1 Introduction

The strategic and extensive forms can be described as canonical, because they are such useful pedagogic tools for understanding the structure of games, and for explaining the principles used to solve them. We used these two representations to define and solve finite games of manageable size. However the canonical forms are cumbersome when a game has many players who make simultaneous moves, big choice sets, chance events that have many possible outcomes, or long sequences of moves that might be defined recursively. This chapter departs from the strategic and extensive form representations to analyze the games that have elements of repetition, or more generally, recursion.

In the next section we show how to represent recursive games, by a sequence of stages and transition probabilities that link them together. The finite horizon games we present in Markov form in Section 2 illustrate three types of games we analyze in this chapter and the next. In repeated games and stage games, the transitions between rounds are exogenous so do not depend on the choices made within the stage. These two types of Markov games are investigated more extensively in the next chapter. The remaining sections of this chapter focus on Markov games in which current choices partly determine the stages follow or the duration of the game.

It is more parsimonious to define recursive games by presenting their stages coupled to their respective probability transitions, rather than using either canonical form. However the special structure of finite horizon Markov games does not facilitate the derivation of their solution. They are solved using the techniques exposited in Chapters 5 through12, and as such warrant no extra discussion in this chapter. But to solve infinite horizon Markov games, we can adapt recursive methods developed for dynamic programming models in Chapter 3. This is the second reason for using the Markov form. So although it is impossible to even present the strategic or extensive form of infinite horizon games, the special structure of the Markov form can be exploited to both present and solve them.

The solutions we search for belong to the class of Markov strategies, characterized by the property that all previous choices in the dynamic games affect current choices only indirectly, if at all, through the current state of play. Thus the defining feature of Markov solutions is that current moves do not depend on what happened in previous stages. Section 3 begins the discussion of Markov strategy solutions to infinite horizon Markov games with several applications.

Then in Section 4 we provide a formulationfor any infinite horizon Markov game with a finite number of stages and discuss some solution techniques. Deriving a solution is quite cumbersome for games with a large number of stages, choices and players, but checking whether experimental subjects are playing a solution is more straightforward. This raises the possibility of discovering the solution to Markov games through experimental methods, a prospect we explore in this section.

The techniques we develop can be adapted to games where both the choices and the stages are uncountable. We provide some examples in Section 5 to illustrate that point, before summarizing our discussion in Section 6.

The applications we investigate in this chapter include competition between firms in the form of research fro new products versus advertising the existing product line, patent races, timing new product releases, competitive entry into an industry, cultivating and degrading shared resources within the industry, and arms races between rival powers.

2 Finite Horizon Games

Any number of finite games for the same group of players can be joined together by a deterministic or a stochastic law of motion to indicate the order in which they are to be played. In this way we could define a new larger game, called a Markov game. Rather than present a Markov game in canonical form, it is more parsimonious to present the original games as stages, along with the stage transitions.

This broad definition encompasses repeated games and stage games. In repeated games there is only one stage, repeated at least once. More generally, stage games may be formed from two or more distinct games with the same number of players, so are a natural extension of repeated games. In both repeated games and stage games the choices players make do not affect the number of repetitions or the stage that will be visited in the future. Both are subsets of Markov games. Markov games also include those formed from games that linked together with laws of motion determined by player choices within the stages. The three examples below illustrates these points.

2.1 A repeated game

Repeated games are formed from an original game, called a stage, by assigning the same players in the same roles play the game once or more. Figure 13.1 is a template for a 2 player game repeated once in which each player chooses between one of two actions. The players collectively chooses one of the four cells in the center blocks of 4, and then move to the outside block indicated by the arrow to make their choice in the second period. The choices in the second period might depend on what happened in the first round even though the payoffs in each block of 4 must be identical (by the definition of a repeated game). This distinction lies at the heart of discussing repeated games.

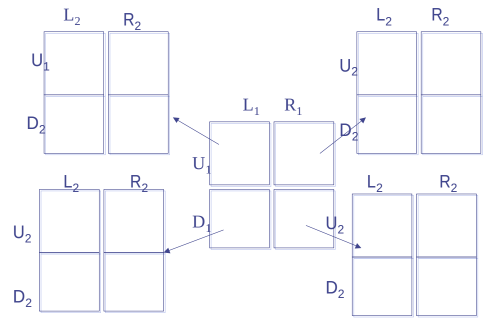

Figure 13.1

Template for a 2 period simultaneous move games for two players

Writing down the strategy space for repeated games is typically a very tedious task. To give the strategic form for the repeated game in Figure 13.1 the strategies for the second period are augmented to the strategies for the first. Instructions for the second period are contingent all that is observed up until the beginning of the period, and the event of racing in the second period. There are four possible profiles of choices in the first period, and this leads to either success or failure at the end of the first period. Hence there are 8 outcomes from the first period on which. Therefore there are two outcomes. Therefore there are 2^4 or 16 different instructions that could be given to each player at the beginning of the second period. However only half of them apply contingent on what happened in the first period. In summary a dimensional bi-matrix is required to state the strategic form of this repeated game. Hence there are 64 strategies in total. This discussion also suggests that the number of strategies and possible game histories proliferate at a much faster rate than the number of repetitions. For this reasons writing down the strategic form of the repeated game is a very cumbersome and unwieldy exercise for all but the simplest of games.

The extensive form for this game is partially constructed in Figure 13.2. The display shows there are 64 outcomes to compare. Clearly the extensive form of more complicated games than this simple game would be cumbersome to analyze.

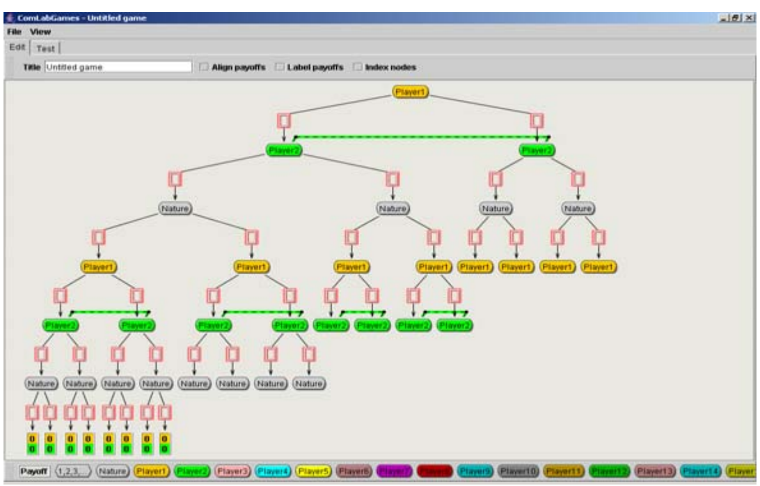

Figure 13.2
Extensive Form of Joint Venture

An alternative convention, that we shall adopt in chapter, is to only write down the extensive or normal form of a stage, along with an indication that it should be repeated once. See Figure 13.3.

