and no fixed costs.

As before we let $Q(p)$ denote the demand curve, and write $P(q)$ for its inverse, meaning $q = Q(P(q))$ for all positive $q.$ We let

$$q = \sum_{j=1}^{J} q_j$$

The k^{th} firm chooses q_k to maximize

$$P\left(\sum_{j=1}^{J} q_j\right)q_k - cq_k$$

The first order condition for an interior maximization is

$$P'\left(\sum_{j=1}^{J} q_j\right)q_k + P\left(\sum_{j=1}^{J} q_j\right) = c$$

In a symmetric equilibrium all firms produce teh same quantity, meaning $Jq_j^o = q^o$ for all $j \in \{1,\ldots,J\}$ and some positive $q^o.$ In that case teh equilibrium quantity produced by each firm can be solved from the equation

$$P'(q^o)q^o + JP(q^o) = Jc$$

The limiting properties of this expression are noteworthy. In teh case of monopoly $J = 1$ and the industry equilibrium quantity is simply teh monopolist's procudtion. We obtain the standard condition that marginal revenue is equated with marginal cost:

$$P'(q_1^o)q_1^o + P(q_1^o) = c$$

A natural question to ask is where this process would converge, and whether there is an easy way to model what would happen in the limit. We may deduce what happens to industry outuput as the number of firms increase by first dividing the equation determining output by the number of firms, and then taking the limit as J increases wtithout bound. Dividing the first order condition by J and rearranging, we see that

$$P(q_J^o) - c = \frac{P'(q_J^o)q_J^o}{J}$$

Taking the limit we now obtain the result that the the equilibrium quantity sold

$$\lim_{J\to\infty}\left\{\frac{P'(q^o)q^o}{J}\right\} = 0$$

For example, supposing the demand curve is linear in price, we may write

$$Q(p) = \alpha_0 - \alpha_1 p$$

in which case

$$P(q) = \frac{\alpha_0}{\alpha_1} - \frac{1}{\alpha_1}q$$
$$P'(q) = -\frac{1}{\alpha_1}$$

Substituting these expressions for $P'(q^o)$ and $P(q^o)$ into the solution for q^o we thus obtain

$$-\frac{1}{\alpha_1}q^o + J\left[\frac{\alpha_0}{\alpha_1} - \frac{1}{\alpha_1}q^o\right] = Jc$$

which simplifies to

$$J(\alpha_0 - c\alpha_1)/2 = q^o$$

Do our experiments suggest that the limit point depends on the cost structure?

Another question is how many firms are required to reach this limit (that is when it exists).

Again we suppose firms set prices, which consumers take as given.

How are things affected?

Exercise *This set of experiments explores different ways in which the competitive limit might be reached. Suppose there are K retailers selling identical products, and consumers shop between the outlets. Conduct experiments as K increases from one to five, reporting on price dispersion in the market, quantity sold,* *and the distribution of the valuations to consumers versus non consumers.*

1. *Suppose each outlet uses a simultaneous discriminatory sealed bid auction.*

2. *Suppose the retailers have a sequential strategy, using a sequential discriminatory auction.*

3. *Suppose each outlet allows consumers to post limit orders to buy which are visible to everyone before the order phase closes and each firm decides who to service.*

4. *Suppose firms post limit orders to sell and consumers submit market orders.*

Exercise *We consider variations on the games that are discussed in this section. What happens if one seller takes a pricing strategy approach and the other one sets a quantity.*

Do not play weakly dominated strategies

Sample test

1. Price competition

Firm1 had two choices: high price and low price. Firm two's choices were also high price and low price. Neither firm knew what the other firm decided when they were making the decision. This is an example of a simultaneous move game.

