Spatial Hotelling Problem with Sequential Selection

written by Cyrus Maz, November 3rd, 2020

Problem Description

Three players A, B, C play the following game. First, A picks a real number between 0 and 1 (both inclusive), then B picks a number in the same range (different from A's choice) and finally C picks a number, also in the same range, (different from the two chosen numbers). We then pick a number in the range uniformly randomly. Whoever's number is closest to this random number wins the game. Assume that A, B and C all play optimally and their sole goal is to maximise their chances of winning. Also assume that if one of them has several optimal choices, then that player will randomly pick one of the optimal choices.

- 1. If A chooses 0, then what is the best choice for B?
- 2. What is the best choice for A?
- 3. Can you write a program to figure out the best choice for the first player when the game is played among four players?

Notation

 A^* : A's choice, B^* : B's choice, C^* : C's choice

1 If A chooses 0, then what is the best choice for B?

The best choice for B is 2/3.

Proof by contradiction: If B chooses 2/3, then $C^* \sim Uniform(0, 2/3)$. Then $P(B \text{ winning}|B^* = 2/3, C^*) = 1/3 + (2/3 - C^*)/2$ and therefore, $\mathbb{E}[P(B \text{ winning}|B^* = 2/3, C^*)] = 1/3 + 1/3 - \mathbb{E}[C^*]/2 + 2/3 - 1/6 = 1/2$ Now, for a contradiction, suppose the optimal choice for B is $2/3 + \alpha$ Case 1: $\alpha \in (0, 1/3]$

Then $C^* \sim Uniform(0, 2/3 + \alpha)$ and $P(B \text{ winning}|B^* = 2/3 + \alpha) = 1/3 - \alpha + (2/3 + \alpha - C^*)/2.$ Therefore,

$$\mathbb{E}[P(B \text{ winning}|B^* = 2/3 + \alpha, C^*)] = 1/3 - \alpha + (2/3 + \alpha - \mathbb{E}[C^*])/2$$

= 1/3 + 1/3 - \alpha + \alpha/2 - \mathbb{E}[C^*]/2
= 2/3 - \alpha/2 - (2/3 + \alpha)/4
= 2/3 - 1/6 - 3/4\alpha
= 1/2 - 3/4\alpha < 1/2 = \mathbb{E}[P(B \text{ winning}|B^* = 2/3, C^*)] (1)

where the final inequality holds because $\alpha > 0$.

Case 2: $\alpha \in (-2/3, 0)$ Then $C^* = 2/3 + \alpha + \epsilon$ for some very small $\epsilon > 0$ and

$$\mathbb{E}[P(B \text{ winning}|B^* = 2/3 + \alpha, C^*)] = P(B \text{ winning}|B^* = 2/3 + \alpha)$$

= $B^*/2 + (C^* - B^*)/2$
= $C^*/2$
= $(2/3 + \alpha + \epsilon)/2$
= $1/3 + \alpha/2 + \epsilon/2$
< $1/2 = \mathbb{E}[P(B \text{ winning}|B^* = 2/3, C^*)]$ (2)

Where the finally inequality holds because $\alpha \in (-2/3, 0)$.

Since other values of α would make B^* fall outside of the unit interval [0, 1] we conclude that the optimal choice for B is 2/3.

Note: as a sanity check, I used the algorithm outlined in my solution to question 3 to solve for the optimal choice for player B, given that A plays 0 in a three player setting, and it yielded the same answer as the one shown here.

2 What is the best choice for A?

The best choice for A is drawing randomly from a (discrete) Unifrom $\{1/4, 3/4\}$ distribution. *Proof:*

First, by contradiction we show that choosing 1/2 is sub-optimal for A.

Suppose A chooses 1/2. Then B must choose from the following set: $[0, 1/2) \cup (1/2, 1]$.