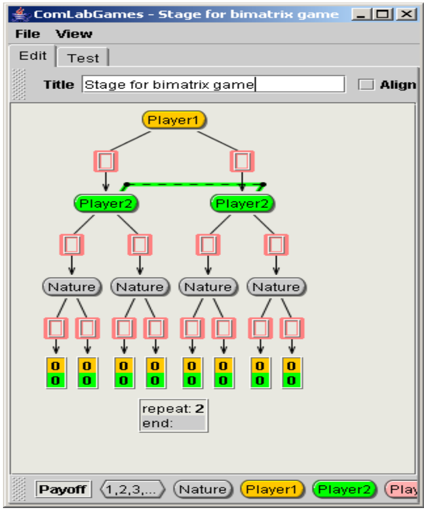

Figure 13.3
Stage for Joint Venture

2.2 A stage game

Repeated games are a special type of stage game, distinguished by the feature that they contain only one stage. Defining a repeated game amounts to defining the stage and a mechanism for when the game stops. Playing the stage once is called a round. Thus the number of rounds count the repetitions of the single stage. A repeated game might last for a fixed number of rounds, or be repeated indefinitely (perhaps ending with a random event). Whereas repeated games have only one stage, stage games may have many. Thus stage games encompass repeated games. The example is a two period game with three stages that closely resembles the repeated game described above. Indeed the first stage, played in the first period, is identical to the first period of the repeated game example, and the other two stages only differ in their payoffs. In the second stage all the payoffs are halved, while in the third stage they are all doubled. Figure 13.4 displays the game. In other wordsThe probability transition below the first stage indicates that after the first stage is played once, there is an equal chance the play will move either the second or the third stage in the second period.

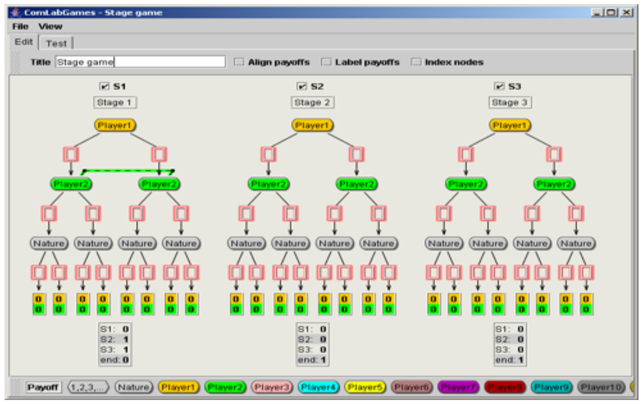

Figure 13.4
A Stage Game of joint Venture

2.3 Endogenous probability transitions

Since the probability of moving to the latter stages does not depend on teh choices that are made in the first period, one transition under the stage first suffices to characterize the law of motion between stages. This convention holds for all stage games, but not for Markov games in general. We now consider a second variation on the original game. We now suppose that if both partners cooperate the the probability business doubles, but if they

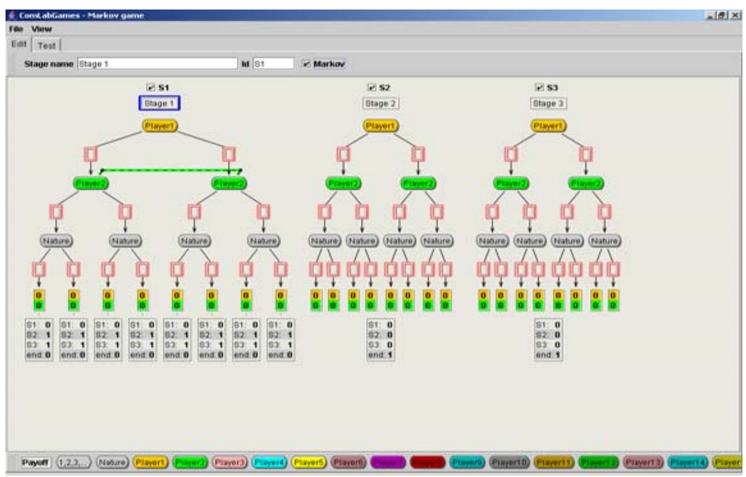

Figure 13.5
Markov game for partnership

2.4 Solving finite horizon games

The solution methods of Chapters 5 through 12 can be used to solve these finite horizon games too.

An experiment was conducted to see whether the outcomes

3 Infinite Horizon Games

We begin by this section with some examples of Markov games, and derive a subset of their solutions, strategies that satisfy all the conditions we require do solutions and in addition have limited history dependence which we define and explain below.

Recall from Chapter 6 that, in games with perfect information, each information set is a singleton. Thus every move begins a new subgame, and can be considered as a separate stage. This section extends our analysis in Chapter 6 to infinite horizon games with perfect information. To solve infinite horizon perfect information games we adapt the techniques developed for the dynamic programming models discussed in Chapter 3 on investment. Beginning with two examples, we then provide a general treatment for discrete games.

3.1 Joint production

The working life of office equipment depends on how well it is serviced and the way it is used. Employees should be alert to malfuntioning to avert potential serious problems, and should take care not to damage the equipment by Each player has the choice of taking more of less care of the office equipment. Taking more care typically requires more concentration and effort than taking less care, thus diverting their attention away from other tasks within the firm and also their leisure activities. For this reason, the immediate benefit from maintaining the equipment is lower than the benefit from mistreating it. This long term cost associated with mistreatment is the higher risk that the equipment will fail, but in this game that replacement cost has not been internalized by the workers, who are not held accountable for their actions.

The game captures these features in a simple manner. It is a one stage game repeated several times, depending on the choices the two players make. Each round the players simultaneously decide whether to maintain or mistreat the machine (when they use it that day, say). The payoff from mistreating it is 5, but the payoff from maintaining it is only 3.

Responsibility for maintaining and investing resources that are shared between divisions, or workers within the same office, are hard for central management to value, a problem to which we will return in Chapter 15. At the macro level, the difficulty of eliciting preferences for common projects and coordinating activities between different divisions takes on larger dimensions, as issues about cross subsidies, control over resources used by several divisions, and translating the strategic vision of the firm into operational plans are fiercely debated between competing divisional managers.

