Online Matching Coins

Steven Finch

March 7, 2014

The first game we discuss originated in [1, 2], although we mostly follow [3] in our exposition. The second and third games appear in [4].

Nature tosses a fair coin repeatedly and independently, yielding an infinite sequence N_0, N_1, N_2, $\ldots$ of 1s and 2s. Just prior to each toss, Alice and Bob simultaneously declare their guess A and B for the resulting N. They win the toss if both guessed correctly. Their goal is to maximize the probability of winning. They are permitted to strategize no later than one hour beforehand; after the game starts, any communication between them is only via the As, Bs and Ns. With no further information, if they agree beforehand to always both guess 1 (for example) then the probability of winning is 1/2. No improvement is possible.

Suppose now that, during their strategizing, Alice and Bob are told that Alice will be given the full sequence N_0, N_1, N_2, $\ldots$ one minute before the game! To improve their odds, Alice must pass relevant information she knows to Bob in an agreed-upon manner. Setting $A = N$ always does not help their cause! At toss 0, Alice might declare

$$A_0 = N_2$$

(sacrificing her knowledge of N_0) so that Bob understands to declare $B_1 = B_2 = A_0$. They will win toss 2 since Alice will declare $A_2 = N_2$. At toss 1, Alice might declare

$$A_1 = N_4$$

(sacrificing her knowledge of N_1) so that Bob understands to declare $B_3 = B_4 = A_1$. They will win toss 4 since Alice will declare $A_4 = N_4$. At toss 3, Alice might declare

$$A_3 = N_6$$

(sacrificing her knowledge of N_3) so that Bob understands to declare $B_5 = B_6 = A_3$. They will win toss 6 since Alice will declare $A_6 = N_6$, and so forth. In summary, Alice and Bob will score one win out of two whenever $\{N_{2t+1}, N_{2t+2}\} = \{1, 2\}$ or $\{2, 1\}$. When $\{N_{2t+1}, N_{2t+2}\} = \{1, 1\}$ or $\{2, 2\}$, they will score one win out of two half the time and two out of two the remaining half, giving odds of

$$\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot \frac{\frac{1}{2} + 1}{2} = \frac{5}{8} = 0.625.$$

Instead of partitioning time into blocks modulo 2, let us do so modulo 3. Define the mode M_t of $\{N_{3t+1}, N_{3t+2}, N_{3t+3}\}$ to be the most common element in the set. At toss 0, Alice might declare

$$A_0 = M_0$$

(sacrificing her knowledge of N_0) so that Bob understands to declare $B_1 = B_2 = B_3 = M_0$. Assume that indices $1 \leq i, j, k \leq 3$ are distinct. Alice's next three declarations might be

$$A_i = A_j = M_0 \text{ and } A_k = M_1 \qquad \text{if } N_k \neq M_0$$

and

$$A_1 = A_2 = M_0 \text{ and } A_3 = M_1 \qquad \text{if } N_1 = N_2 = N_3 = M_0$$

(sacrificing her knowledge of N_3 for the latter) so that Bob understands to declare $B_4 = B_5 = B_6 = M_1$. In summary, Alice and Bob will score two wins out of three whenever $\{N_1, N_2, N_3\}$ contains two 1s and one 2, or two 2s and one 1. When $\{N_1, N_2, N_3\}$ contains all 1s or all 2s, they will score two wins out of three half the time and three out of three the remaining half, giving odds of

$$\frac{3}{4} \cdot \frac{2}{3} + \frac{1}{4} \cdot \frac{\frac{2}{3} + 1}{2} = \frac{17}{24} = 0.7083....$$

A more sophisticated strategy allows the win probability to approach $x = 0.8107103750...$ as closely as desired, where x is the unique solution of the equation [1]

$$-x \ln(x) - (1 - x) \ln(1 - x) + (1 - x) \ln(3) = \ln(2).$$

No further improvement is possible beyond this point.

0.1. Symmetric Online Matching Coins. The preceding game is asymmetric – Alice knows everything and Bob knows nothing – for the following game, information is distributed equally among the players and they will both need to send signals to each other. Imagine here that a fair coin has four equally-likely sides, not two. (A regular tetrahedral die would be a better metaphor.) Also define

$$f(N) = \begin{cases} 1 & \text{if } N = 1 \text{ or } 3, \\ 0 & \text{if } N = 2 \text{ or } 4, \end{cases} \qquad g(N) = \begin{cases} 1 & \text{if } N = 1 \text{ or } 2, \\ 0 & \text{if } N = 3 \text{ or } 4 \end{cases}$$

for convenience, that is, $f(N)$ anwers the question "Is N odd?" and $g(N)$ answers the question "Is $N \leq 2$?"