Then the optimal choices for B and C are $1/2 \pm \epsilon$ and $1/2 \mp \epsilon$ for a very small $\epsilon > 0$, and:

 $\mathbb{E}[P(\mathbf{A} \text{ winning}|A^* = 1/2)] = P(\mathbf{A} \text{ winning}|A^* = 1/2) = \epsilon$

Then, considering the scenario in question 1, we note that if $A^* = 0$, then $B^* = 2/3$, and subsequently $C^* \sim Uniform(0,2/3)$. Therefore, $P(A \text{ winning}|A^* = 0) = C^*/2$ and thus

$$\mathbb{E}[P(\mathbf{A} \text{ winning}|A^* = 0)] = \mathbb{E}[C^*]/2$$

$$= 1/6$$

$$> \epsilon = \mathbb{E}[P(\mathbf{A} \text{ winning}|A^* = 1/2)]$$
(3)

Hence, 1/2 is a sub-optimal choice for A.

Without loss of generality, assume that $A^* < 1/2$. Then,

- i. if $B^* \leq 1/2$ and $A^* > B^*$ then $C^* = A^* + \epsilon$, for a very small $\epsilon > 0$ and $P(B \text{ winning}|A^*, B^*) = B^* + (A^* B^*)/2 = (B^* + A^*)/2$ $P(A \text{ winning}|A^*, B^*) = (B^* - A^*)/2 + \epsilon/2$
- ii. if $B^* \leq 1/2$ and $A^* < B^*$ then $C^* = B^* + \epsilon$, for a very small $\epsilon > 0$ and $P(B \text{ winning}|A^*, B^*) = (B^* A^*)/2 + \epsilon/2$ $P(A \text{ winning}|A^*, B^*) = A^* + (B^* - A^*)/2 = (B^* + A^*)/2$
- iii. if $B^* > 1/2$ and $max(A^*, (B^* A^*)/2, 1 B^*) = 1 B^*$ then $C^* = B^* + \epsilon$ for a very small $\epsilon > 0$ and $P(B \text{ winning}|A^*, B^*) = (B^* A^*)/2 + \epsilon/2$ $P(A \text{ winning}|A^*, B^*) = A^* + (B^* - A^*)/2 = (B^* + A^*)/2$
- iv. if $B^* > 1/2$ and $max(A^*, (B^* A^*)/2, 1 B^*) = A^*$ then $C^* = A^* \epsilon$ for a very small $\epsilon > 0$ and $P(B \text{ winning}|A^*, B^*) = (1 B^*) + (B^* A^*)/2 = 1 (B^* + A^*)/2$ $P(A \text{ winning}|A^*, B^*) = \epsilon/2 + (B^* - A^*)/2$
- v. if $B^* > 1/2$ and $max(A^*, (B^* A^*)/2, 1 B^*) = (B^* A^*)/2$ then $C^* \sim Uniform(A^*, B^*)$, $P(B \text{ winning}|A^*, B^*) = (1 - B^*) + (B^* - C^*)/2 = 1 - 1/2B^* - C^*/2$; and thus, $\mathbb{E}[P(B \text{ winning}|A^*, B^*)] = 1 - 1/2B^* + \mathbb{E}[C^*]/2 = 1 - 1/2B^* - (A^* + B^*)/4 = 1 - 3/4B^* - A^*/4$ $P(A \text{ winning}|A^*, B^*) = A^* + (C^* - A^*)/2 = A^*/2 + C^*/2$; and thus, $\mathbb{E}[P(A \text{ winning}|A^*, B^*)] = A^*/2 + (B^* + A^*)/4$

Let us revise the enumerated list above into explicit equations for expected conditional probabilities of A and B winning.

$$\mathbb{E}[P(\text{B winning}|A^*, B^*)] = \begin{cases} (B^* + A^*)/2 + \epsilon/2 \text{ if } B^* \leq 1/2 \text{ and } A^* > B^* \\ (B^* - A^*)/2 + \epsilon/2 \text{ if } B^* \leq 1/2 \text{ and } A^* < B^* \\ (B^* - A^*)/2 + \epsilon/2 \text{ if } B^* > 1/2 \text{ and } \max(A^*, (B^* - A^*)/2, 1 - B^*) = 1 - B^* \\ (B^* + A^*)/2 + \epsilon/2 \text{ if } B^* > 1/2 \text{ and } \max(A^*, (B^* - A^*)/2, 1 - B^*) = A^* \\ (1 - 3/4B^* - A^*/4 + \epsilon/2 \text{ if } B^* > 1/2 \text{ and } \max(A^*, (B^* - A^*)/2, 1 - B^*) = (B^* - A^*)/2 \end{cases}$$