Consider for example a service which is jointly produced between two division of a division

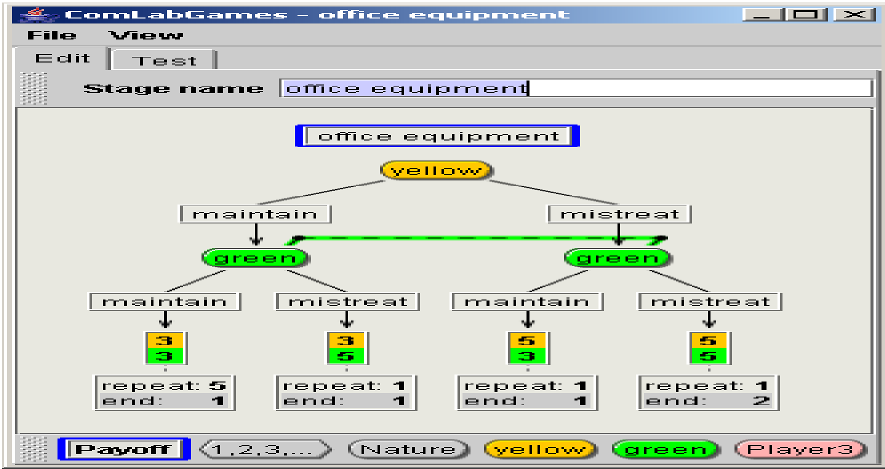

Figure12.8
Office Equipment

As an intermediate step to deriving the unique Markov solution to this game, we first determine the value of the office equipment when both players cooperate by maximizing the sum of the discounted benefits from using the machine. The solution to this simple dynamic optimization problem is found by comparing four benefit streams. If both players mistreat the machine, together they receive 10 with probability 3^{1-t} in the t^{th} round of the game. Summing the geometrically declining sequence over $t \in \{1, 2, \dots\}$ we obtain

$$10 \sum_{t=1}^{\infty} 3^{1-t} = \frac{10}{1 - 1/3} = 15$$

If both players maintained the machine, then together they would receive only 6 in the t^{th} round of the game if it was still functioning, but the probability of survival is higher, $(5/6)^{t-1}$. Computing the discounted sum we obtain 36. Using the same method we find the value from one player mistreating the machine and the other maintaining it is 16. Thus if both players are jointly maximizing the value from the it is optimal treatment for the machine, it is optimal for both players to maintain it.

The noncooperative solution to this game is derived from these four numbers. If both players maintain the machine, they split the expected value evenly for an expected value of 18. The expected value from maintain If one player maintains the machine and the other mistreats it, the expected value to the former is 6 and the the expected value to the latter is 10. Finally the expected value to each player when both mistreat it is 7.5. Figure 12.9 displays the Markov strategic form, a bi-matrix showing the expected payoffs from the player pursuing the different combinations of pure Markov strategies.

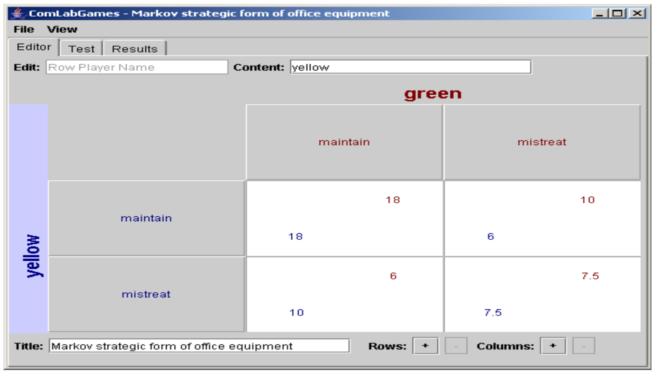

Figure 12.9

Markov strategic form of office equipment game

It is evident from the bi-matrix that mistreat is a dominant strategy (once we pull down the value of maintenance). Hence the unique Markov solution to this game is (mistreat,). We remark that since neither player uses the history of the game to formulate his strategy, it is optimal for the other player to ignore the history too. This proves that the Markov strategy is a solution to the larger game. Although it is clearly the unique Markov solution, it is not the only solution to the game. For example a trigger strategy solution of playing maintain unless someone has already played mistreat is also a solution to this game.

The variation in the solution outcome depends on the Do we buy better or worse machines? We now discuss the virtues of buying high maintenance equipment versus equipment that is fail-safe.

3.2 Advertising versus research

Being awarded a patent race can be modeled as a two stage process, developing the new drug and s cna be regarded as having two stagesBeyer and Merck compete against each other to develop a new treatment for headaches. There are three stages to this game, each of which might be visited an indefinite number of times. The line of reasoning justifies eliminating the choice of either firm to discard a drug that it has just developed in favor of seeking approval from the FDA. The Markov form of the reduced game is displayed in figure 13.8.

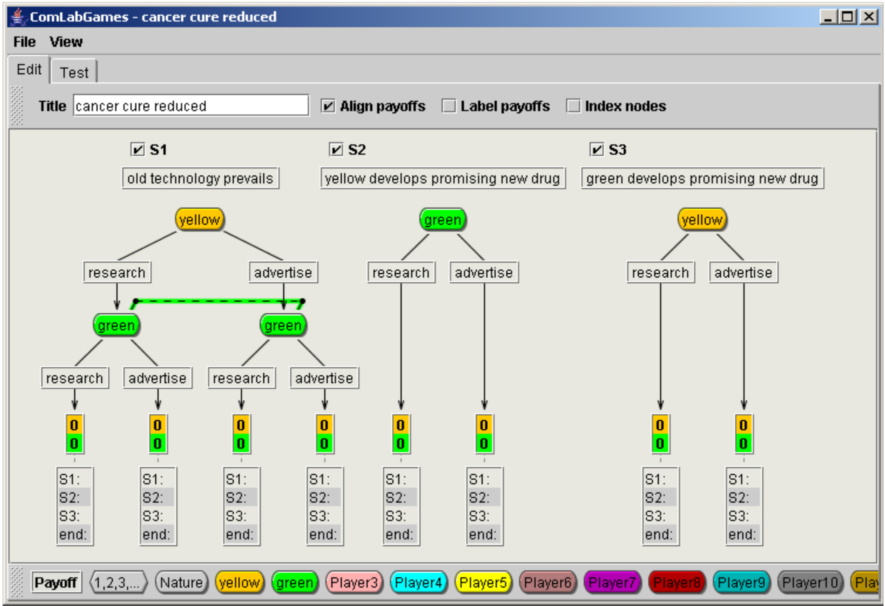

Figure 12.1

Markov Form of Pain Relief Game

In the Markov form of this game, each firm has two information sets with two choices at each, implying a total of four strategies, which we display with their associated payoffs in figure 13.9.

