Coin-Flipping Magic*

Nadia Benbernou† Erik D. Demaine† Martin L. Demaine† Benjamin Rossman†

Prepared in honor of Martin Gardner for Gathering for Gardner 8

Abstract

This paper analyzes a variety of generalizations of a coin-flipping magic trick invented independently by Martin Gardner and Karl Fulves. In the original trick, a blindfolded magician asks the spectator to flip three coins, forcing them into an all-equal state by surprisingly few moves. We generalize to any number of coins, coins/dice with more than two sides, and multiple flips at once. Next we consider a generalization described by Martin Gardner in which the spectator can re-arrange the coins in certain ways in between each flip. Finally we consider the variation in which the magician equalizes the number of heads and tails, which can be achieved exponentially faster.

1 Introduction

Figure 1: *Three pennies.*

The trick. Thank you, dear reader, for volunteering for a magic trick. May I ask, do you have a few coins that we could use for the trick? If not, you can borrow mine from Figure 1. Please get out three coins; they can be different denominations. Now please arrange the coins in a line from left to right. Very good. Now I will blindfold myself and look away. I guarantee that I cannot see the coins.

To get started, I would like you to flip over some of the coins. You can flip over as many or as few as you like. The only rule is that the coins should not be all heads or all tails. Let me know when you are finished. Good, let us proceed.

I am now visualizing your coins in my mind. With you acting as my hand, I will make the coins all the same: all heads or all tails.

Please flip over the left coin. Are the coins now all the same? One third of my readers will shout "yes!" and be blown away by my omniscience. For the rest of you, the trick continues.

Please flip over the middle coin. Very good. Now are the coins the same? Another third of my readers will be surprised by my fast success. For the rest of you, the trick continues.

Let me see; I must have made a mistake in visualizing your coins. Ah yes, I see. I shouldn't have flipped the left coin in the first place. Please flip over the left coin, back the way it was. Now, I tell you, the coins are all the same. Feel free to check my blindfolds.

Background. This self-working magic trick appears in Karl Fulves's book *The Children's Magic Kit* [Ful80, page 15]. According to that book, the trick was independently devised by Martin Gardner and Fulves, based on an idea of Sam Schwartz. It works with coins or cards, over the telephone or the radio.

At first we, and presumably many spectators, find it surprising that just three blindfolded flips are enough to equalize the three coins. Indeed, there are $2 \times 2 \times 2 = 8$ possible states of the coins (HHH, HHT, HTH,

*Preliminary draft. The latest version of this paper can be found at http://erikdemaine.org/papers/CoinFlipping/

†MIT Computer Science and Artificial Intelligence Laboratory, 32 Vassar St., Cambridge, MA 02139, USA, {nbenbern,edemaine,mdemaine,brossman}@mit.edu

THH, HTT, THT, TTH, TTT). How do we navigate to two of these states (HHH or TTT) using just three moves (and often fewer)? Motivated by this simple question, this paper studies several generalizations and variations of this magic trick.

Results. We begin in Section 2 with a simple generalization of the trick to n coins. This generalization, and the original trick, turn out to be easy to analyze: they are equivalent to Hamiltonian paths in an $(n-1)$-dimensional hypercube graph, and the classic Gray code gives one such solution. Not surprisingly, the number of moves required grows exponentially with n. More interesting is that we save an entire factor of two by having two goal states, all heads and all tails. Namely, the worst-case optimal number of blindfolded flips is $2^{n-1}-1$. The analysis also easily generalizes to k-sided dice instead of coins: in Section 3, we show that the worst-case optimal number of blindfolded operations is $k^{n-1}-1$. This family of tricks is thus really most impressive for $n=3$ coins, where the number 3 of flips is really quite small; beyond $n=3$ or $k=2$, the number of flips grows quickly beyond feasibility. For sake of illustration, however, Figure 2 shows a magic-trick sequence for four coins.