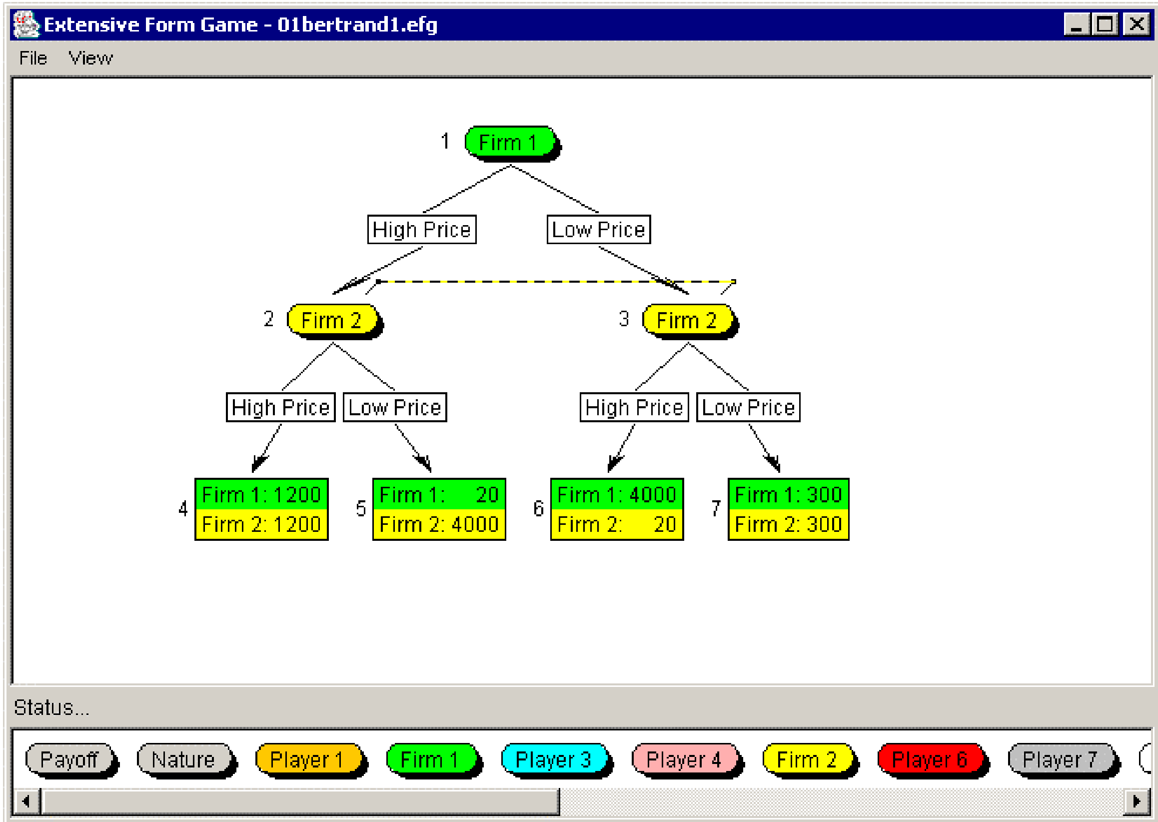

Comments and Solution

Firm 1 has a strictly dominant strategy to select "Low Price": 4000 > 1200 and 300 > 20 and similarly for Firm 2. Solution to this game is: (low price, low price) or node number 7.

The problem is graded in the following way: Each correct choice is one point. If one player selects low price and the other player selects high price, the player who selected a low price receives 1 point and the player who selected a high price receives 0 point for that round. The points are summed up and divided them by the number of potential points that could have been received in n rounds correctly.

2. Should homeowners own a gun?

In the United States many homeowners own guns for self-defense. In Britain almost no one owns a gun. Cultural differences provide one explanation. The possibility of strategic moves provides another. In both countries, a majority of homeowners prefer to live in an unarmed society. But they are willing to buy a gun if they have a reason to fear that criminals will be armed. Many criminals prefer to carry a gun as one of the tools of their trade. The matrix below suggests a possible ranking of outcomes.