Figure 12.2

Markov Strategic Form for Pain Relief Race

The Markov solution to this game can now be derived in a straightforward manner.

3.3 Outsourcing

Firms periodically face the question of how far back to integrate their activities versus how much to contract out to suppliers. Outsourcing can lead to increased uncertainty about delivery times and the quality of the goods and services produced, but undertaking the activities internally creates organizational challenges and may lead the firm away from areas where it has a comparitive advantage. How much outsourcing to do depends on many factors, and may vary with economic conditions, including the demand for the firm’s main product. For example, in periods of low product demand the costs of diverting workers from the firm’s core activities to producing input components that could be purchased from an ouside supplier might be lower than when demand for its product is high.

Figure 13.6 depicts the Markov form of a game between a manufacturer and a component supplier. The cost of outsourcing is independent of the demand for the firm’s product, but the cost of building the component is less when demand is low. Demand is determined endogenously. Conditional on product quality, a period of high demand is more likely to follow if there was high demand in the preceding period than if there was low period. Products made with high quality components stimulat demand,

and products made with low quality components repress it. The probability of the game ends after the latest round is also decreasing in component quality, reflecting an increased risk of bankruptcy and liquidation

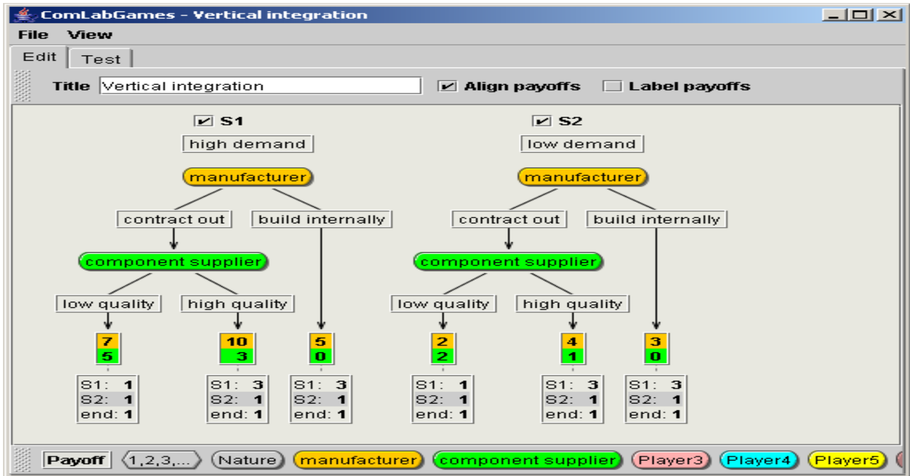

Figure 13.6

Outsourcing

The current value of the remainder of the game to the two players depends on whether they are in the high demand or the low demand stage, and can be simply computed as a function of their pure Markov strategies. Each player has four. If the manufacturer always builds internally, then the component supplier never receives any contracts and the value of the manufacturing firm is

Now suppose the manager always outsources.

There are 64 payoffs to compute, 32 for each player, because at each stage the payoffs depend on how both players behave at the other stage. These are depicted in Tables 13.1 and 13.2.

High demand stage

		supplier	
		sturdy	*poor*
manufacturer	*build*	[ ]	[ ]
	outsource	[ ]	[ ]

Table 13.1

Pure Strategy Payoffs for High demand stage

To interpret this figure, note that . Figure 13 is computed in a similar way.

Low demand stage

supplier

manufacturer		*sturdy*	*poor*
	build		
	outsource		

Figure 13
Pure Strategy Payoffs for Low demand stage

A systematic way of finding any pure solution strategies is to check for conformity that a solution imposes.

3.4 Supply chain management

Companies often ally themselves with partners in vertically integrated chains. For example in the automobile industry dealers, manufacturers and their components suppliers rely on on each other for business. The following example captures the nature of competition between rival suppliers to a manufacturing firm. A supplier contracts with the manufacturer on many of the product specifications, but not all, and in this example the flexibiity of delivery schedule is left unspecified. We assume that delivery is either tardy or punctual. At the end of each period the manufacturer decides between recontracting with his existing supplier versus approaching an alternative source. If the manufacturer seeks to change suppliers, there are setup costs of $1 million and he is only successful three quarters of the time. Attempting to switch costs chooses bete

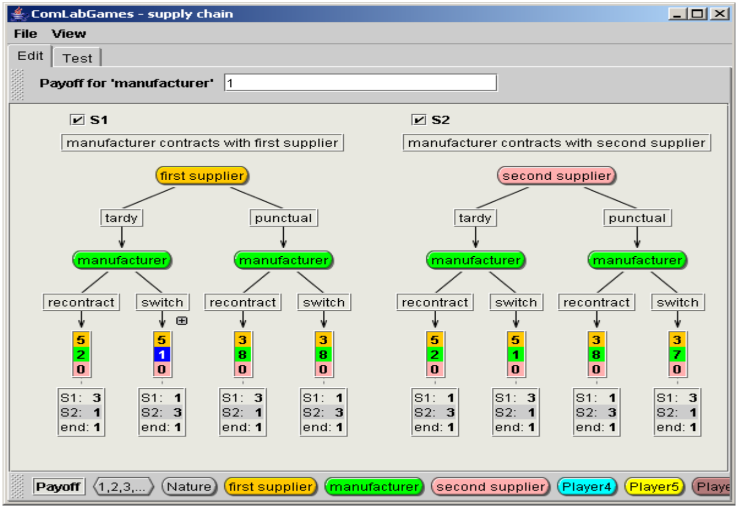

Figure 13.6
Supply Chain

In this game there are two strategies for the first supplier (trady or punctual), the same two for the second, and sixteen for the manufacturer. Since the costs and benefits do not depend on which supplier the manufacturer contracts with, we assume the manufacturer adopts a symmetric strategy with respect to the suppliers. There are four such strategies (always recontract, always switch, recontract if punctual but switch if tardy, switch if punctual but recontract if tardy).

for this game are found by computing the values

3.5 Product cycle

Year after year two hiking boot manufacturers, Alpina and Planetka, compete against each other for footwear demand from hikers trekking through the alps, by developing and marketing boots that are lighter, tougher and more comfortable. During the summer a consensus forms amongst hikers about which boot is superior. Over the winter months the manufacturers simultaneously decide whether to produce a new model or not in preparation for the next hiking season. Developing a new boot is costly, but the market rewards relative quality, and this stimulates new boot designs.