Equalizing four coins:
1. Flip the first coin.
2. Flip the second coin.
3. Flip the first coin.
4. Flip the third coin.
5. Flip the first coin.
6. Flip the second coin.
7. Flip the first coin.

Figure 2: *Equalizing four coins with at most seven flips.*

One solution to this exponential growth is to change the goal from the two all-heads and all-tails states to some larger collection of states. In Section 4, we consider a natural such goal: equalize the number of heads and tails. This goal is exponentially easier to achieve. Within just $n-1$ coin flips, the magician can force the numbers of heads and tails to be equal. The algorithm is simple: just flip one coin at a time, in any order, until the goal has been reached. Figure 3 shows an example for $n=6$. Although not obvious, this algorithm will succeed before every coin has been flipped once. Furthermore, by randomly flipping all the coins in each move, the magician expects to require only around $\sqrt{n}$ moves. The practicality of this type of trick clearly scales to much larger n.

Balancing six coins:
1. Flip the second coin.
2. Flip the fifth coin.
3. Flip the third coin.
4. Flip the first coin.
5. Flip the fourth coin.

Figure 3: *Equalizing the numbers of heads and tails in six coins using at most four flips.*

Returning to the goal of equalizing all the coins, the next generalization we consider is to allow the magician to flip more than one coin at once, between asking for whether the coins are yet all the same. This flexibility cannot help the magician to equalize the coins any faster than $2^{n-1}-1$ moves, but it can help obscure what the magician is doing. For example, if the magician had to repeat the three-coin trick several times, it might help to try some of the variations in Figure 4. Under what conditions can the magician still equalize n coins, ideally in the same number $2^{n-1}-1$ of moves? Obviously, flipping $n-1$ of the coins is equivalent to flipping just one coin. On the negative side, in Section 5, we show that the sequence of flips cannot vary arbitrarily: if the spectator is allowed to choose how many coins the magician should flip in each move, then the magician can succeed only if $n \leq 4$ (no matter how many moves are permitted). On the positive side, in Section 6, we show that it is possible to flip most fixed numbers of coins in every move and achieve the optimal worst case of $2^{n-1}-1$ moves. This result is fairly technical but interesting in the way that it generalizes Gray codes.

Varying number of coins flipped:
1. Flip the left and middle coins.
2. Flip the left coin.
3. Flip the left and middle coins.

Exactly two coins flipped:
1. Flip the left and middle coins.
2. Flip the middle and right coins.
3. Flip the left and middle coins.

Figure 4: *Alternate solutions to equalizing three coins with at most three moves, flipping more than one coin in some moves.*

The final variation we consider allows the spectator to re-arrange the coins in certain ways after each move. Again this can only make the magician's job harder. In each move, the magician specifies a subset of coins to flip, but before the spectator actually flips them, the spectator can re-arrange the coins according to certain permutations. Then, after flipping these coins, the spectator reveals whether all coins are the same, and if not, the trick continues. In Section 7, we characterize the exact group structures of allowed re-arrangements that still permit the magician to equalize the coins. For example, if 2^k coins are arranged on a table, then the spectator can rotate and/or flip the table in each move, and still the magician can perform the trick. Figure 5 shows the solution for four coins arranged in a square which the spectator can rotate by 0, 90°, 180°, or 270° during each move.

Equalizing four rotating coins:

1. Flip the north and south coins.
2. Flip the north and east coins.
3. Flip the north and south coins.
4. Flip the north coin.
5. Flip the north and south coins.
6. Flip the north and east coins.
7. Flip the north and south coins.