$$\mathbb{E}[P(A \text{ winning}|A^*, B^*)] = \begin{cases} (B^* - A^*)/2 + \epsilon/2 \text{ if } B^* \leq 1/2 \text{ and } A^* > B^* \\ (B^* + A^*)/2 + \epsilon/2 \text{ if } B^* \leq 1/2 \text{ and } A^* < B^* \\ (B^* + A^*)/2 + \epsilon/2 \text{ if } B^* > 1/2 \text{ and } max(A^*, (B^* - A^*)/2, 1 - B^*) = 1 - B^* \\ (B^* - A^*)/2 + \epsilon/2 \text{ if } B^* > 1/2 \text{ and } max(A^*, (B^* - A^*)/2, 1 - B^*) = A^* \\ A^*/2 + (B^* + A^*)/4 + \epsilon/2 \text{ if } B^* > 1/2 \text{ and } max(A^*, (B^* - A^*)/2, 1 - B^*) = (B^* - A^*)/2 \end{cases}$$

We now have everything we need to find the optimal choice for A numerically. Let

$$\tilde{B} = \left\{ (a,b) : b = \underset{b}{\operatorname{argmax}} \mathbb{E}[P(\mathbf{B} \text{ winning}|A^* = a, B^* = b)], a \in [0, 1/2) \right\} \text{ and let}$$
$$\tilde{A} = \left\{ a : a = \underset{a}{\operatorname{argmax}} \mathbb{E}[P(\mathbf{A} \text{ winning}|A^* = a, B^* = b)], (a,b) \in \tilde{B} \right\}$$

Then \tilde{A} is the set of optimal plays by A in the interval [0, 1/2). We will approximate \tilde{B} by grid search over $(A^* \times B^*) \in [0, 1/2] \times [0, 1]$. Next, we will calculate (approximately) \tilde{A} by searching through the approximation of \tilde{B} .



Figure 1: (Question 2) Left: the surface of the expectation of the conditional probability of B winning represented as a heatmap. Right: Approximation of \tilde{B} computed by grid search.

Using the Python code included in Appendix A, with a grid step size of 0.001, we find that $\tilde{A} = \{0.25\}$ Then, by symmetry, 0.75 is also an optimal choice for A. Thus, we have that the optimal choice for A is drawing randomly from a (discrete) Unifrom $\{1/4, 3/4\}$ distribution.

Note: as a sanity check, I used the algorithm outlined in my solution to question 3 to solve for the optimal choice for player A in a three player setting, and it yielded the same answer as the one calculated here. Output from question 3's algorithm used to solve this question is included at the end of Appendix B.

3 Can you write a program to figure out the best choice for the first player when the game is played among four players?

3.1 Caveats

We must first note that since no player may choose a number already chosen by any of the previous players, the space of potential choices for all but the first player are non-compact sets. In a strictly mathematical sense, this makes "optimal" choices to *not exist* in certain cases. To demonstrate this subtle point, consider the following scenario in a two player setting: player A chooses 1/3; then the choice for player B that *maximizes* their probability of winning is $1/3 + \epsilon$, for a very small $\epsilon > 0$ which yields $P(B \text{ winning}|A^* = 1/3) = 1 - (1/3 + \epsilon) = 2/3 - \epsilon$. Note that

$$\lim_{\epsilon \to 0^+} P(\mathbf{B} \text{ winning} | A^* = 1/3) = 2/3$$

in a monotonically increasing manner, and thus $P(B \text{ winning}|A^* = 1/3)$ cannot be maximized over (1/3, 1].

To circumvent this mathematical nuance, we will impose a minimum acceptable value for ϵ . Let ϵ_0 denote this minimum acceptable value for ϵ .

Furthermore, to solve the puzzle algorithmically, we will approximate the exact solution, by discretizing [0, 1] into M equally-lengthed intervals, thereby restricting the set of possible choices to $\{i/(M+1) : i = 0, 1, ..., M-1, M, M+1\}$ for some large integer M. In this discrete setting, ϵ_0 must be $\leq 1/M$.

