ON URN EARNINGS

Baruch College Mathematics Department
MS in Applied Mathematics for Finance
http://math.baruch.cuny.edu/masters_finance.html

The Game. An urn contains six marbles: you know that three are green and three are red. You randomly remove them from the urn one at a time, revealing their colors. When you pull out a green marble I pay you one dollar, but if it's red you must pay me one dollar. You can stop removing marbles at any time and simply walk away with your winnings. How much is this game worth to you?

Solution. To analyze the game, it is helpful to have a strategy that tells you when to stop pulling marbles out of the urn. Let's look at three such strategies:

- The "Cautious" strategy: stop when you're first ahead by a dollar, assuming that happens at some point in the game.
- The "Greedy" strategy: stop only after you've removed all the green marbles (getting every last dollar from me).
- The "Best" strategy: stop (i) if the first two are green OR (ii) if two of the first three are green OR (iii) you've removed all the green marbles.

The table below shows, for these three strategies and all 20 possible orderings of the marble colors, the marbles you would draw and the payoff of the strategy. For example, with the color order *GGGRRR*: under the Cautious strategy you stop after removing the first marble because you are then 1 dollar ahead; under the Greedy strategy you stop after the third marble (when you've gotten all the greens) and your payoff is 3 dollars; under the Best strategy you stop after two marbles (for reason (i)) and your payoff is 2 dollars. Other lines may be similarly verified; for the Best strategy we also report the reason (i), (ii), or (iii) for stopping.

Color Order	Cautious	Greedy	Best
GGGRRR	1 - *G*	3 - *GGG*	2 - *GG* reason (i)
GGRGRR	1 - *G*	2 - *GGRG*	2 - *GG* (i)
GGRRGR	1 - *G*	1 - *GGRRG*	2 - *GG* (i)
GGRRRG	1 - *G*	0 - *GGRRRG*	2 - *GG* (i)
GRGGRR	1 - *G*	2 - *GRGG*	1 - *GRG* (ii)
GRGRGR	1 - *G*	1 - *GRGRG*	1 - *GRG* (ii)
GRGRRG	1 - *G*	0 - *GRGRRG*	1 - *GRG* (ii)
GRRGGR	1 - *G*	1 - *GRRGG*	1 - *GRRGG* (iii)
GRRGRG	1 - *G*	0 - *GRRGRG*	0 - *GRRGRG* (iii)

Payoff table continued:

Color Order	Cautious	Greedy	Best
GRRRGG	1 - *G*	0 - *GRRRGG*	0 - *GRRRGG* (iii)
RGGGRR	1 - *RGG*	2 - *RGGG*	1 - *RGG* (ii)
RGGRGR	1 - *RGG*	1 - *RGGRG*	1 - *RGG* (ii)
RGGRRG	1 - *RGG*	0 - *RGGRRG*	1 - *RGG* (ii)
RGRGGR	1 - *RGRGG*	1 - *RGRGG*	1 - *RGRGG* (iii)
RGRGRG	0 - *RGRGRG*	0 - *RGRGRG*	0 - *RGRGRG* (iii)
RGRRGG	0 - *RGRRGG*	0 - *RGRRGG*	0 - *RGRRGG* (iii)
RRGGGR	1 - *RRGGG*	1 - *RRGGG*	1 - *RRGGG* (iii)
RRGGRG	0 - *RRGGRG*	0 - *RRGGRG*	0 - *RRGGRG* (iii)
RRGRGG	0 - *RRGRGG*	0 - *RRGRGG*	0 - *RRGRGG* (iii)
RRRGGG	0 - *RRRGGG*	0 - *RRRGGG*	0 - *RRRGGG* (iii)

Notice that the Best strategy does not produce the highest payoff for every possible ordering of the colors. So in what sense is it the best? It's the best in the sense that it produces the largest *expected* payoff. This is an idea from probability theory, but when the outcomes of a random process are equally likely (as they are here), the expected payoff is just an average of payoffs across the possible outcomes. So to compute the expected value of a strategy, simply add up the payoffs for that strategy in the 20 different color orderings and divide by 20. This produces $15/20 = .75$ for the Cautious and Greedy strategies and $17/20 = .85$ for the Best strategy.

The Best strategy produces the highest expected payoff, at least among the three strategies we've considered. But how can we be sure there isn't a better strategy for deciding when to stop? To answer this requires a more sophisticated approach. Let's pretend that instead of having one urn before us we have sixteen. Each one is marked to indicate how many red marbles r and how many green marbles g are in the urn. Pretend we have all combinations of r and g with $r = 0$, 1, 2, or 3 and $g = 0$, 1, 2, or 3 (hence 16 urns). Since the urns are marked, you know how many marbles are in each urn and how many are red and how many are green. Let $v(r, g)$ denote the value of the game to you if we play it with the urn marked r reds and g greens. We wish to calculate $v(3, 3)$ and show that it is .85.

Clearly, if there are no green marbles in the urn then the value of the game is 0. This is because, since you are assured of only pulling out red marbles, you would simply walk away and not even start playing the game. Hence $v(0,0) = v(1,0) = v(2,0) = v(3,0) = 0$. On the other hand, if there are only green marbles in the urn, you would pull them all out and earn as many dollars as there are green marbles. Hence $v(0,1) = 1$, $v(0,2) = 2$, and $v(0,3) = 3$. We summarize this in the table below:

		Number of Greens		
	0	1	2	3
Number of Reds 0	0	1	2	3
1	0			
2	0			
3	0			???