Nature tosses a fair coin repeatedly and independently, yielding an infinite sequence N_1, N_2, N_3, ... of 1s, 2s, 3s and 4s. Just prior to each toss, Alice and Bob simultaneously declare their guess A and B for the resulting N. They win the toss if both guessed correctly. Their goal is to maximize the probability of winning. They are permitted to strategize no later than one hour beforehand; after the game starts, any communication between them is only via the As, Bs and Ns. With no further information, if they agree beforehand to always both guess 1 (for example) then the probability of winning is 1/4. No improvement is possible.

Suppose now that, during their strategizing, Alice and Bob are told that Alice will be given the sequence $f(N_1)$, $f(N_2)$, $f(N_3)$, $\ldots$ and Bob will be given the sequence $g(N_1)$, $g(N_2)$, $g(N_3)$, $\ldots$ one minute before the game! At toss 1, Alice might declare

$$A_1 = \begin{cases} 1 & \text{if } f(N_1) = 1 \text{ and } f(N_2) = 0, \\ 2 & \text{if } f(N_1) = 0 \text{ and } f(N_2) = 1, \\ 3 & \text{if } f(N_1) = 1 \text{ and } f(N_2) = 1, \\ 4 & \text{if } f(N_1) = 0 \text{ and } f(N_2) = 0 \end{cases}$$

and Bob might declare

$$B_1 = \begin{cases} 1 & \text{if } g(N_1) = 1 \text{ and } g(N_2) = 0, \\ 2 & \text{if } g(N_1) = 0 \text{ and } g(N_2) = 1, \\ 3 & \text{if } g(N_1) = 1 \text{ and } g(N_2) = 1, \\ 4 & \text{if } g(N_1) = 0 \text{ and } g(N_2) = 0 \end{cases}$$

so that they will win toss 2. The odds here are

$$\frac{1}{2}\left(\frac{1}{2} \cdot \frac{1}{2} + 1\right) = \frac{5}{8} = 0.625.$$

Instead of devoting resources to guessing N_1, let us shift emphasis entirely to signaling ahead for N_2 and N_3. At toss 1, Alice might declare

$$A_1 = \begin{cases} 1 & \text{if } f(N_2) = 1 \text{ and } f(N_3) = 0, \\ 2 & \text{if } f(N_2) = 0 \text{ and } f(N_3) = 1, \\ 3 & \text{if } f(N_2) = 1 \text{ and } f(N_3) = 1, \\ 4 & \text{if } f(N_2) = 0 \text{ and } f(N_3) = 0 \end{cases}$$

and Bob might declare

$$B_1 = \begin{cases} 1 & \text{if } g(N_2) = 1 \text{ and } g(N_3) = 0, \\ 2 & \text{if } g(N_2) = 0 \text{ and } g(N_3) = 1, \\ 3 & \text{if } g(N_2) = 1 \text{ and } g(N_3) = 1, \\ 4 & \text{if } g(N_2) = 0 \text{ and } g(N_3) = 0 \end{cases}$$

(both sacrificing their partial knowledge of N_1) so that they will win tosses 2 and 3. The odds here are

$$\frac{1}{3}\left(\frac{1}{4} \cdot \frac{1}{4} + 1 + 1\right) = \frac{33}{48} = 0.6875.$$

A more sophisticated strategy allows the win probability to approach $\kappa = 0.7337221510...$ as closely as desired [4]. The formulas underlying this constant are more elaborate than before. Define a hyperplanar region in $\mathbb{R}^8$:

$$\Delta(8) = \left\{ (x_1, x_2, \ldots, x_8) : \sum_{\ell=1}^{8} x_\ell = 1 \text{ and } x_\ell \geq 0 \text{ for all } \ell \right\}$$

and a real-valued function on $\Delta(8)$:

$$h(x) = -\frac{1}{\ln(2)} \sum_{\ell=1}^{8} x_\ell \ln(x_\ell)$$

with the convention that $0 \cdot \ln(0) = 0$. Let $\varphi : [0, 3] \to \mathbb{R}$ be given by

$$\varphi(r) = \max \left\{ \sum_{\ell=1}^{4} x_\ell^2 : x \in \Delta(8) \text{ and } h(x) \geq r \right\}$$

and let $\psi : [0, 3] \to \mathbb{R}$ be the minimal concave function $\geq \varphi$. The desired probability κ is $\psi(1)$, which numerically appears to be equal to $\varphi(1)$. No further improvement is possible beyond this point. It also appears that the minimizing vector x can be taken such that $x_1 = x_2 = x_3$ and $x_5 = x_6 = x_7$, which would simplify our presentation.