3.2 General Procedure

Let N be the number of players. Let M be the number of equally spaced intervals [0, 1] is discretized into. Consequently, the set of possible choices is restricted $\{i/(M+1) : i = 0, 1, ..., M - 1, M, M + 1\}$. Let $\epsilon_0 = 1/M$. Start with every possible (optimal and non-optimal) sequence of choices by the first N - 1 players. Then for each sequence, we find the optimal choice for the Nth player. Algorithm 1, *optimal_path_calculator(sequence)*, outlines the procedure for this task. A sequence of one choices for the first N - 1 players followed by the corresponding optimal choice for the N-th player (conditional on the first N - 1 choices), shall be referred to as a *path* henceforth. A set of paths such that an initial sequence (of some length) is shared amongst the paths shall be referred to as a *group* henceforth. The set of all groups shall be referred to as the *path-space* henceforth. We will regroup and prune the path-space, in essence optimizing the choice of each player, according to what the subsequent player(s) will play conditional on the choice of the player in question as well the choice(s) of any preceding player(s). We repeat this procedure iteratively starting with the Nth player and working backwards until we find the optimum choice(s) for the first player. The pseudocode in Algorithm 2, *backward_solver(M,N)*, describes this procedure in more detail. Working Python code as well as example output for various values of M are included in Appendix B.

Note that in the code in the appendix, we restrict the path-space to those whose choice of player A is in $[0, 1/2] \cap \{i/(M+1) : i = 0, 1, ..., M-1, M, M+1\}$. This is to save on computational cost and is justified by the fact that if some value a_* is optimal for player A on [0, 1/2], then by symmetry, $1 - a_*$ is also optimal.

3.3 Results & Discussion

Running the Python script with N = 4 and M = 100, 127, 200 gives optimal choices for A approximately equal to 0.16 and 0.83. It is plausible that the exact answer is 1/6 and 5/6 as the estimates are close to these fractions.

Figure 2 summarizes these results for N = 4 and M = 100, 127, 200. The plots on the left show the expected probability of A winning for each $A^* \in [0, 1/2] \cap \{i/(M+1) : i = 0, 1, ..., M - 1, M, M + 1\}$. The plots on the right show the possible sequence of optimal choices by each player. (Note that in our script, whenever D has an interval of optimal choices available to them, we choose the middle point for accuracy and savings on computational cost. More on this in section 3.4.)

Discretizing the interval [0, 1] results in this solution being an approximation. I believe that for sufficiently large M, the solution produced by this algorithm can be arbitrarily close to the exact solution. However, the computational cost for large choices of M and N are significant. For this reason I was unable to arrive at an answer with M > 200, in the four player setting.

3.3.1 Plots



Figure 2: The general shape of the mapping $[0, 1/2] \mapsto \mathbb{E}[P(A \text{ winning}|A^*)]$ is consistent for all values of M tested. Interestingly the mapping appears to be less smooth for M=200 than M=100 or 127.

3.4 Key Algorithms

Algorithm 1: optimal_path_calculator(sequence)

Let ϵ_0 denote the minimum acceptable ϵ Let $\{x_1, x_2, ..., x_{N-1}\}$ denote the sequence of *sorted* choices of the first N-1 players

Let $\{x_1, x_2, ..., x_{N-1}\}$ denote the sequence of *sorted* choices of the first N-1 players Let $d = \{x_1, max(x_j - x_{j-1})/2, 1 - x_{N-1} : j = 1, ..., N\}$ Let $d^* = max(d)$ Let p = [] be an empty list to be filled with optimal choices by the Nth player for each $j \in \{2, ..., N-1\}$ do if $(x_j - x_{j-1})/2 = d^*$ then append $(x_j + x_{j-1})/2$ to p end if end for if $x_1 = d^*$ then append $x_1 - \epsilon_0$ to p end if if $x_{N-1} = d^*$ then append $x_{N-1} + \epsilon_0$ to p end ifReturn p