Now I claim we can calculate $v(1,1)$, the value of the game if we play it with the urn marked one red and one green. Consider what can happen when you pull out a marble. With probability $1/2$ you'll remove a red marble, in which case you pay me a dollar and you are left with an urn with 0 reds and 1 green. On the other hand, with probability $1/2$ you'll remove a green marble, in which case I pay you a dollar but you're left with an urn with 1 red and 0 greens. Hence:

$$\begin{aligned} v(1,1) &= (1/2)\big(-1+v(0,1)\big)+(1/2)\big(1+v(1,0)\big) \\ &= (1/2)(-1+1)+(1/2)(1+0) \;=\; 0.5. \end{aligned}$$

Adding this to our table gives

		Number of Greens		
	0	1	2	3
Number of Reds 0	0	1	2	3
1	0	0.5		
2	0			
3	0			???

But now we can calculate $v(2,1)$, the value of the game if we play it with the urn marked two reds and one green. Consider again what can happen when you pull out a marble. With probability $2/3$ you'll remove a red marble, in which case you pay me a dollar and you are left with an urn with 1 red and 1 green. On the other hand, with probability $1/3$ you'll remove a green marble, in which case I pay you a dollar but you're left with an urn with 2 red and 0 greens. Hence:

$$\begin{aligned} v(2,1) &= (2/3)\big(-1+v(1,1)\big)+(1/3)\big(1+v(2,0)\big) \\ &= (2/3)(-1+0.5)+(1/3)(1+0) \;=\; 0.0. \end{aligned}$$

Our table becomes

		Number of Greens		
	0	1	2	3
Number of Reds 0	0	1	2	3
1	0	0.5		
2	0	0.0		
3	0			???

Now we calculate $v(3,1)$, the value of the game if we play it with the urn marked three reds and one green. Here, **if** you remove a marble, with probability 3/4 you'll get a red marble, in which case you pay me a dollar and you are left with an urn with 2 reds and 1 green. On the other hand, with probability 1/4 you'll get a green marble, in which case I pay you a dollar but you're left with an urn with 3 reds and 0 greens. Hence:

$$\begin{aligned} v(3,1) &= (3/4)\big(-1+v(2,1)\big)+(1/4)\big(1+v(3,0)\big) \\ &= (3/4)(-1+0.0)+(1/4)(1+0.0) \;=\; -0.5. \end{aligned}$$

But wait a minute! That was **if** you remove a marble. You don't have to draw any marbles from the urn. In this case the value *if* you choose to play is negative, so you would simply decline to play the $(3,1)$ game yielding that $v(3,1)=0$. Our table is now

		Number of Greens		
	0	1	2	3
Number of Reds 0	0	1	2	3
1	0	0.5		
2	0	0.0		
3	0	0.0		???

We may now go on and compute, in this order, the quantities $v(1,2)$, $v(2,2)$, $v(3,2)$, $v(1,3)$, $v(2,3)$, and, finally, $v(3,3)$. At each step, say the urn with r red and g green marbles, we compute the value *if* you choose to remove a marble as

$$\tilde{v} \;=\; \big(r/(r+g)\big)\big(-1+v(r-1,g)\big) \;+\; \big(g/(r+g)\big)\big(1+v(r,g-1)\big).$$

If $\tilde{v}$ is a positive number, you elect to remove a marble and have $v(r,g)=\tilde{v}$. Otherwise, $\tilde{v}\le 0$ and you walk away from the game so $v(r,g)=0$. When we fill in our table this way we get

Number of Reds \ Number of Greens	0	1	2	3
0	0	1	2	3
1	0	0.5	1.3333	2.25
2	0	0.0	0.6667	1.5
3	0	0.0	0.2000	0.85

This analysis not only produces the answer, but it reveals the optimal strategy for stopping when your are playing with the urn containing 3 reds and 3 greens. Clearly one should stop when the value of playing, as given in the table above, is 0. This happens when there is: (1) three reds and one green left in the urn (i.e., the first two marbles removed were green – case (i) earlier) OR (2) there are two reds and one green left in the urn (i.e., two of the first three marbles removed were green – case (ii) earlier) OR (3) there are no greens left in the urn (case (iii) earlier).

A Little C Program. Try running the following C program:

```
#include "stdio.h"
main () {
    double v(int, int);
    printf ("Game value is %8.5f", v(3,3));
}
double v(int r, int g) {
    double v0, probgreen, probred;
    if (r == 0) {
      v0 = (double) g;
    } else if (g == 0) {
      v0 = (double) 0;
    } else {
      probred = ((double) r) / (r+g)
      probgreen = ((double) g) / (r+g);
      v0 = probred * (-1 + v(r-1, g)) + probgreen * (1 + v(r, g-1));
      if (v0 < 0) {
        v0 = 0;
      }
    }
    return v0;
}
```

References. For more information on this and some related problems see:

Boyce, W.M., *Stopping rules for selling bonds*, Bell Journal of Economics and Management Science **1** (1970), 27–53.

Chen, R.W., A. Zame, A.M. Odlyzko, and L.A. Shepp, *An optimal acceptance policy for an urn scheme*, SIAM Journal on Discrete Mathematics **11** (1998), 183–195.

Shepp, L.A., *Explicit solutions to some problems of optimal stopping*, Annals of Mathematical Statistics **40** (1969), 993–1010.