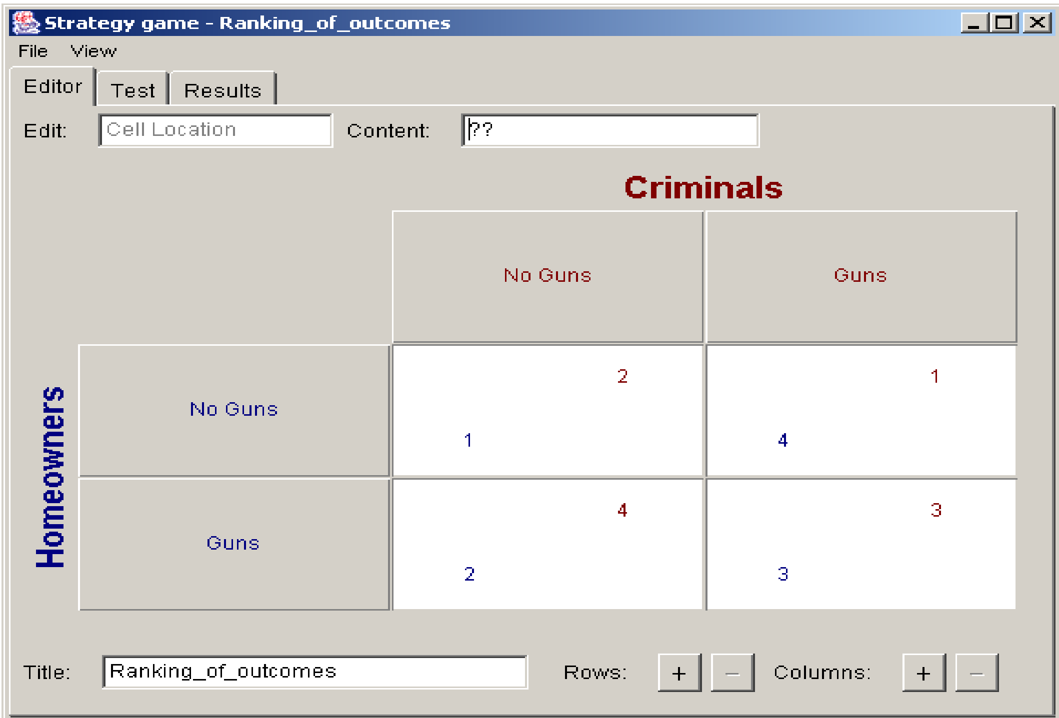

Comment:

Criminals have a strictly dominant strategy: Guns

Homeowners do not have a dominant strategy; they prefer to respond in kind. If criminals are unarmed, a gun is not needed for self-defense. What is the predicted outcome when we model the game as one with simultaneous move?

Solution:

Following Rule 2, we predict that the side with a dominant strategy uses it; the other side chooses its best response to the dominant strategy of its opponent. Since guns is the dominant strategy for criminals, this is their predicted course of action. Homeowners choose their best response to guns; they too will own a gun. The resulting equilibrium is ranked [3,3], the third best outcome for both parties.

The follow up discussion (not part of the test)

In spite of their conflicting interests, the two sides can agree on one thing. They both prefer the outcome in which neither side carries guns [1,2] to the case in which both sides are armed [3,3]. What strategic moves make this possible and how could it be credible?

3. Cournot market or common-pool resource dilemma

This game can be interpreted as either a Cournot market or common-pool resource dilemma. The main difference in interpretation is that in a common-pool problem, the joint maximum for two players is also a social optimum; but in the case of a Cournot market, the collusive joint profit maximum for sellers is harmful to consumers.

