Problems on Repeated Games

Part I: Finite Dynamic Programming

Consider investing in a project that may be in one of two states: active or bankrupt. If the project is bankrupt, it remains that way forever, and pays out 0 per period. If it is active it you must choose whether or not to invest in the project. If you invest, you receive a net profit of one, and the project will remain active next period. If you do not invest, you receive a net profit of two and the project has a fifty percent chance of going bankrupt next period. If the project is initially active, for what values of the (fixed) subject discount factor δ will you choose to invest?

Part II: Long Run and Short Run

1. Easy Long versus Short-Run

Consider the following chain store game played between a patient player one (chain store) with discount factor δ and a sequence of short-run myopic player 2's (entrants – with discount factor 0)

	out	**in**
fight	3,0	-2,-2
give in	4,0	2,2

a. What is the Nash equilibrium if the game is played once?

b. What is the Stackelberg equilibrium in which player 1 gets to commit if the game is played once?

c. What is the subgame perfect equilibrium if the game is repeated $T < \infty$ times?

d. If the game is infinitely repeated, find a δ and strategies for both players such that the long-run player gets 3.

2. Hard Long versus Short-Run

Consider the following game played between a patient player one with discount factor δ and a sequence of short-run myopic player 2's with discount factor 0

	a	**b**	**c**	**d**	**e**
A	0,4	3,3	0,0	-1,0	-2,-1
B	0,0	4,3	0,4	-1,0	-2,-1
C	0,0	0,0	0,0	0,2	-2,1

a. What are all the mixed Nash equilibrium if the game is played once?

b. What is the pure Stackelberg equilibrium in which player 1 gets to commit if the game is played once?

c. What is the mixed Stackelberg equilibrium in which player 1 gets to commit if the game is played once?

d. If the game is infinitely repeated, and δ is "close enough to one" find the best equilibrium payoff for the long-run player and the worst.

e. How close is "close enough to one" in part d?

3. Long-run Short-Run with Noise

Consider the following quality game. A short-run Player 2 must decide whether or not to purchase a good. After the decision is made, a long-run Player 1 must decide whether to produce high or low quality. If high quality is chosen, there is a 90% chance the good is acceptable, and 10% that it is defective. If low quality is chosen there is a 10% chance the good is acceptable and 90% chance it is defective. If no sale is made, both players get nothing. If the high quality good is sold, the producer (player 2) gets 9, if the low quality good is sold, he gets 10. If an acceptable good is purchased the consumer (player 1) gets a utility net of purchase cost of 10; if a defective good is purchased, he gets a utility net of purchase cost of -10.

a. Find the normal form of this game.

b. Find all Nash equilibria and the minmax and pure and mixed Stackleberg payoffs for player 1.

c. If the game is repeated between a long-run player 1 and short-run player 2 find the set of perfect public equilibria with public randomization when quality can be observed.

d. If the game is repeated between a long-run player 1 and short-run player 2 find the set of perfect public equilibria with public randomization when quality cannot be observed.

e. How would your answer change if there were many types of long-run player?

4. Greenspan

A long-lived central bank faces a short-run representative consumer. The bank must decide whether or not to inflate; the consumer must decide whether or not to expect inflation. If the consumer guesses correctly, she gets 1; incorrectly she gets 0. Central bank payoffs are

	Guess inflate	**Guess not**
inflate	0	2
not	-10	1

As a result of whether or not the central bank chose to inflate, economic activity is determined: there are two possibilities hyperinflation or price stability. If the bank chose to inflate the probability of hyperinflation is 1; if the bank chose not to inflate, the probability of hyperinflation is 10%. In all that follows, equilibrium means perfect public equilibrium of the infinitely repeated game with public randomization.

a. Find the extensive and normal forms of the stage-game.

b. For the long-run player, find the minmax, the static Nash, mixed precommitment and pure precommitment payoffs.

c. Find the worst equilibrium for the long-run player, and describe in general terms the set of equilibrium payoffs for the long-run player.

First assume that the consumer can observe whether or not the central bank inflates.

d. Find the best equilibrium for the central bank as a function of the discount factor.

Now assume that the consumer cannot observe whether or not the central bank inflates but can observe whether or not there is hyperinflation.

e. Find the best equilibrium for the central bank as a function of the discount factor.