Algorithm 2: *backward_solver(M,N)*

Find all permutations of possible (optimal and non-optimal) plays by the first N-1 players For each permutation, find the optimum play(s) by the N-th player The sequence of each permutation, and the corresponding optimal choices by the N-th player constitutes a group Set k = Nwhile k > 0 do # (find the optimal paths for the (k-1)th player) # Average for each group in the path-space do compute the average payoffs of all paths in the group end for # Regroup if k-2 > 0 then Regroup: One group for each unique initial permutation sequence of length k-2 i.e. the first k-2permutation sequence is shared among members of each group else Regroup: Take the union of all remaining groups so that there is only one group in the path-space end if # Optimize for each group in the path-space do Discard every path that is sub-optimal (in terms of the group payoff average) for the (k-1)th player end for end while Return every unique play by the first player in the remaining paths

It should be noted that $optimal_path_calculator(sequence)$ does not explicitly account for the different choices the final player can make when an interval of optimal choices is available to them. This is justified by the fact that when estimating the expected probabilities of winning for the previous players, using the midpoint of an interval that is optimal for the final player as their choice, rather than averaging the probabilities of winning corresponding to multiple discrete points on the interval, as we do for the preceding players, has two advantages: 1. more accurate estimation for the expected conditional probabilities of player N-1 winning (this boost of accuracy is preserved and transferred backward at each iteration of $backward_solver(M,N)$), and 2. significant savings on computational cost. Furthermore, when multiple optimal intervals are available to the final player the algorithm records the midpoint of each interval as an optimal choice; and when singular optimal choices are available, $optimal_path_calculator(sequence)$ does in fact record each one as an optimal choice.

Appendix A Code for Question 2 (hotelling_q2.py 🖸)

```
import itertools
1
2
   def P_B_wins(A,B):
3
        if (B<=1/2) and (A>B):
4
            return (B+A)/2
5
        if (B \le 1/2) and (A \le B):
6
            return (B-A)/2
7
        if (B>1/2) and (max(A, 1-B, (B-A)/2)==1-B):
8
            return (B-A)/2
9
        if (B>1/2) and (max(A,1-B,(B-A)/2)==A):
10
            return 1-(B+A)/2
11
        if (B>1/2) and (max(A, 1-B, (B-A)/2) == (B-A)/2):
12
            return 1-3/4*B-A/4
13
        if (B==A): return 0
14
15
   def P_A_wins(A,B):
16
        if (B<=1/2) and (A>B):
17
            return (B-A)/2
18
        if (B<=1/2) and (A<B):
19
            return (B+A)/2
20
        if (B>1/2) and (max(A, 1-B, (B-A)/2)==1-B):
21
            return (B-A)/2
22
        if (B>1/2) and (max(A, 1-B, (B-A)/2)==A):
23
            return (B-A)/2
24
        if (B>1/2) and (max(A, 1-B, (B-A)/2) == (B-A)/2):
^{25}
            return A/2+(B+A)/4
^{26}
        if (B==A): return 0
27
^{28}
   mesh_width = 1000
^{29}
   A_vals = [x/mesh_width for x in range(int(mesh_width/2)+1)]
30
   B_vals = [x/mesh_width for x in range(int(mesh_width)+1)]
31
   grid = list(itertools.product(A_vals, B_vals))
32
33
   Tilde_B = []
34
   for A in A_vals:
35
        val=[P_B_wins(A,B) for B in B_vals]
36
        max_val = max(val)
37
        max_index = [i for i, j in enumerate(val) if j == max_val]
38
        Tilde_B.append(B_vals[max_index[0]])
39
40
   Tilde_A = [dict(A=a,P_A_wins=P_A_wins(a,b)) for a,b in zip(A_vals, Tilde_B)]
41
   seq = [p['P_A_wins'] for p in Tilde_A]
42
   max_seq=max(seq)
^{43}
   Tilde_A = list(filter(lambda x: x['P_A_wins']==max_seq, Tilde_A))
44
   print(Tilde_A)
45
46
    # [{'A': 0.25, 'P_A_wins': 0.37525}]
47
```