Figure 5: *Equalizing four coins that the spectator can rotate at each stage, using at most seven moves.*

The four-coin magic trick of Figure 5 goes back to a March 1979 letter from Miner S. Keeler to Martin Gardner [Gar91], and was disseminated more recently by Eric Roode [Roo02]. Keeler's letter was inspired by an article of [Gar91] about a somewhat weaker magic trick where, after the spectator turns the table arbitrarily, the magician can feel two coins before deciding which to flip. This weaker trick has been generalized to n coins on a rotating circular table and k hands: the trick can be performed if and only if $k \geq (1 - 1/p)n$ where p is the largest prime divisor of n [LW80, LR81]. The fully blind trick we consider, without the ability to feel coins, was first generalized beyond Keeler's four-coin trick to n dice, each with k sides, on a rotating circular table: the trick can be performed if and only if k and n are powers of a common prime [YEM93]. Interestingly, this characterization remains the same even if the magician can see the dice at all times (but the spectator can still turn the table before actually flipping coins); however, the worst-case number of moves reduces from $k^n - 1$ to $n + (\alpha - 1)(n - p^{\beta-1})$ where $k = p^\alpha$ and $n = p^\beta$. (Interestingly, the optimal number of moves for a specific sequence of coins gives rise to the notion of "word depth", now studied in the context of linear codes in information theory [Etz97, LFW00].)

Our general scenario considers an arbitrary "group" of permutations, instead of just a rotating circular table. This scenario was also essentially solved by Ehrenborg and Skinner [ES95]: they characterize performability in terms of the chain structure of the group. Our characterization is simpler: the group must have a number of elements equal to an exact power of two. Our proof is also simpler, showing a connection to invariant flags from group representation theory. It also uses the most sophisticated mathematical tools among proofs in this paper.

2 n Coins

The simplest generalization is to consider n coins instead of three. The goal is to make the coins all the same (all heads or all tails) by a sequence of single coin flips, where after each flip the magician asks "are the coins all the same yet?"

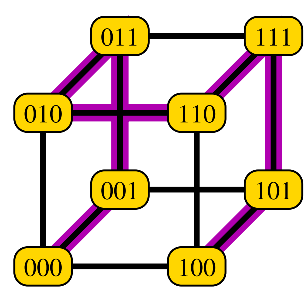

Figure 6: *The graph corresponding to the three-coin trick: the 3-dimensional binary cube.*

We can visualize this problem as navigating a graph, where each vertex corresponds to a possible state of all the coins, and an edge corresponds to flipping a single coin. This graph is the well-known *n-dimensional binary hypercube*; Figure 6 shows the case $n = 3$. In general, the n-dimensional binary hypercube has 2^n vertices, one for each possible binary string of length n (where 0 bits correspond to heads and 1 bits correspond to tails), and has an edge between two vertices whose binary strings differ in exactly one bit.

In the magic trick, the spectator chooses an arbitrary start vertex, and the magician's goal is to reach one of two goal vertices: $00 \cdots 0$ (all heads) or $11 \cdots 1$ (all tails). At each step, the magician can pick which edge to move along: flipping the ith coin corresponds to flipping the ith bit. The only feedback is when the magician hits one of the goal vertices.

An equivalent but more useful viewpoint is to map coin configurations onto the binary hypercube by defining a bit in a binary string to be 0 if that coin is the same orientation (heads/tails) as the spectator's original choice, and 1 if the coin is different from the spectator's original choice. In this view, the magician always starts at the same vertex $00 \cdots 0$. The two goal configurations g and $\overline{g}$ are now the unknown part; the only knowledge is that they are inversions of each other (with 0s turned into 1s and vice versa).

In order for the magician to be sure of visiting one of the two solution states, the chosen path (sequence of coin flips) must visit either v or its inversion $\overline{v}$ for every vertex v in the hypercube. There are 2^n total vertices, so the path must visit $2^n/2 = 2^{n-1}$ different vertices. This argument proves a worst-case lower bound of $2^{n-1} - 1$ flips in the worst-case execution of the magic trick.

To see that $2^{n-1} - 1$ flips also suffice in the worst case, it suffices to find a Hamiltonian path in any $(n-1)$-dimensional subcube of the n-dimensional cube, dropping whichever dimension we prefer. (A Hamiltonian path visits each vertex exactly once.) The spectator sets this dimension arbitrarily to heads or tails, and the Hamiltonian path explores all possible values for the remaining $n - 1$ bits, so eventually we will reach the configuration in which all bits match the dropped bit. The subcube has 2^{n-1} total vertices, so the presumed Hamiltonian path has exactly $2^{n-1} - 1$ edges as desired.