5. Auto Repair

A long-lived auto repair shop with discount factor $\delta > 0$ faces a sequence of short-lived car owners. The car owners must each decide whether to have their cars repaired or not. If they do, the repair shop must decide whether to repair the car or not. If the car is not repaired, the probability it will work is $1 > \pi > 0$. If it is repaired, the probability it will work is $1 \geq \theta > \pi$. The price of the repair is $p > 0$; the cost of repair to the shop is $0 < c < p$. A car that does not work is worth nothing. A car that works is worth v. Assume that $(\theta - \pi)v > p$. Car owners can only observe whether or not the car works, not whether or not the shop repaired it. In all that follows, *equilibrium* means perfect public equilibrium of the infinitely repeated game with public randomization.

a. Find the extensive and normal forms of the stage-game.

b. For the long-run player, find the minmax, the static Nash, mixed precommitment and pure precommitment payoffs.

c. Find the worst equilibrium for the long-run player, and describe in general terms the set of equilibrium payoffs for the long-run player.

d. Find the best equilibrium for the repair shop as a function of the parameters.

Part III: Repeated Games

1. Folk Theorem

Consider the following Prisoner's dilemma game played with public randomization

3,3	0,8
8,0	1,1

a. Sketch the socially feasible, individually rational set.

b. How is the sum of player payoffs maximized?

c. Find a discount factor and subgame perfect strategies such that each player receives half the maximum sum of player payoffs.

2. Not So Nice Folk Theorem

Consider the following coordination game:

2,2	1,0
0,1	0,0

a. What is the unique static Nash equilibrium?

b. Sketch the socially feasible, individually rational set.

c. Find a discount factor and subgame perfect strategies such that each player receives 1.5.

d. Can you find an information system for which this is an equilibrium in a matching protocol?

Part IV: Information Conditions

1. Short Answers

In each of the following games, determine whether which pure strategy profiles are enforcible, and which are pairwise identifiable? Which games admit a strategy profile that satisfies the pairwise full rank condition? What implications do your calculations have for the set of perfect public equilibrium payoffs with equally (and very) patient players?

a. In a two person partnership, each person may provide either one or zero units of effort, and k levels of output are possible $(0, y, 2y, 3y, \ldots (k\text{-}1)y)$. Output depends only on the total amount of effort (0, 1 or 2). Every level of output has positive probability, and the distribution of output for a higher total amount of effort stochastically dominates that for any lower level of effort. Output is equally shared between the partners, and utility is linear in output and effort.

b. In a principal-agent game, the agent secretly chooses either one or zero units of effort, and k levels of output are possible. The principal gets all the output and may pay either zero or w dollars to the agent as a function of the output level. The principal's utility is linear in output and dollars, the agent is risk averse for dollars, and has utility linear in effort. The distribution of output for high effort stochastically dominates that for low effort. Every level of output has positive probability.

c. In a duopoly, two firms produce goods that are perfect substitutes and choose secretly between one and two units of output. Marginal cost is zero, and demand is linear with slope -1, and a random intercept that can take on k different (as above) levels each with positive probability. Assume that the monopoly solution facing the expected demand curve is to produce two units of output. Firms do not observe demand, or their rivals output, they only see the market price.

2. Enforceability

Consider the insurance game with a single consumption good in which a player receives an endowment of one and draws independently every period a marginal utility of either

$\overline{\eta}$ (hungry) or $\underline{\eta}$ (not hungry), where $\overline{\eta} > \underline{\eta} > 0$. The probability of being hungry is $1 > \pi > 0$. A strategy for the player is an announcement of his type as a function of his true type, hence there are four strategies: tell the truth, tell the opposite of the truth, always say "hungry" and never say "hungry." Whenever the player says he is hungry he has a $1-\pi$ chance of getting two units of consumption and a π chance of autarky (one unit of consumption). If he says he is not hungry he has a π chance of getting nothing and a $1-\pi$ chance of autarky.

a. If $\pi = 1/2$ show that the strategy of telling the opposite of the truth is not enforceable.

b. If $\pi > 1/2$ show that the strategy of telling the opposite of the truth is not enforceable.

c. If $\pi < 1/2$ show that the strategy of telling the opposite of the truth is not enforceable.

Part V: Reputation

1. The Chain Store Paradox-Paradox

Consider the Kreps-Wilson version of the chain store paradox: An entrant may stay out and get nothing (0), or he may enter. If he enters, the incumbent may fight or acquiesce. The entrant gets *b* if the incumbent acquiesces, and *b-1* if he fights, where $0<b<1$. There are two types of incumbent, both receiving $a>1$ if there is no entry. If there is a fight, the *strong* incumbent gets 0 and the weak incumbent gets -1; if a strong incumbent acquiesces he gets -1, a weak incumbent 0.

Only the incumbent knows whether he is weak or strong; it is common knowledge that the entrant a priori believes that he has a π0 chance of facing a strong incumbent. Define

$$\gamma = \frac{p_0}{1-p_0}\frac{1-b}{b}$$

a. Sketch the extensive form of this game.

b. Define a sequential equilibrium of this game.

c. Show that if $\gamma \neq 1$, there is a unique sequential equilibrium, and that if $\gamma > 1$ entry never occurs, while if $\gamma < 1$ entry always occurs.

d. What are the sequential equilibria if $\gamma = 1$?

e. Now suppose that the incumbent plays a second round against a different entrant who knows the result of the first round. The incumbent's goal is to maximize the sum of his payoffs in the two rounds. Show that if $\gamma > 1$ there is a sequential equilibrium in which the entrant enters on the first round and both types of incumbents acquiesce. Be careful to specify both the equilibrium strategies and beliefs.