Appendix B Code for Question 3 (hotelling_q3.py 🗹)

```
from copy import deepcopy
1
   from itertools import permutations, groupby
2
   from functools import reduce
3
   from operator import add
4
   from collections import Counter
5
   from math import isclose
6
   import json
8
    # Given a path, calculate the probability of each player winning
9
    # INPUT:
10
         path: a path of the form {'a':0.2, 'b': 0.5, 'c':0.8,... }
    #
11
    # OUTPUT: the payoff for each player, where payoff is defined
12
              as the probability of the player winning the game
    #
13
   def payoff_calculator(path):
14
        sorted_keys=sorted(path, key=path.get)
15
        output=dict()
16
        for k in range(len(sorted_keys)):
17
18
            if k==0:
19
                lowbo=0
20
                upbo=(path[sorted_keys[k]]+path[sorted_keys[k+1]])/2
21
            elif k==len(sorted_keys)-1:
22
                lowbo=(path[sorted_keys[k-1]]+path[sorted_keys[k]])/2
23
                upbo=1
24
            else:
25
                lowbo=(path[sorted_keys[k-1]]+path[sorted_keys[k]])/2
^{26}
                upbo=(path[sorted_keys[k]]+path[sorted_keys[k+1]])/2
27
            output[sorted_keys[k]]=upbo-lowbo
^{28}
        return output
^{29}
30
    # generate every possible sequence of choices by the first N-1 players
31
    # INPUTS:
32
        M: determines the choices for players are \{i/(M+1) : i=0,1,\ldots,M-1,M,M+1\}
33
       N: the number of players
    #
34
    # OUTPUT: all possible permutations of choices to play by the
35
              first N-1 players
    #
36
   def play_generator(M, N):
37
        discretized = [x/(M) \text{ for } x \text{ in } range(0, M+1)]
38
        perms = list(permutations(discretized, N-1))
39
        output = []
40
        for perm in perms:
41
            processed_perm = {alphabet[i]:perm[i] for i in range(N-1)}
42
            output.append(processed_perm)
^{43}
        return output
44
^{45}
    ### ALGORITHM 1:
46
    # compute the optimal choice for the N-th player, given a sequence of choices by the first
47
    # N-1 players
48
    # INPUTS:
49
        path: a path of N-1 plays by the first N-1 players, {'a':0.2, 'b': 0.5, ... }
    #
50
    # OUTPUT: a dictionary consisting of two key/value pairs:
51
   #
        'path_group': a list of dicitonaries of paths that are equal to the input
52
53
    #
                       path in the first N-1 plays, and have optimal plays by the Nth player
   #
        'payoff_group': a list of dictionaries of payoffs corresponding to the respective path
54
    #.
                         in the path_group list
55
```

```
def optimal_path_calculator(path):
56
         letter = alphabet[len(path)]
57
         vals = list(path.values())
58
        vals.sort()
59
         intervals = []
60
        max_len = 0
61
         for i in range(len(vals)+1):
62
             if i==0: lobo=0
63
             else: lobo=vals[i-1]
64
65
             if i==(len(vals)): upbo=1
66
             else: upbo=vals[i]
67
68
             if not ((i==0) or (i==(len(vals)))):
69
                 length=(upbo-lobo)/2
70
             elif ((i==0) or (i==(len(vals)))):
71
                 length=upbo-lobo
72
73
             if i==0:
74
                 lobo='zero'
75
             elif i==len(vals):
76
                 upbo='one'
77
78
             intervals.append([lobo, upbo, length])
79
             if length>max_len: max_len=length
80
81
         intervals=list(filter(lambda x: isclose(x[2],max_len,abs_tol=r), intervals))
82
 83
        path_list = []
84
         for interval in intervals:
85
             if interval[0] == 'zero':
86
                 path_list.append(min(vals)-e)
 87
88
             elif interval[1]=='one':
89
                 path_list.append(max(vals)+e)
90
             else:
91
                 path_list.append((interval[0] + interval[1])/2)
^{92}
93
        path_group=[]
94
        payoff_group=[]
95
         for n in path_list:
96
             newdict = deepcopy(path)
97
             newdict[letter]=n
98
             path_group.append(newdict)
99
             payoff_group.append(payoff_calculator(newdict))
100
101
        return {'path_group':path_group, 'payoff_group':payoff_group}
102
103
    # compute the average payoff for all paths in the group
104
    def payoff_average_calculator(group):
105
        payoff_group = group['payoff_group']
106
        payoff_average = dict(reduce(add, map(Counter, payoff_group)))
107
        payoff_average = {k:v/len(payoff_group) for k,v in payoff_average.items()}
108
        group['payoff_average'] = [payoff_average]*len(payoff_group)
109
        return group
110
111
    # regroup path_space such that the first k-2 in each group are identical
112
```

```
def regroup(path_space, k):
113
        path_group_agg = []
114
        payoff_group_agg = []
115
        payoff_average_agg = []