The final piece of the puzzle is that n-dimensional cubes actually have Hamiltonian paths. This fact is well-known. One such path is given by the *binary Gray code*, also known as the reflected binary code [Gra53]. This code/path can be constructed recursively as follows. The n-bit Gray code first visits all strings starting with 0 in the order given by the $(n-1)$-bit Gray code among the remaining bits; then it visits all strings starting with 1 in the reverse of the order given by the $(n-1)$-bit Gray code among the remaining bits. For example, the 1-bit Gray code is just 0, 1; the 2-bit Gray code is 00, 01, 11, 10; and the 3-bit Gray code is 000, 001, 011, 010, 110, 111, 101, 100. Figure 6 illustrates this last path.

Theorem 1 *The optimal sequence of flips guaranteed to eventually make n coins all heads or all tails uses exactly $2^{n-1} - 1$ flips in the worst case.*

3 n Dice

One natural extension of flipping coins is to rolling k-sided dice. Suppose we have n dice, each die has k faces, and each face is labeled uniquely with an integer $0, 1, \ldots, k-1$. The spectator arranges each die with a particular face up. As before, the magician is blindfolded. In a single move, the magician can increment or decrement any single die by ± 1 (wrapping around from k to 0). At the end of such a move, the magician asks whether all the dice display the same value face up. The magician's goal is to reach such a configuration.

As we show in this section, our analysis extends fairly easily to show that the magician can succeed in $k^{n-1} - 1$ steps. A configuration of n dice becomes a k-ary string of n digits between 0 and $k - 1$. In the most useful viewpoint, a digit of 0 represents the same as the original state chosen by the spectator, and a digit of i represents that the die value is i larger (modulo k) than the original die value. Thus $(0, 0, \ldots, 0)$ represents the initial configuration chosen by the spectator, and the k goal states $g_0, g_1, \ldots, g_{k-1}$ have the property that g_i corresponds to adding i to each entry of g_0 (modulo k).

The analogous graph here is the k-ary n-dimensional torus. Figure 7 shows the case $n = k = 3$. In general, the vertices correspond to k-ary strings of length n, and edges connect two vertices $a = (a_1, a_2, \ldots, a_n)$ and $b = (b_1, b_2, \ldots, b_n)$ that differ by exactly ± 1 (modulo k) in exactly one position: $b = (a_1, a_2, \ldots, a_{i-1}, a_i \pm 1, a_{i+1}, \ldots, a_n)$.

Again we drop an arbitrary digit/dimension, and focus on the resulting k-ary $(n-1)$-dimensional subtorus. The magic trick becomes equivalent to finding a Hamiltonian path in this subtorus. Such a path exists by a natural generalization of the Gray code [Gua98]. Visiting all configurations of the other dice will eventually match the value of the dropped dimension.

Theorem 2 *The optimal sequence of die increments/decrements guaranteed to eventually make n k-sided dice all the same uses exactly $k^{n-1} - 1$ moves in the worst case.*

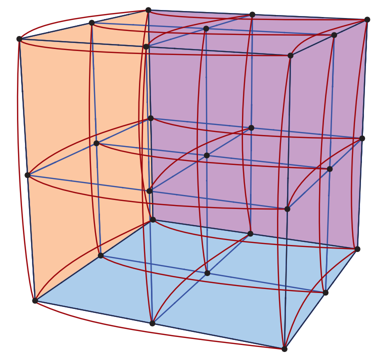

Figure 7: 3-*ary* 3-*dimensional torus.*

4 Equal Numbers of Heads and Tails

In this section, we explore the variation of the magic trick in which the goal is to equalize the numbers of heads and tails, instead of equalizing the coins themselves. We consider two strategies: a fast randomized strategy, and a simple deterministic strategy that achieves a balanced configuration in linear time.