Figure 13.9 displays the essence of this rivalry. After one season of competition, there are two states of the world; either Planetka's boot is superior or Alpina's boot is superior. The transition probability to the next state is determined by three factors, which firm had the most popular boot in the current season, and who innovated. For example if neither firm develops a new boot, then the manufacturer which previously had the best boot will certainly retain its competitive advantage. Similarly if both firms develop a new boot it is more likely that the manufacturer which previously had the best boot will retain its edge. But supposing the manufacturer with less popular boot develops a new model, then the probability of it winning the quality contest next year increases if its rival does not develop a new model too. The game starts when nature determines which manufacturer initially has a better boot.

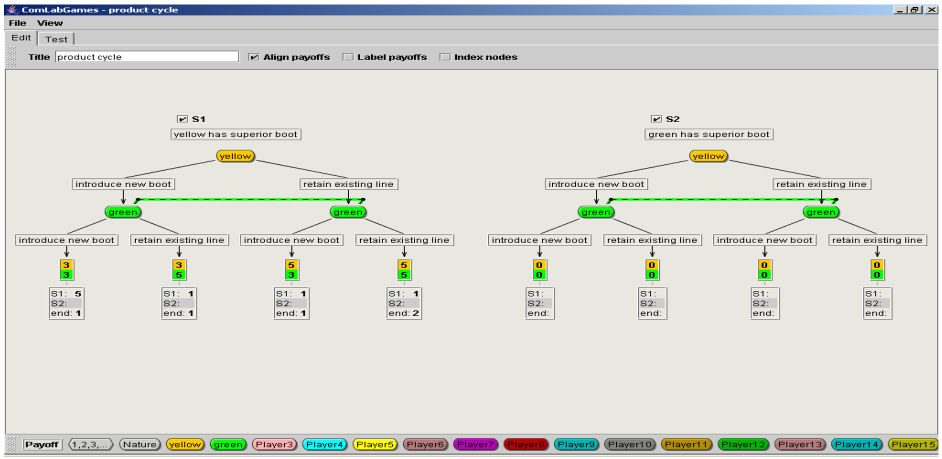

Figure 13.9

Product Cycle

When players make a move within a stage they recognize the expected payoff is partly determined by their move in the other stage. This observation is teh key to investigating the solutions to this game. Suppose for example that Alpina currently has a superior boot, and that when Planeka has the superior boot, both manufacturers (for whatever reason) retain their respective product lines. We now compute the expected present values of both firms. In this case the equilibrium for this part of the game is that. Can this be an equilbrium? We now ask . . .

A systematic way of finidng all the solutions to this game, that is both pure and mixed strategies is to compute the various payoff senarios that are possible from palying different combinations of moves. the payoffs are shown in Figure 13. below.

Alpina has the superior boot

Stage 1		*Alpina*	
		introduce	*retain*
Planeka	*introduce*	[ ]	[ ]
	retain	[ ]	[ ]

Figure 13.

Pure Strategy Payoffs for Stage 1

To interpret this figure, note that . Figure 13 is computed in a similar way.

Planeka has the superior boot

Stage 2

		Alpina	
		introduce	*retain*
Planeka	*introduce*		
	retain		

Figure 13
Pure Strategy Payoffs for Stage 2

We now check for conformity that a solution imposes

Having computed the pure strategy solutions for this game, we now investigate the possibilities for mixed strategy solutions. There are three to consider. Either both firms mix in both stages, or both mix in one stage, the first or the second. If mixing is occurring in both stages there are four equations in four unknown probabilities. If mixing is only occurring in the first stage, then it must be compatible with a pure strategy in teh second stage, and vice versa in the other case.

Each player has four Markov strategies, to always introduce a new product, to never introduce one, to introduce a new product only when the other company is acclaimed to have a better boot, and to introduce a new boot only when the other company has a worse boot. The Markov strategic form is depicted in Figure 13.10.

4 A General Framework

For the sake of completeness we now provide a general framework for considering infinite horizon Markov games where the perfect information, a finite number of stages, and a finite number of choices at each stage. We label each period of play by $t \in \{1,2,\ldots\}$, the N players by $n \in \{1,\ldots,N\}$, the S stages by $s \in \{1,\ldots,S\}$, and denote by s_t the stage visited in period t. Without loss of generality we suppose that at each stage one player makes a single move. (If the representation does not have this property, the stages where multiple players move can be split.) Thus each stage is a choice node for one player, and we denote by K_s the size of the choice set at stage s. The particular choice selected helps determine the next stage that is visited, or alternatively whether play ends at that point.

Given a choice $k_s \in \{1,\ldots,K_s\}$ assigned for every stage $s \in \{1,\ldots,S\}$, we can define a probability transition matrix that stochastically determines how the game history unfolds. The S rows of the transition matrix represent the current stage, the S columns the stage next period, and the (i,j) element in the matrix, denoted p_{ijk}, is the probability that the j^{th} stage will directly follow the i^{th} when choice $k \in \{1,\ldots,K_i\}$is selected at node i.

Players are motivated by payoffs additively acruing to them throughout the course of play, writing $u_n(s_t)$ for the current payoff to player n when stage s is reached at time t and choice k_s is selected by the player assigned to making choices at that stage. Following the assumptions made in Chapters 3 and 4 we assume each player obeys the expected utility hypothesis. In that case we can represent the n^{th} player's preferences by the expected utility formulation:

$$E_0\left[\sum_{t=0}^{\infty} \beta^t u_n(s_t, k_s)\right]$$

where β has the same interpretation as we gave it in Chapter 3, as a discount parameter, or as an exogenous probability (compounded by the probability of transiting to s_0 where appropriate) that the game ends after each round. We can express this objective function making explicit use of the probabiilty transition matrix P

Since Markov strategies are defined without reference to the history of play, we can define a Markov strategy for the game by a choice $k_s \in \{1, \ldots, K_s\}$ for each stage $s \in \{1, \ldots, S\}$. A Markov solution $\{k_1^e, \ldots, k_S^e\}$ is a Markov strategy from which no player can profitably deviate when it is his turn to move.

The existence of a solution is found by taking the limit of finite horizion games. Using the same argument we made in Chapter 3, they converge, so that for T sufficiently large, the first move made at any given intial state does not change by expanding the horizion by one more period. We then show teh vector of first moves constitutes a Markov strategy solution to the infinite horizon game.

Similarly, a variation of the contraction mappping algorithm can also applied to Markov solutions to infinite horizon games with perfect information.

4.1 Notation

Then we lay out the terminology and define Markov games.