116
117
        for e in path_space:
118
             for path, payoff_group, payoff_average in zip(
119
                 e['path_group'],
120
                 e['payoff_group'],
121
                 e['payoff_average']
122
                 ):
123
                 path_group_agg.append(path)
124
                 payoff_group_agg.append(payoff_group)
125
                 payoff_average_agg.append(payoff_average)
126
127
         trip = zip(path_group_agg, payoff_group_agg, payoff_average_agg)
128
         l=alphabet[0:k-2]
129
         print('regrouping such that choices for ({}) in each group are equal'.format(1))
130
        regrouped_raw=[
131
             [*j] for i, j in
132
                 groupby(
133
                      sorted(trip, key=lambda x: [x[0][i] for i in l]),
134
                     key=lambda x: [x[0][i] for i in 1])
135
                 ]
136
137
        regrouped = []
138
         for group in regrouped_raw:
139
             path_group = []
140
             payoff_group = []
141
             payoff_average = []
142
             for path in group:
143
                 path_group.append(path[0])
144
                 payoff_group.append(path[1])
145
                 payoff_average.append(path[2])
146
             d=dict(
                        path_group=path_group,
147
                        payoff_group=payoff_group,
148
                        payoff_average=payoff_average)
149
             regrouped.append(d)
150
        return regrouped
151
152
    # for each group in path_space, discard every path that is suboptimal for player k-1
153
    def optimize_groups(path_space, k):
154
         letter = alphabet[k-2]
155
         print('optimizing {}'.format(letter.upper()))
156
         optimized = []
157
158
        for i in range(len(path_space)):
159
             group = deepcopy(path_space[i])
160
             seq=[g[letter] for g in group['payoff_average']]
161
             max_pay = max(seq)
162
             filtered_group = list(filter(
163
                 lambda x: isclose(x[2][letter],max_pay,abs_tol=r),
164
                 zip(group['path_group'], group['payoff_group'],group['payoff_average'])
165
                 ))
166
             filtered_group = list(zip(*filtered_group))
167
             optimized.append(
168
                 {'path_group': list(filtered_group[0]), 'payoff_group': list(filtered_group[1])})
169
```

```
return optimized
171
    # recursively carry out the while loop in Algorithm 1
172
    def recursive_solver(path_space, k):
173
        print("k={}".format(k))
174
        path_space_withaverage = [payoff_average_calculator(o) for o in path_space]
175
        path_space_regrouped = regroup(path_space_withaverage, k)
176
        path_space_optimized_by_group = optimize_groups(path_space_regrouped, k)
177
        if k==2:
178
            return [payoff_average_calculator(o) for o in path_space_optimized_by_group]
179
        else:
180
            k-=1
181
            return recursive_solver(path_space_optimized_by_group,k)
182
183
    ### ALGORITHM 2:
184
    # wrapper for recursive_solver
185
    def backwards_solver(M,N, question_1=False, write=False):
186
        print('M={}'.format(M))
187
        print('N={}'.format(N))
188
        path_space = play_generator(M,N)
189
190
        if question_1 is True:
191
            path_space = list(filter(lambda x: x['a']==0, path_space))
192
        else:
193
            path_space = list(filter(lambda x: x['a']<=1/2, path_space))</pre>
194
195
        path_space = [optimal_path_calculator(play) for play in path_space]
196
        optimal_paths = recursive_solver(path_space,N)
197
198
        path_group=optimal_paths[0]['path_group']
199
        average_payoff = optimal_paths[0]['payoff_average'][0]
200
201
        if question_1 is False:
202
            A = list(set([x['a'] for x in path_group]))
203
            print('optimal choice(s) for A: {}'.format(A))
204
            print('average payoff: {}'.format(average_payoff))
205
            print('optimal paths: {}'.format(path_group))
206
207
        elif question_1 is True:
208
            B = list(set([x['b'] for x in path_group]))
209
            print('optimal choice(s) for B: {}'.format(B))
210
            print('average payoff: {}'.format(average_payoff))
211
            print('optimal paths: {}'.format(path_group))
212
213
        if write is True:
214
            with open('optimal_paths_M{}_N{}_r{}.txt'.format(M,N,r), 'w') as filehandle:
215
                json.dump(optimal_paths, filehandle)
216
217
    218
    # e: the minimum allowable epsilon
219
    # r: absolute tolerance when comparing floating points
220
    e = 1/100000
221
    r = e/10
222
    alphabet = ['a', 'b', 'c', 'd']
223
    224
    # QUESTION 3: Optimal choie for player A in four player setting.
225
    backwards_solver(M=100, N=4,question_1=False, write=False)
226
```

