Food Court - SOLUTION

The food court diagrams represent two Discrete-Time Markov Chains (DTMC). The arrows indicate positive transition probabilities. For instance, starting at restaurant 13, there is a probability p5 of going to 7 in the next time period, and 1-p5 of going to 9. The idea is to solve for p1,...,p12, and then determine steady-state solutions for each food court.

Most of the transition probabilities can be calculated with simple formula that can be found on Wikipedia. The exceptions can be reasoned out.

p1. Alice received 15 votes for best chili in a landslide victory over Bob, who only got 4 votes. The ballots were counted one-by-one in random order. What is the probability that Alice was always strictly ahead of Bob in votes during the counting process?

Solution. $\frac{15-4}{15+4} = \frac{11}{19}$. See Bertrand's Ballot Theorem (`https://en.wikipedia.org/wiki/Bertrand%27s_ballot_theorem`).

p2. To decide who gets to manage the cash register, Alice, Bob, and Carol take turns flipping a coin, in alphabetical order (Alice→ Bob→ Carol→ Alice, etc). The first to flip heads wins. What is the probability that Carol wins?

Solution. $\frac{1}{7}$. Carol gets a chance to flip only if Alice and Bob both flip tails, so 1/4 of the time. Carol's odds of winning, p_C, are thus 1/4(1/2+ 1/2(odds of winning after Carol flips tails the first time)). Then from conditional independence we have $p_C = 1/4(1/2 + 1/2p_C)$, so $p_C = 1/7$.

p3. Alice picks pizza toppings by spinning a bottle. The bottle can point to one of six different toppings, spaced evenly around the bottle. Suppose four pizza toppings are chosen in this way. What is the probability that (at least) one topping is chosen twice or more?

Solution. $1 - (1 - \frac{1}{6})(1 - \frac{2}{6})(1 - \frac{3}{6}) = \frac{13}{18}$. Several approaches are possible; for instance, we can view it as a generalized birthday/collision problem (`https://en.wikipedia.org/wiki/Birthday_problem#Cast_as_a_collision_problem`).

p4. Alice absentmindedly cut a carrot at some random point along its length, yielding two pieces. What is the probability that the longer piece is at least 8 times the length of the shorter piece?

Solution. $\frac{2}{9}$. Suppose the carrot is oriented horizontally, and is cut into a left and right piece. The probability the left half is at least 8 times the length of the right is the probability the cut was applied in the first ninth of the carrot, so 1/9. By symmetry, the right half is at least 8 times the left with the same probability.

p5. To pass the time during a lunch break, Alice plays a (fair) game of Three-Card Monte with Bob and loses track of the target card, and so chooses a card at random. Bob reveals to Alice that one of the other two cards is NOT the target, and gives Alice the opportunity to switch to the remaining card instead. Suppose Bob would've made such an offer regardless of Alice's initial choice. What is the probability Alice wins by switching?

Solution. $\frac{2}{3}$. See the Monty Hall problem (`https://en.wikipedia.org/wiki/Monty_Hall_problem`).

p6. Alice, Bob, and Carol draw straws to decide which one of them has to write the new menu. They have, respectively, a 1/10, 2/5, and 3/5 chance of making typos on a menu. If the menu ends up with typos, what is the (conditional) probability that Alice wrote it?

Solution. P(Alice | typos) =P(typos | Alice)*P(Alice)/P(typos) $= (1/10) * (1/3)/[(1/3) * (1/10 + 2/5 + 3/5)] = \frac{1}{11}$. See Bayes' Theorem (`https://en.wikipedia.org/wiki/Bayes%27_theorem`).

p7. Alice noticed that it takes, on average, 1 minute to serve a customer; moreover, customers arrive, on average, every 2 minutes. Alice can only serve one customer at a time, and so additional customers wait in line. Suppose that both the time between customer arrivals as well as the service time are exponentially distributed. What is the probability that (at a randomly selected point in time) there are exactly 2 customers waiting in line and/or being served?
Solution. $(1-\frac{1}{2})(\frac{1}{2})^2 = \frac{1}{8}$ See M/M/1 queues (`https://en.wikipedia.org/wiki/M/M/1_queue#Average_number_of_customers_in_the_system`).

p8. Alice got 16 coins from tips, and Bob got 3. They decide to play the following game. Each player chooses one of their own coins. They flip their coins until one is heads, the other is tails. The person with heads keeps both coins. This is repeated until one player gets all the coins. What is the probability that Alice loses her tips to Bob?
Solution. $\frac{3}{16+3} = \frac{3}{19}$ See Gambler's Ruin `https://en.wikipedia.org/wiki/Gambler%27s_ruin#Fair_coin_flipping`.