This section develops a representation that models the choice set and the transition probabilites as a Markov process. The defining characteristic of a Markov process is that a finite dimensional vector defines the probabilty distribution. Much of the section is devoted to illustrating and defining a natural extension of the representation we developed for analyzing stage games, to broader classes of games. called Markov games. Its main components are the phases of the game that histories might pass through, the moves that occur within phases taken by players, the payoffs that occur at the end of each phase, the transition probabilities that govern phases, and an auxiliary vector of state variables that are arguments in the payoffs and the transition probabilities. As we demonstrate, each of the examples can be parsimoniously presented within this form. We also discuss how the solution concepts developed in the previous chapters can be applied to Markov games, before turning to focus upon Markov solutions, strategy profiles that recursively solve the game. Solving these games can require the application of sophisticated nymerical algorithms. We also provide simpler, partial characterizations of the solution that can be checked in experimental play.

4.2 Markov solutions

Much of the terminolgy for prefect information games carries over to the imperfect information case. we label time, the players, the stages, the probability transitions, and utility functions exactly the same way. There are only two essential and related differences taht arise from the fact that not all the information sets within an imperfect information game are singletons. are not is that wtihin a stage The mechanism for selecting which stage to select is defined by matrix.

This assumption means that all differences throughtout the game between evaluating come through the state and the choices. Finally we assume that the player obeys the expected utility hypothesis.

Consider the strategy space in each game S_p. A Markov strategy is a strategy for each of teh P phases $(s_{1n},\ldots,s_{Pn}) \in S_1 \times\ldots\times S_P$ such that whenever the p^{th} phase is reached the n^{th} player invariably selects s_{pn}. The significance of a Markov strategy is that the player selects the same strategy regardless of the game history that has led to this state of affairs. In a single player game, that is a decision theoretic problem of the type studied in the previous chpater, is a that is played the same

Through this section we will focus on solutions that have special properties. First they are symmetric, meaning that if two players confront the same set of payoffs, we only look at solutions in which they make the same choices. This implies that the solution concept itself is not the source of differences in behavior. By a Markov solution we mean that the only source of differences in a behavioral response is attributable to diffences in a finite vector of state variables.

Denote by $s^{(n)} = \left(s_1^{(n)},\ldots,s_J^{(n)}\right)$ a Markov strategy of the n^{th} player, and let $s^{(-n)} = \left(s_{\tilde{n}1},\ldots,s_J^{(-n)}\right)$ denote strategies for the other $N-1$ other players $m \in \{1,\ldots,n-1,n+1,\ldots,N\}$. Let $p_{ij}(s_i)$ denote the probability that play transits from stage i to stage j when $s_i = \left(s_i^{(1)},\ldots,s_i^{(N)}\right)$ is played in stage i. Also define the probability transition mastrix as

$$P(s) = \begin{bmatrix} p_{11}(s_1) & \cdots & p_{J1}(s_J) \\ \vdots & \ddots & \vdots \\ p_{1J}(s_1) & \ldots & p_{JJ}(s_J) \end{bmatrix}$$

Noting that $1\{k=j\}$ is an indicator function for stage j, taking a value of one if $k=j$ and zero otherwise,Then the probability that that play is in the k^{th} stage at date t after starting at the first stage is

$$\begin{bmatrix} p_{11}(s_1) & \cdots & p_{1J}(s_1) \end{bmatrix} P(s)^t \begin{bmatrix} 1\{k=1\} \\ \vdots \\ 1\{k=J\} \end{bmatrix}$$

Note that the last expression formed from the indicator functions simply picks out the k^{th} column of the exponentiated matrix $P(s)^t$. The payoff to the n^{th} player at time t if play has moved to stage j is $\beta^t u_n(s_j)$. Therefore the

$$E\left[\sum_{t=0}^{T} \beta^t u_{nt}\right] = u_n(s_1) + \left[\begin{array}{ccc} p_{11}(s_1) & \cdots & p_{1J}(s_1) \end{array}\right] \sum_{t=0}^{T-1} \beta^{t+1} P(s)^t \left[\begin{array}{c} u_n(s_1) \\ \vdots \\ u_n(s_J) \end{array}\right]$$

In the case of an infinite horizon problem, the formula for an infinite geometric series implies

$$E\left[\sum_{t=0}^{T} \beta^t u_{nt}\right] = u_n(s_1) + \beta\left[\begin{array}{ccc} p_{11}(s_1) & \cdots & p_{1J}(s_1) \end{array}\right] [1_{JJ} - \beta P(s)]^{-1} \left[\begin{array}{c} u_n(s_1) \\ \vdots \\ u_n(s_J) \end{array}\right]$$

where 1_{JJ} denotes the $J \times J$ identity matrix.

Dominance relations and Nash equilibrium can be defined for teh Markov strategies in the same way as we defined these concepts in Chapters 8 though 10. For example a Nash equilibrium is defined as a strategy $s^e = (s_n^e, s_{\tilde{n}}^e)$ such that s_n^e maximizes the expected utility of player n from the game when the other players collectively choose $s_{\tilde{n}}^e$, where $s_m^e \in S_m$ for all $m \in \{1, \ldots, N\}$. When we speak of dominance amongst Markov strategies, the comparison set is with Markov strategies alone, but in a Nash equilbrium the best response over all strategies is a Markove strategy. Consequently the maximization is over a much bigger set of all history dependent strategies $h_n \in H_n$.

$$s_{nj}^e = \underset{s_{nj} \in S_{nj}}{\arg\max} \left\{ u_n\left(s_{nj}, s_{\tilde{n}j}^e\right) + \left[\begin{array}{ccc} p_{j1}\left(s_{nj}, s_{\tilde{n}j}^e\right) & \cdots & p_{jJ}\left(s_{nj}, s_{\tilde{n}j}^e\right) \end{array}\right] \sum_{t=0}^{T-1} \beta^{t+1} P(s_n, s_{\tilde{n}}^e)^t \left[\begin{array}{c} u_n(s_{n1}^e, s_{\tilde{n}1}^e) \\ \vdots \\ u_n\left(s_{n,j-1}^e, s_{\tilde{n},j-1}^e\right) \\ u_n\left(s_{nj}, s_{\tilde{n}j}^e\right) \\ u_n\left(s_{n,j+1}^e, s_{\tilde{n},j+1}^e\right) \\ \vdots \\ u_n(s_{nJ}^e, s_{\tilde{n}J}^e) \end{array}\right]\right.$$

In games with many stages numerical procedures are typically required to find the set of Markov equilibrium.

4.3 Ergodicity and the probability transition matrix

Here we discuss the ergodic set and absorbing stages.