170

```
# optimal choice(s) for A: [0.16]
227
    # average payoff:
228
    #
         { 'a': 0.2875, 'd': 0.169999999999999998, 'c': 0.25499999999999, 'b': 0.287500000000001}
229
    # optimal paths: [{'a': 0.16, 'b': 0.84, 'c': 0.5, 'd': 0.33},
230
                      {'a': 0.16, 'b': 0.84, 'c': 0.5, 'd': 0.6699999999999999]]
231
232
    # backwards_solver(M=127, N=4, question_1=False, write=False)
233
    # optimal choice(s) for A: [0.16535433070866143]
234
    # average payoff:
235
         # {'a': 0.29035433070866146,
236
         # 'c': 0.24999999999999992,
237
         # 'd': 0.1692913385826772,
238
         # 'b': 0.29035433070866146}
239
    # optimal paths: [
240
         {'a': 0.16535433070866143,
    #
241
           'b': 0.8346456692913385,
    #
242
           'c': 0.49606299212598426,
    #
^{243}
           'd': 0.6653543307086613},
    #
244
    #
         { 'a': 0.16535433070866143,
^{245}
           'b': 0.8346456692913385,
    #
246
           'c': 0.5039370078740157,
    #
247
    #
            'd': 0.3346456692913386}]
248
249
    # backwards_solver(M=200,N=4,question_1=False, write=False)
250
    # optimal choice(s) for A: [0.165]
251
    # average payoff: {'a': 0.290625, 'd': 0.167499999999999998, 'c': 0.25125, 'b': 0.290625}
252
    # optimal paths: [{'a': 0.165, 'b': 0.835, 'c': 0.5, 'd': 0.3325},
253
                      {'a': 0.165, 'b': 0.835, 'c': 0.5, 'd': 0.6675}]
    #
254
255
    256
    # QUESTION 1: Optimal choice for player B when A plays 0.
257
    # backwards_solver(M=100,N=3,question_1=True, write=False)
258
    # optimal choice(s) for B: [0.67]
259
    # average payoff: {'a': 0.1675, 'c': 0.335000000000001, 'b': 0.49749999999999994}}
260
    # optimal paths: [{'a': 0.0, 'b': 0.67, 'c': 0.335}]
261
262
    263
    # QUESTION 2: Optimal choie for player A in three player setting.
264
    # backwards_solver(M=100,N=3,question_1=False, write=False)
265
    # optimal choice(s) for A: [0.25]
266
    # average payoff: {'c': 0.2499966666666666664, 'a': 0.37500166666666667, 'b': 0.3750016666666667}
267
    # optimal paths: [{'a': 0.25, 'b': 0.75, 'c': 0.24999},
268
                     {'a': 0.25, 'b': 0.75, 'c': 0.5},
    #
269
                      {'a': 0.25, 'b': 0.75, 'c': 0.75001}]
270
    #
```