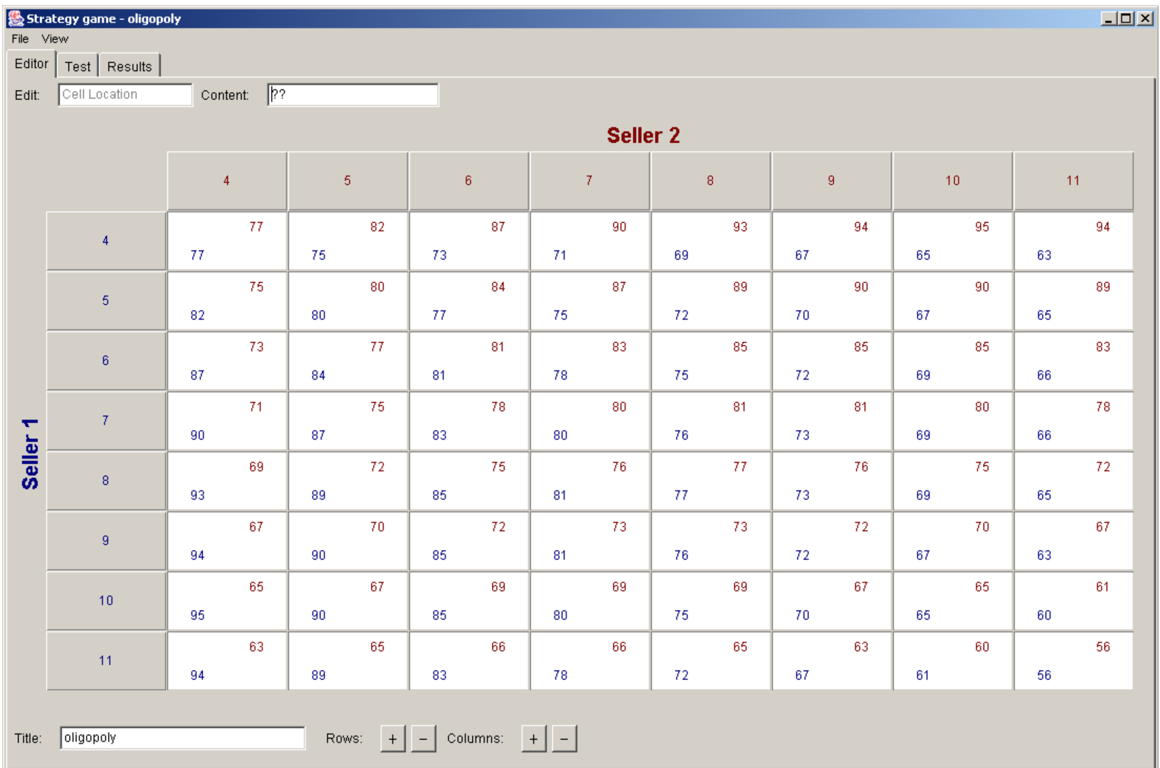

Seller 1 \ Seller 2	4	5	6	7	8	9	10	11
4	77 / 77	82 / 75	87 / 73	90 / 71	93 / 69	94 / 67	95 / 65	94 / 63
5	75 / 82	80 / 80	84 / 77	87 / 75	89 / 72	90 / 70	90 / 67	89 / 65
6	73 / 87	77 / 84	81 / 81	83 / 78	85 / 75	85 / 72	85 / 69	83 / 66
7	71 / 90	75 / 87	78 / 83	80 / 80	81 / 76	81 / 73	80 / 69	78 / 66
8	69 / 93	72 / 89	75 / 85	76 / 81	77 / 77	76 / 73	75 / 69	72 / 65
9	67 / 94	70 / 90	72 / 85	73 / 81	73 / 76	72 / 72	70 / 67	67 / 63
10	65 / 95	67 / 90	69 / 85	69 / 80	69 / 75	67 / 70	65 / 65	61 / 60
11	63 / 94	65 / 89	66 / 83	66 / 78	65 / 72	63 / 67	60 / 61	56 / 56

Solution

By iteratively removing weakly dominated strategies, both players end up offering q=8 or an effort level 8. Notice that if the column player chooses an effort level 8, then the best response of the row player is also 8, which is the only symmetric Nash equilibrium. The payoffs to each player are: [77,77].

The follow up discussion (not part of the test)

However, if the players could make binding agreement, they would both choose a common effort level of 6, which has a higher payoff of 81 for each.

4. Pricing strategies

Background:

The client (C) sold a highly profitable product that dominated its category, but was bracing for the introduction of a therapeutic substitute by another major pharmaceutical company. As a late mover, the entrant (E) was expected to launch its product at a very large discount despite its greater therapeutic benefits. It was unclear, however, exactly how low E's launch price would be and whether C should reduce its own price in anticipation or reaction. The cash flows involved were large enough to compel consideration of C's options.

The analysis began by specifying four options, involving different levels of discounting, for E's launch price. In addition, it identified four options for C's own (relative) price that were bracketed by the alternatives of holding C's price level constant and of neutralizing E's price advantage. Experts helped gouge the market share implications of each pair of prices. These market shares were then combined with knowledge of C's costs and estimates of E's costs to calculate the net present value of the tow products for their respective companies in the payoff matrix.