The refinements of subgame perfection can also be applied to Markov games, and afforded byas we have seen from chpater . simplificationsWe define the ergodic sets of the game as groups of stages that play never moves out of, once entering. For example the second and third stages in the pain reliever game are both ergodic sets. More generally, a group of stages, say $k \in \{K, \ldots, N\}$ is an ergodic set if we can reorder teh stage so that there is a lower triangular block of zeros

$$P(s) = \begin{bmatrix} p_{11}(s_1) & & \cdots & & p_{J1}(s_J) \\ \vdots & & & & \\ & & \ddots & & \vdots \\ & & 0 & & \\ p_{1J}(s_1) & \ldots & 0 & 0 & p_{JJ}(s_J) \end{bmatrix}$$

Having identified the ergodic sets, Transition probabilities between the stages: solve for ergodic set: substitute back into previous stages

5 Continuous State Spaces and Choice Sets

Convenient to extend space to deal with choice sets that are closed intervals

The transition probability is a conditional probability distribution which generates next period's the random variable p_{t+1} conditional on the phase, the vector of supplementary state variables, and the choice of the player in period t. That is p_{t+1} is a random variable generated by the distribution conditional on (p_t, c_t). We denote this conditional distribution by $F(p|p_t, c_t)$. The state of play at each point during teh game is vector of variables labeled by $s \in S$, where S might be a finite set, a proability simplex, or a Euclidean space. Thus the state of play at time t can always be described by $s \in S$. Suppose the game is in a state s. Then the choice set may by written as $C(s)$. At that state of play in the game, the decision maker chooses some $c \in C(s)$. The transition from one period to the next determines how the value of the state changes. For convenience let us now subscript the state of play and the choice by t to indicate the period. Accordingly let s_t denote the state in period t, and c_t the choice made in t. The transition probability is a conditional probability distribution which generates next period's state variable s_{t+1} conditional on the state variables, and the choice of the player in period t. That is s_{t+1} is a random variable generated by the distribution conditional on (s_t, c_t). We denote this conditional distribution by $F(s|s_t, c_t)$. Without loss of generality we may define T by any state that repeats itself and at which no choices are made, and call this a terminal state.

while this class of solutions is more manageble than a more general approach, the computational complexity of locating Markov solutions are quite onerous.

In the remainder of this chapter we present several applilcations of Markov games and and analyze their equilibrium.

5.1 Entry

Consider an industry where there N potential entrants. Each firm $n \in \{1, \ldots, N\}$ has one opportunity to enter the industry at the time s_n it chooses. To simplify the analysis, we ignore exit from teh industry. The cost of entering teh industry is γ_n, an independent random variable continously distributed between $\underline{\gamma}$ and $\overline{\gamma}$ with probability distribution function $F(\gamma)$ where $F(\underline{\gamma}) = 0$ and $F(\overline{\gamma}) = 1$. Net revenue per period to each producer depends inversely on the number of firms in the industry, N_t, and the state of demand, D_t. We assume that each firm in the industry nets D_t/N_t. The stochastic process for demand is the autoregressive process:

$$D_t = \rho D_{t-1} + \varepsilon_t$$

where ε_t is an independently distributed random variable and $0 < \rho < 1$. This formulation implies that the expected value of demand in a future period converges to:

$$D_t = \rho D_{t-1} + \varepsilon_t$$

$=$ or alternatively $\varepsilon_t + \rho\varepsilon_{t-1}$ or Weiner?

This is an incomplete information Markov game. Firms know their own entry cost parameter, the number of entrants in each period, but do not know the entry costs of any other firm. At the beginning of the game the objective function of the n^{th} firm is

$$E_0\left[-\gamma_n \beta^{s_n} + \sum_{t=s_n}^{\infty} \beta^t D_t N_t^{-1}\right]$$

Each period the firms keep track of the state of demand, the number of firms that have yet to enter the industry. The n^{th} firm sequently chooses the entry date s_n and exit date t_n to maximize its value, knowing the number of firms in teh industry already N_t, their own entry cost γ_n, and the state of demand d_t. In a symmetric Markov pure strategy solution, differences in γ_n fully explain why firms do not enter simultaneously.

The last remaining firm to enter chooses a minimal level of demand that determines when it will enter. If it enters at period t its expected current value at that time is

$$E_t\left[N^{-1} \sum_{s=t}^{\infty} \beta^{s-t} D_t\right] - \gamma_n = E_t\left[N^{-1} \sum_{s=0}^{\infty} \beta^s \rho^s D_t + \ldots + \sum_{s=0}^{\infty} \beta^s \rho^s \varepsilon_{t+s}\right] - \gamma_n = \frac{D_t}{1-\beta\rho} - \gamma_n$$

Let $v(D)$ denote the value of the problem to the firm, and denote teh threshhold level of demand by D_n. There is a value D_n such that entering now is as good as entering the next time demand hits some level of demand at least as good as D_n. It faces an optimal stopping problem of the type analyzed in Chapter 3. Then for all $D > D_n$ it immediately follows that

$$v(D) = \frac{D}{1-\beta\rho} - \gamma_n$$

We see that

$$v(D_n) = \frac{D_n}{1-\beta\rho} - \gamma_n = \beta\left[\Pr(D > D_n) E\left(\frac{D_n}{1-\beta\rho} - \gamma_n | D > D_n\right) + \Pr(D \le D_n) E(v(D) | D \le D_n)\right]$$

Recalling our analysis of search in Chapter 3 we can solve for $v(D)$ to obtain D_n.

Consider now at the beginning of the game and in all priods before either firm has entered the industry. We focus on tow cases. If each firm knows the costs of both, then the order of entry is determined.

5.2 Patent Race

We now consider a game between several firms to develop a cure to a disease. Analogous to the dynamic optimization problem in Chapter 12, the probability of discovery in any one period by the n^{th} firm is denoted by $\lambda(k_{nt})$ where

$$\Pr[c_{nt} = 1 | c_{t-1} = 0] = \lambda(k_{nt})$$

where

$$k_{nt} = \delta k_{n,t-1} + q_{nt}$$

Suppose the firm reaps a benefit of v when the cure is discovered. Then the present value of the firm is

$$\sum_{t=1}^{\infty} \beta^t \Big[\lambda(k_{nt}) \prod_{s=1}^{t-1} \Big(1 - \sum_{n'=1}^{N} \lambda(k_{n's})\Big) v - c(q_t) \Big]$$

Following the approach established in the previous sections of this chapter, we consider a finite horizon problem with terminal date T, derive an approximate solution to this finite horizon analogue, then iterate on the policy or value function, and last bound the size of the approximation error to within tolerable limits.