p9. Alice plays rock-paper-scissors with Bob to decide who has to take out the garbage. Bob flips a coin, and, if it is heads Bob throws rock; otherwise if the coin is tails Bob employs some mixed strategy. Alice knows about this coin strategy but not the outcome of the coin flip. Suppose the players are in Nash equilibrium. What is the probability that Alice throws paper?
Solution. $\frac{2}{3}$. **Update: thanks to Sune Reeh for catching a wrong assumption, solution is fixed.** Let r_A, p_A, s_A denote the probability Alice throws rock, paper, or scissors respectively; likewise for Bob, conditioned on the coin flipping tails, assign probabilities r_B, p_B, s_B. Then, Alice's payoff z_A is

$$z_A := r_A(\frac{1}{2}(s_B - p_B)) + p_A(\frac{1}{2}(r_B - s_B) + \frac{1}{2}) + s_A(\frac{1}{2}(p_B - r_B) - \frac{1}{2}),$$

where the $\frac{1}{2}$ values reflect the outcome of the coin. We can work through the complementarity conditions to attain a solution (i.e. testing cases where probabilities are set to 0 or 1). First-order conditions give us the following

$$s_B < p_B \implies r_A = 0, p_B < 1 \implies s_A = 0.$$

Now if $p_B = 1$, then $s_B = r_B = 0$ and so Alice's payoff is $z_A = -\frac{1}{2}r_A + \frac{1}{2}p_A$ implying a best response $r_A = 0, p_A = 1, s_A = 0$. Bob's payoff z_B is (this is zero-sum so $z_B = -z_A$ but this rearrangement is helpful to look at):

$$z_B := \frac{1}{2}(1 + r_B)(s_A - p_A) + \frac{1}{2}p_B(r_A - s_A) + \frac{1}{2}s_B(p_A - r_A).$$

If $p_A = 1$, then Bob's best response is $s_B = 1$, so $p_B = 1$ cannot occur at equilibrium. Hence at equilibrium we have $p_B < 1, s_A = 0$. Updating the payoffs:

$$z_A = r_A(\frac{1}{2}(s_B - p_B)) + p_A(\frac{1}{2}(r_B - s_B) + \frac{1}{2}),$$

$$z_B = \frac{1}{2}(1 + r_B)(-p_A) + \frac{1}{2}p_B r_A + \frac{1}{2}s_B(p_A - r_A).$$

By similar reasoning, if $p_A = 0$ (in addition to $s_A = 0$), then Bob's best response is $p_B = 1$, to which Alice's best response is $p_A = 1$. So at equilibrium $p_A > 0$, which implies Bob's best response (at equilibrium) is $r_B = 0$. Updating the payoffs again, we have

$$z_A = r_A(\frac{1}{2}(s_B - p_B)) + p_A(-\frac{1}{2}s_B + \frac{1}{2}),$$

$$z_B = -\frac{1}{2}p_A + \frac{1}{2}p_B r_A + \frac{1}{2}s_B(p_A - r_A).$$

Now, using the same logic as before, we can rule out equilibria where any of the remaining probabilities are set to zero, e.g. $p_A = 0 \implies p_B = 1 \implies p_A = 1$. Since all remaining probabilities are strictly positive, we have from Alice and Bob's first-order conditions (respectively):

$$\frac{1}{2}(s_B - p_B) = -\frac{1}{2}s_B + \frac{1}{2} \iff 2s_B = 1 + p_B,$$

$$\frac{1}{2}r_A = \frac{1}{2}(p_A - r_A) \iff p_A = 2r_A$$

Together with $r_A + p_A + s_A = 1, r_B + p_B + s_B = 1$, we attain the unique Nash solution

$$r_A = \frac{1}{3}, p_A = \frac{2}{3}, s_A = 0,$$

$$\implies r_B = 0, p_B = \frac{1}{3}, s_B = \frac{2}{3}.$$

An alternative route is to modify Alice's payoff matrix for standard rock-paper-scissors:

$$\begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix}.$$

Half the time they are playing standard RPS, half the time Alice faces only rock. Recall that in a Nash equilibrium, players are maximizing their expected outcome. Thus Alice should be considering the average of the two payoffs:

$$\begin{bmatrix} 0 & 1 & -1 \\ -1/2 & 1/2 & 0 \\ 1/2 & 0 & -1/2 \end{bmatrix}.$$

This is still a zero-sum game, since Alice wins only if Bob loses, and vice versa. Hence one can plug this payoff matrix into a zero-sum Nash equilibrium solver or apply linear programming, i.e. let the computer slog through complementarity conditions.

p10. Alice knows there is exactly one mouse in the food court. On a given day, the mouse can decide: to leave (10% chance), stay (60% chance), or invite an additional mouse friend to stay with them (30% chance). All mice (independently) act in the same manner. What is the probability that (absent outside intervention) the food court will ever be free of mice?
Solution. $\frac{1}{3}$. Taken from the extinction problem example on Wikipedia (`https://en.wikipedia.org/wiki/Branching_process#Example_of_extinction_problem`).