Client's (C's) Price \ Entrant's (E's) Price	Very Low	Low	Moderate	High
No Price Change	190.0 / 350.0	168.0 / 507.0	129.0 / 585.0	116.0 / 624.0
E has large price advantage	163.0 / 418.0	168.0 / 507.0	-99.99 / -99.99	-99.99 / -99.99
E has small price advantage	155.0 / 454.0	138.0 / 511.0	126.0 / 636.0	-99.99 / -99.99
C neutralizes E's advantage	50.0 / 428.0	124.0 / 504.0	129.0 / 585.0	128.0 / 669.0

Solution

1. Moderate weakly dominates High for Entrant, so E will never play High
2. Then "E has small price advantage" becomes strictly dominant strategy for C.
3. E responds with "Very Low".

The follow up discussion (not part of the test)

This schematic became the centerpiece of the pricing study. First it raised questions about the existing business plan, which assumed that E would launch with high price and that C would not change its price at all. The modeling reveals that this “base case” was a highly unlikely outcome.

This equilibrium point was highly unattractive to C’s managers who saw it as career-threatening. Instead, they began to explore whether they could change the game by credibly pre committing to a (relative) pricing strategy for their product….

5. A first price sealed bid auction with know valuations: Valuation 3 and Valuation 4.

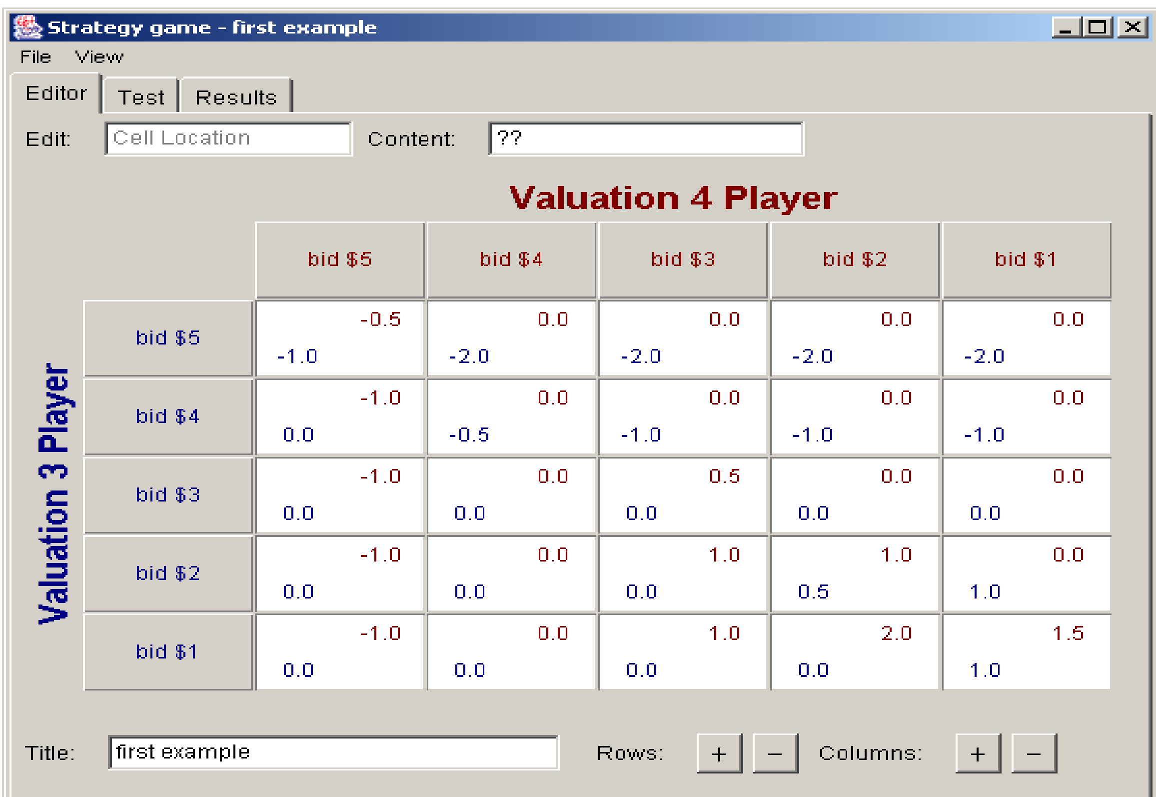

Valuation 3 Player \ Valuation 4 Player	bid $5	bid $4	bid $3	bid $2	bid $1
bid $5	-0.5 -1.0	0.0 -2.0	0.0 -2.0	0.0 -2.0	0.0 -2.0
bid $4	-1.0 0.0	0.0 -0.5	0.0 -1.0	0.0 -1.0	0.0 -1.0
bid $3	-1.0 0.0	0.0 0.0	0.5 0.0	0.0 0.0	0.0 0.0
bid $2	-1.0 0.0	0.0 0.0	1.0 0.0	1.0 0.5	0.0 1.0
bid $1	-1.0 0.0	0.0 0.0	1.0 0.0	2.0 0.0	1.5 1.0

Solution:

1. “Bid $4” weakly dominates “bid $5” for Valuation 4 player.
2. “Bid $3” weakly dominates “bid $4” for Valuation 4 player.
3. “Bid $2” weakly dominates any other bid for Valuation 2 player.
4. Valuation 4 Player will bid $3 and Valuation 3 Player will “bid $2” because “bid $3” weakly dominates “bid $2” for Valuation 4 player knowing that Valuation 3 player will not play “bid $1”.

Bibliography

Ahn, Toh-Kyeong, Elinor Ostrom, David Schmidt, Robert Shupp, and James Walker (1999) “Cooperation in PD Games: Fear, Greed, and History of Play,” Public Choice, forthcoming.