Consider a one period problem. The firm maximizes

$$\lambda(\delta k_{t-1} + q_t)v - c(q_t)$$

which has a first order condition

$$v\lambda'(\delta k_{t-1} + q_t) = c'(q_t)$$

Let $\bar{q}_1$ solve $v\lambda'(q) = c'(q)$. Then solving for q_t we obtain

$$q_1^o = \max\{\bar{q}_1 - \delta k, 0\}$$

and the one period valuation function is

$$v_1(k) = \lambda(\delta k + \max\{\bar{q}_1 - \delta k, 0\})v - c(\max\{\bar{q}_1 - \delta k, 0\})$$

To find the steady state find the k such that

$$\delta k = \arg\max_q \Big\{ \sum_{t=1}^{\infty} \beta^t \Big[\lambda(k_t) \prod_{s=1}^{t-1} (1 - \lambda(k_s))v - c(q_t) \Big] \Big\}$$

5.3 Detente

Suppose there are two countries who get opportunities to dominate or fight the other one. If a country declares war, then the other one can fight or sue for peace. If both countries fight, then a probability distribution determines which one wins. Let k_{1t} denote the size of the war force of the first power and k_{2t} the size of the second force. These two variables follow the law of motion

$$k_{2,t+1} = \gamma k_{2t} + f(i_{2t})$$

Military build up is achieved at the cost of forgone consumption

$$\sum_{t=1}^{s} \beta^t (y_t - i_{2t}) + \beta^s d_{jt}$$

The loser in a war receives nothing from the time hositilites end. The payoff to the winner depends on who declares war. If winner had been the sole power to declare war, then its payyoff is w_2, if both parties had declared war, then the winner's payoff is w_1, but if only the enemy country had declared war, then the winner receives w_0, where we assume $w_2 > w_1 > w_0$. That is, in the event of winning, striking the other power unprepared is more profitable than if they make their own delcaration, and if the winner prevails over the other power who had the initial advantage, the spoils of war are less.

In the event of armed conflict, the probability of victory depends on the relative size of the military power. We assume that for each country $n \in \{1,2\}$

$$\Pr[d_{jt} \neq 0] = k_{jt}(k_{1t} + k_{2t})^{-1}$$

At the beginning of each period t, each power decides whether to declare war on the other one, and if not how much to increase the defence force.

5.4 Investment in plant capacity

The next example considers N firms who compete on quantity basis. Units of capital cost are measured in dollars. Each firm chooses how much of is profits to plough back into the firm versus release to shareholders as dividends. Thus

$$k_{n,t+1} = \delta k_{nt} + g(i_{nt})$$

and

$$\pi_{nt} = f(k_{nt})\left[p\left(\sum_{n=1}^{N} f(k_{nt})\right) - c\right] - i_{nt}$$

The state variables in this model is the vector of capital stocks $(k_{1t}, \ldots, k_{Nt})$. The value of the firm at time s is the expected value of the dividend stream from that period onwards:

$$E_s\left[\sum_{t=s}^{\infty} \beta^t \pi_{nt}\right]$$

We consider the life cycle of product of a firm which produces revenue from its plant of size k_t at period t. Given capital of k_t it produces material of $f(k_t)$ which can be sold as or transformed into capital for next period, denoted . The output which is marketed

$$f(k_t) = y_t + k_{t+1}$$

generates revenue of $\pi(y_t)$. We assume that both $f(k_t)$ and $\pi(y_t)$ are increasing concave twice differentiable function; this implies $\pi'(y_t) \geq 0$ and $\pi''(y_t) < 0$ and similarly for $f(k_t)$. Given an initial start up capital of k_0 the firm chooses $\{k_t\}_{t=1}^{T}$, or equivalently sales $\{y_t\}_{t=1}^{T}$to maximize its present value

$$\sum_{t=0}^{T} \frac{\pi(y_t)}{1 + r_t}$$

subject to the internal financing constraint.

To solve this problem we substitute y_t in the firm's objective function to obtain

$$\sum_{t=0}^{T} \frac{\pi(f(k_t) - k_{t+1})}{1 + r_t}$$

differentiate with respect to k_t and substitute y_t and y_{t-1} back into the expressions for output to obtain the first order condition

$$\frac{\pi'(y_{t-1})}{1 + r_{t-1}} = \frac{\pi'(y_t) f'(k_t)}{1 + r_t}$$

out. This condition has an intuitive interpretation. The expression $\pi'(y_{t-1})$ is the increment in profits due to marketing another unit of output in period $t - 1,$ and dividing this term by $(1 + r_{t-1})$ puts it in present value terms. Rather than sell the output to consumers in period $t - 1$ the firm could use the units to produce $f'(k_t)$ units of capital which in turn increase profits by $\pi'(y_t)$ in period $t.$ Dividing these two expressions by $(1 + r_{t-1})$ and $(1 + r_t)$ respectively put the expressions in present value terms. If the firm is following an optimal investment plan the value created from these two activities are equated.

To make progress solving this problem, we now focus on some special cases. Suppose the production function is $f(k_t) = k_t^{\alpha}$ for some $\alpha \in (0, 1),$ the revenue function is $\pi(y_t) = \log y_t$, and the interest rate is a constant $r \in (0, 1)$. Then $(1 + r_t) = (1 + r)^t$ and the first order condition derived above reduces to

$$(1 + r)(k_t^{\alpha} - k_{t+1}) = \alpha(k_{t-1}^{\alpha} - k_t) k_t^{\alpha - 1}$$

We may rewrite this difference equation as

6 Summary

This chapter introduced the dual concepts of stages and transition probabilities as tools for parsimoniously representing games where there are elements of repetition. We began by asking whether solutions to stage and Markov games can be found by simply joining the solutions for each stage played independently. In stage games (where the probability transition is independent of the choices made), strategy profiles formed this way are indeed a subset of the solution set for the whole game. In Markov games (where the probability transitions are affected by choices), the arguments developed in Chapter 3 on investment and Chapter 6 on perfect information games apply here with equal force: we would not expect this approach generate a solution to the game.

Not only does the Markov form provide a convenient way of defining games that have no fixed ending time. It also suggests that recursive algorithms might for generating a solution which we formally estabvlished in the case of perfect information games, by modifying some resutls we obtained in Chapter 3. Sequential games with imperfect information are trickier: we found that checking the Markov solutions to an infinite horizion game for two players who simlutaneously select one of two actions wadh period when there are only two stages is quite an onerous taks. Numerical techniques are often required to solve more complex games. Simplifications are available when the stages in each ergodic set are considerably fewer than the total number of stages.

One feature of the Markov solutions that we have not emphasized in this chpater is that the solution concept in some sense uses the least information possible that is compared with other solutions ot the game. An underlying theme of this chapter is to identify those games in which the Markov solution mimics a coalition of interest that maximizes the portential gains from cooperation. A question we address in the next chapter is whether there are other solutions that use more information about play, and when such solutions can enforce cooperative play, pariticularly in cases where the Markov solutions do not.