p11. Alice wants to know if a bag of onions is still fresh. A fresh bag of 7-week-old onions spoils within 2 weeks 19% of the time, and has a 73% chance of spoiling within 3 weeks. Furthermore, a fresh bag of 8-week-old onions has a 30% chance of remaining fresh for at least 2 weeks. What is the probability that a fresh bag of 8-week-old onions will remain fresh for at least 1 week?
Solution. $\frac{9}{10}$. (This question was derived from the Waterloo Euclid Math Contest 2011 question 7a). A fresh 7-week-old bag has a 100%-19%=81% chance of staying fresh after 2 weeks, and is fresh for at least 3 weeks with 27% probability. Thus a fresh 9-week-old bag survives at least one week with probability 27/81 = 1/3. By the same reasoning, a fresh 8-week-old bag has a 30% chance of remaining fresh for at least 2 weeks, and so a fresh 8-week-old bag is fresh for at least 1 week with probability 0.3/(1/3)=0.9.

p12. Alice has a 40% chance of working a given shift at the food court, Bob has a 30% chance. The chance that either Alice and/or Bob are working a shift is 50%. What is the probability that Alice and Bob are working together?

Solution. $0.4+0.3-0.5 = \frac{1}{5}$. See Inclusion-Exclusion `https://en.wikipedia.org/wiki/Inclusion%E2%80%93exclusion_principle`.

After solving for the above probabilities, we obtain the following From/To transition probability matrices for each Markov chain:

	1	3	5	7	8	9	12	13	15	20
1	0	0	0	0	0	0	0	0	0	1
3	0	0	0	0	0	1	0	0	0	0
5	0	0	0	11/19	0	0	0	0	0	8/19
7	1/7	0	0	0	6/7	0	0	0	0	0
8	0	0	13/18	0	0	0	2/9	0	1/18	0
9	0	0	1	0	0	0	0	0	0	0
12	0	0	0	0	0	0	0	1	0	0
13	0	0	0	2/3	0	1/3	0	0	0	0
15	0	0	0	0	0	0	1	0	0	0
20	0	1/11	0	0	0	0	0	10/11	0	0

	2	4	6	10	11	14	16	17	18	19
2	7/8	0	0	1/8	0	0	0	0	0	0
4	0	0	0	0	0	0	1	0	0	0
6	0	0	0	0	3/19	0	0	16/19	0	0
10	0	0	0	0	0	0	0	0	0	1
11	0	0	0	0	0	0	1/3	0	2/3	0
14	0	0	0	0	0	0	0	2/3	1/3	0
16	0	0	1	0	0	0	0	0	0	0
17	0	9/10	0	0	0	1/10	0	0	0	0
18	1/5	0	0	0	0	4/5	0	0	0	0
19	0	0	0	0	0	0	0	0	1	0

The steady-state of a Markov chain is the long-run probability (i.e. as time goes to infinity) of ending up in a given state (regardless of the initial starting point). The puzzle's chains have been designed to ensure such states exist. Given transition matrix P, the steady state π is the solution to $\pi P = \pi$. They can be found by computing a left eigenvector of the (unique) eigenvalue equal to 1, and then scaling it to be a probability vector (divide by the sum of its entries). Note that it is also possible that a puzzle-solver attempts to find the continuous-time steady state and gets nonsense. Hopefully if they get so far, they also try the aforementioned discrete-time formula. The steady state probabilities are:

$$\pi_1 = 0.03$$
$$\pi_3 = 0.01$$
$$\pi_5 = 0.19$$
$$\pi_7 = 0.21$$
$$\pi_8 = 0.18$$
$$\pi_9 = 0.06$$
$$\pi_{12} = 0.05$$
$$\pi_{13} = 0.15$$
$$\pi_{15} = 0.01$$
$$\pi_{20} = 0.11$$

$$\pi_2 = 0.08$$
$$\pi_4 = 0.18$$
$$\pi_6 = 0.19$$
$$\pi_{10} = 0.01$$
$$\pi_{11} = 0.03$$
$$\pi_{14} = 0.06$$
$$\pi_{16} = 0.19$$
$$\pi_{17} = 0.20$$
$$\pi_{18} = 0.05$$
$$\pi_{19} = 0.01$$

If done correctly, all steady-state probabilities can be exactly expressed with two digits of precision — this is hinted in the puzzle with probabilities of 1.00 in the chain. Converting these two digits to letters (0.01=A, 0.02=B, etc.), and using the lexicographic order of nodes yields "CHARS SURFACE OF A STEAK", which clues the final answer SEARS.