Andreoni, James, and Hal Varian (1999) “Pre-Play Contracting in the Prisoner’s Dilemma,” Proceedings of the National Academy of Sciences, 96(September), 10933-10938. Abstract: The prisoner’s dilemma is modified by adding an initial stage in which players can precommit to reward the other one for a cooperative decision in the second stage. This modification greatly increases the rate of cooperation.

Beckman, Steven R. (1990) “Producer’s Dilemma Experiments,” Journal of Economic Behavior and Organization, 27-46.

Beckman, Steven R. (1990) “Producer’s Dilemma Experiments: Reply,” Journal of Economic Behavior and Organization, 14:2 (October), 287.

Bohnet, Iris, and B. S. Frey (1999) “The Sound of Silence in Prisoner’s Dilemma and Dictator Games,” Journal of Economic Behavior and Organization, 38:1 (January), 43-57. Abstract: The paper examines the effects of relaxing the strict anonymity conditions that are standard in most prisoner’s dilemma and dictator experiments, Knowing the other player’s identity increases cooperative behavior, even in the absence of cooperation.

Bonacich, Phillip (1970) “Putting the Dilemma Back into Prisoner’s Dilemma,” Journal of Conflict Resolution, 14379-387.

Flood, Merrill M. (1954) “On Game-Learning Theory and Some Decision-Making Experiments,” in Decision Processes, edited by R. M. Thrall, C. H. Coombs and R. L. Davis, New York: Wiley, 139-158.

Grobelnik, Marko, Vesna Prasnikar, and Charles A. Holt (1999) “Classroom Games: Strategic Interaction on the Internet,” Journal of Economic Perspectives, 13:2 (Spring), 211-220.

Halpin, Stanley M., and Marc Pilisuk (1970) “Prediction and Choice in the Prisoner’s Dilemma,” Behavioral Science, 15141-153.

Holm, H. J. (1995) “The Prisoners’ Dilemma or the Jury’s Dilemma? A Popular Story with Dubious Name,” Journal of Institutional and Theoretical Economics, 151:4 (December), 699-702.

Pruitt, Dean G. (1970) “Motivational Processes in the Decomposed Prisoner’s Dilemma,” Journal of Personality and Social Psychology, 14227-238.

Schmidt, David, Robert Shupp, James Walker, and Elinor Ostrom (1999) “Playing Safe in Coordination Games: The Role of Risk Dominance, Social History, and Reputation,” Indiana University, Discussion Paper.

Straub, Paul G. (1995) “Risk Dominance and Coordination Failures in Static Games,” Quarterly Review of Economics and Finance, 35:4 (Winter), 339-363.