2. Reputation

A sequence of consumers must choose what product to buy from Gigantic Corporation: a mediocre product or a special improved brand. The mediocre product yields a utility to the consumer of 1 and a profit to Gigantic of 1. The special improved brand yields a utility of 2 and a profit of 2. However, Gigantic has the option of producing a cheap imitation brand that is indistinguishable from the special improved brand. This yields a utility of 0 and a profit of 4. If a consumer buys a special improved brand, he finds out whether or not it is the cheap imitation, and reveals this information to later consumers.

a. Show that there is a sequential equilibrium in which Gigantic produces only cheap imitations and consumers always buy the mediocre product.

b. If Gigantic is very patient and there is a positive probability that it is "honest" and does not produce imitations, does this make a difference?

c. Would it make a difference if Gigantic has also the option of producing defective products that are indistinguishable from mediocre products? These yield a utility of -1 and a profit of 0.

d. What if in part c. all pure strategy types have equal probability?

3. *Bad Reputation*

A series of two short-run consumers chooses whether to visit a long-run mechanic with discount factor δ or to buy a new car. The existing care either needs a new engine or a tune-up, each with 50% probability. Only the mechanic knows whether an engine or tune-up is needed. There are two types of mechanics: with probability p the mechanic is honest otherwise the mechanic is dishonest. A dishonest mechanic always says the care needs a new engine.

If the consumer chooses not to go to the mechanic, he gets 0 as does the honest mechanic. If he goes to the mechanic and the mechanic tells the truth, the consumer gets 2 and the honest mechanic 2; if the mechanic lies, the consumer gets –4 and the honest mechanic 0. So the consumer moves first and chooses whether or not to go to the mechanic. The mechanic moves second, observes the true state of the car, then either reports that an engine is needed or a tuneup is needed. Remember, the game is repeated twice each with time with a different consumer. The second consumer does not observe whether the mechanic is honest or dishonest, only whether he recommended a new engine or a tuneup.

a. What are the best and worst sequential equilibria when $p = 1$?

b. Find a value of p such that the consumer would choose to go to the mechanic if the honest mechanic tells the truth, but there is no sequential equilibrium in which the consumer goes to the mechanic.

4. Inference and Martingales

A single decision-maker picks a sequence of actions $a_t \in A$, a finite set. He is drawn from a finite set of types Ω. If $h_t = (a_1, a_2, \ldots a_t)$ is the history of his play through *t* his strategy may be described by a probability distribution over *A* at time *t*, $\sigma_t(h_{t-1}, \omega)$, which depends on the history and his type. You observe the play of this player, and place probability $\mu(\omega) > 0$ on his being type ω.

Consider $\mu(\omega | h_t)$. By Bayes law

$$\mu(\omega \mid h_t) = \frac{\sigma_t(h_{t-1}, \omega)(a_t)\mu(\omega \mid h_{t-1})}{\sum_{\omega'} \sigma_t(h_{t-1}, \omega')(a_t)\mu(\omega' \mid h_{t-1})}$$

Fix a type ω^+, and let $\Omega^+ \equiv \Omega \setminus \omega^+$ be the set of all other types. We may define random variables p_t, q_t by

$$p_t(a) = \sigma_t(h_{t-1}, \omega^+)(a_t), p_t = p_t(a_t)$$

$$q_t(a) = \frac{\sum_{\omega' \in \Omega^+} \sigma_t(h_{t-1}, \omega')(a_t)\mu(\omega' \mid h_{t-1})}{1 - \mu(\omega^+ \mid h_{t-1})}, q_t = q_t(a_t)$$

We also define L_t recursively by

$$L_0 = \frac{1 - \mu(\omega^+)}{\mu(\omega^+)}$$

$$L_t = \frac{q_t}{p_t} L_{t-1}$$

a. What are p_t and q_t .

b. Show by induction that

$$L_t = \frac{1 - \mu(\omega^+ \mid h_t)}{\mu(\omega^+ h_t)}$$

c. Show that

$$E[L_t \mid L_{t-1}, h_{t-1}, L_{t-2}, h_{t-2}, \ldots \omega^+) \leq L_{t-1}$$

This means (by definition) that L_t is a supermartingale; obviously $L_t \geq 0$.

d. It is known that if L_t is a non-negative supermartingale, with probability one, the sequence $(L_0, L_1, L_2, ...)$ converges to a limit. How can you interpret this fact?