0.2. Cross Over Matching Coins. Here the game is symmetric, as for the preceding, but Nature instead tosses a *pair* of distinguishable coins (two sides apiece). Thus we have two infinite sequences N_0^α, N_1^α, N_2^α, $\ldots$ and N_0^β, N_1^β, N_2^β, $\ldots$ of 1s and 2s. Just prior to each toss, Alice and Bob simultaneously declare their guesses A and B for the resulting N^α and N^β, respectively. During their one-hour prior strategizing, they learn that Alice will be given the sequence N_0^β, N_1^β, N_2^β, $\ldots$ and Bob will be given the sequence N_0^α, N_1^α, N_2^α, $\ldots$ at one-minute prior! Their goal is to maximize the average of (the probability of Alice winning) and (the probability of Bob winning). Communication between them, via the As, Bs, N^αs and N^βs, is again critical to their success.

The optimal win probability here is $\lambda = 0.8041565330...$ [4]. Define a line segment in $\mathbb{R}^2$:

$$\Delta(2) = \left\{ (x_1, x_2) : \sum_{\ell=1}^{2} x_\ell = 1 \text{ and } x_\ell \geq 0 \text{ for all } \ell \right\}$$

and a real-valued function on $\Delta(2)$:

$$h(x) = -\frac{1}{\ln(2)} \sum_{\ell=1}^{2} x_\ell \ln(x_\ell).$$

Let $\varphi : [0,1] \to \mathbb{R}$ be given by

$$\varphi(r) = \max\left\{\sum_{\ell=1}^{2} x_\ell^2 : x \in \Delta(2) \text{ and } h(x) \geq r\right\}$$

and let $\psi : [0,1] \to \mathbb{R}$ be the minimal concave function $\geq \varphi$. The desired probability λ is $\psi(1/2)$, which is (in this case) provably equal to $\varphi(1/2)$.

A simpler presentation is hence clear: $\lambda = y^2 + (1-y)^2$ where y is either of the two reals satisfying

$$-2y\ln(y) - 2(1-y)\ln(1-y) = \ln(2).$$

No closed-form expression for this constant (or for other constants in this essay) seems to be available.

It is possible to generalize the symmetric online game to an arbitrary number m of players and a single n^m-sided coin. The real-valued function h on $\Delta(n^{m+1})$ gives rise to a φ (maximum sum of m^{th} powers, indices from 1 to n^m) and a minimal concave $\psi \geq \varphi$. For $m > 2$ or $n > 2$, however, $\psi(\ln(n)/\ln(2))$ is strictly greater than $\varphi(\ln(n)/\ln(2))$. This complicates the numerical calculation of a optimal win probability in the general setting.

It is also possible to generalize the cross over matching game to a pair of n-sided coins. The real-valued function h on $\Delta(n)$ gives rise to a φ (maximum sum of n squares) and a minimal concave $\psi \geq \varphi$. For $n > 2$, however, $\psi(\ln(n)/(2\ln(2)))$ is strictly greater than $\varphi(\ln(n)/(2\ln(2)))$. This again complicates calculations in general.

Related ideas appear in [5] (best strategies) and [6] (maximal convex function $\leq \varphi$).

References

[1] O. Gossner, P. Hernández and A. Neyman, Online matching pennies, Center for the Study of Rationality, Hebrew Univ., 2003, http://ratio.huji.ac.il/sites/default/files/publications/dp316.pdf.

[2] O. Gossner, P. Hernández and A. Neyman, Optimal use of communication resources, *Econometrica* 74 (2006) 1603–1636; MR2268411 (2007e:91112).

[3] P. Winkler, *Mathematical Mind-Benders*, A. K. Peters, 2007, pp. 11, 17–18; MR2334790 (2008f:00002).

[4] A. Shapira, *Communication Games with Asymmetric Information*, Ph.D. thesis, Hebrew Univ., 2008, http://shemer.mslib.huji.ac.il/dissertations/W/JMC/001493090_1.pdf.

[5] S. R. Finch, Optimal stopping constants, *Mathematical Constants*, Cambridge Univ. Press, 2003, pp. 361–363.

[6] S. R. Finch, Shapiro-Drinfeld constant, *Mathematical Constants*, Cambridge Univ. Press, 2003, pp. 208–211.