Psy 5018H: Math Models Human Behavior
Spring 2008
Prof. Paul Schrater
Homework #4, Due May 1st

Problem Set

Submit homework as an electronic file via email. You may submit any common file format.

1) Computing games with social utilities (40%)

Read CamererChap3.pdf.
We will use the Fehr-Schmidt formula:
For N Players, the total payoff is denoted

$$X = \{x_1, \cdots, x_N\}$$

Player i's Utility for the total payoff is

$$U_i(X) = x_i - \frac{\alpha}{N-1}\sum_{k \neq i}\lfloor x_k - x_i \rfloor - \frac{\beta}{N-1}\sum_{k \neq i}\lfloor x_i - x_k \rfloor$$

where $\lfloor y \rfloor = \begin{cases} y & \text{if } y > 0 \\ 0 & \text{else} \end{cases}$

Assume that $0 \leq \alpha, \beta < 1$

Here is a Prisoner's Dilemma (PD) game

2 \ 1	C	D
C	3 ,3	0,5
D	5,0	1,1

(a) Suppose there are players whose utilities are given by the Fehr-Schmidt formula above and selfish players (with $\alpha=\beta=0$). They play the PD simultaneously. What strategy will selfish players use if $\alpha = 0.1$ and $\beta=0.5$? Does the value of α and β matter? Suppose that there is a fraction q of Fehr-Schmidt players who all have the same α and the same β (both strictly positive and weakly less then one). (Treat this fraction as the probability you will face an opponent of Type FS or Regular and compute the expected Utility over Type). Find conditions on q, α, and β so that there are two pure-strategy equilibria: One

in which all players choose D; and one in which Fehr-Schmidt players choose C and selfish players choose D.
(b) Find a condition on q, α, and β such that for Fehr-Schmidt players, the equilibrium in which they play C (and selfish players play D) is better in expected utility terms than the all-D equilibrium.

2) Understanding Markov Chains (60%)

Understanding Markov chains. Download ShooterData.m Contained within is 100 free throws from several basketball players. Successes are coded as '1', misses as '0'

- A streak shooter is one that is more likely to get a free throw correct after a previous success.
- A "Daryl Dawkins" shooter is one that is more likely to get a free throw correct after a miss.
- A random shooter has no serial biases.

1) ***Build a conditional probability table**:* P(s(j) |s(j-1)), where j is the trial number, and s is the variable that encodes the player's success on that trial

For each player's freethrow data that looks like

```
      \     s(j-1)=1 s(j-1)=0
---------------------------------
 s(j)=1 |  p(1,1)  p(1,0)
        |
 s(j)=0 |  p(0,1)  p(0,0)
```

Use Matlab to estimate these probabilities from the given sequences. Compute tables for each player. Label each of the shooters as Streak, Dawkins, or random.

2) ***Stationary distributions of a Markov Chain.*** If you have a Markov process that can be described by a single time invariant probability matrix,
p(s(t) | s(t-1), s(t-2), ... , s(t-k)) then there may exist a 'stationary distribution' that describes the probability of the process being in each state at time t, p(s(t)), The distribution is defined via a limiting process, but is frequently easy to compute within reasonable precision by simulating a finite series of steps.

The steps form a series of marginalizations--
If the initial state probability is p(s(0)) , then p(s1) = sum(p(s1 |s0) p(s0))
And p(s2) = sum(p(s2|s1)p(s1))
and so on.
For the conditional tables you just built, these marginalizations can be computed by matrix multiplication.

Let pLP be the conditional matrix estimated from Larry Parrot. Then
from t1 -> t2, p(s2|s0) = p(s2|s1)*p(s1|s0) = pLP*pLP

In general to go n time steps, compute pLP*pLP*....*pLP, n time steps. In matlab, exponentiating a matrix does exactly this operation, pLP*pLP*....*pLP.
For n time steps => pLP^n

Exponentiate the players conditional matrix to several large numbers (>10) that are not too big. Notice that the results are almost all the same, and that the matrix converges to a repeated vector. This vector is called the stationary distribution, and the result means that no matter what the initial state probability p(s0), after a few free throws, the players performance will be described by the stationary distribution.

For any 2x2 conditional matrix
p(s(t) | s(t-1)) = [1-alpha, beta
alpha 1-beta]

There is a simple relationship between alpha, beta, and the stationary distribution that comes from solving for the stationary distribution. Compute the stationary distribution by this formula and by the procedure described above. Do they agree?

3) Derive an equation for the stationary distribution of 3x3 matrix.

4) **Relabelling**: Relabel states [s(t-1) s(t)] with unique identifiers. For example:
[0 0] -> 1
[0 1] -> 2
[1 0] -> 3
[1 1] -> 4
so that the sequence [0 0 1 1 0 0 0] goes to [1 2 4 3 1 1]
Relabel the states for HW. Compute the 4x4 conditional probability distribution for this new sequence. Does this eliminate dependency across time? (No dependency is when all the columns in the table are identical.)

Next try the following relabelling, in which every pair is labelled
so that the sequence [0 0 1 1 0 0] goes to [1 4 1]
Are dependencies bigger or smaller with this relabelling?