Call a configuration *balanced* if it has an equal number of heads and tails. Throughout, we assume n is even, although we could adapt these results to the case of n odd by expanding the definition of balanced configuration to allow the numbers of heads and tails to differ by at most 1.

4.1 Randomized Strategy

In the randomized strategy, in each move, the magician flips each coin with probability $\frac{1}{2}$. Thus the magician flips around half the coins in each move, effectively randomizing the entire configuration. We show that the magician reaches a balanced configuration in around $\sqrt{n}$ such moves:

Theorem 3 *Using the randomized strategy, the magician balances n coins within $O(\sqrt{n})$ steps with constant probability, and within $O(\sqrt{n}\lg n)$ steps with probability $1 - O(1/n^c)$ for any desired $c > 0$.*

Proof: For any two configurations a and b on n coins, a single move from a reaches b with probability $1/2^n$. Hence each move uniformly samples the configuration space. The number of balanced configurations is $\binom{n}{n/2}$, so the probability of reaching a balanced configuration in each step is $\binom{n}{n/2}/2^n$. We simply need to determine the number of such trials before one succeeds.

First we lower bound the probability of success using Stirling's formula:

$$\sqrt{2\pi n}\,(n/e)^n\, e^{1/(12n+1)} < n! < \sqrt{2\pi n}\,(n/e)^n\, e^{1/(12n)}.$$

Thus

$$\binom{n}{n/2} = \frac{n!}{(n/2)!(n/2)!} \geq \frac{\sqrt{2\pi n}\,(\frac{n}{e})^n\, e^{1/(12n+1)}}{\left(\sqrt{\pi n}\,(\frac{n/2}{e})^{n/2}\, e^{1/(6n)}\right)^2} = \sqrt{\frac{2}{\pi n}} \cdot 2^n \cdot e^{\frac{9n+1}{3n(12n+1)}}.$$

Hence,

$$\Pr\{\text{reaching balanced configuration in one move}\} = \frac{\binom{n}{n/2}}{2^n} \geq \sqrt{\frac{2}{\pi n}} \cdot e^{\frac{9n+1}{3n(12n+1)}} > 0.7/\sqrt{n}.$$

Next we upper bound the probability of never reaching a goal state within t steps:

$$(1 - \Pr\{\text{reaching goal state in one step}\})^t \leq (1 - 0.7/\sqrt{n})^t \leq e^{-0.7\,t/\sqrt{n}}$$

using the fact that $(1-x)^t \leq e^{-xt}$ for all $0 \leq x \leq 1$ and $t \geq 0$. Hence the probability of obtaining a balanced configuration within t steps is at least $1 - e^{-0.7\,t/\sqrt{n}}$. Therefore, within $t = \sqrt{n}$ steps, we reach a balanced configuration with constant probability, and within $t = (c/0.7)\sqrt{n}\ln n$ steps, we reach a balanced configuration with probability $1 - 1/n^c$ for any constant $c > 0$. $\square$

4.2 Deterministic Strategy

At first glance, a fast deterministic strategy may not be obvious. Nonetheless, our deterministic strategy is simple: the magician flips the first coin, then the second coin, then the third, and so on (or in any permutation thereof), until reaching a balanced configuration. With the strategy in hand, its analysis is a straightforward continuity argument:

Theorem 4 *Using the deterministic strategy, the magician balances n coins in at most $n-1$ coin flips.*

Proof: Let d_i denote the number of heads minus the number of tails after the ith coin flip in the deterministic strategy. In particular, d_0 is the imbalance of the initial configuration. If we reach n flips, we would have flipped all coins, so $d_n = -d_0$. Thus d_0 and d_n have opposite signs (or are possibly both 0). We also know that $|d_i - d_{i-1}| = 1$. By the discrete intermediate value theorem [Joh98], $d_i = 0$ for some i with $0 \leq i < n$. Thus the magician reaches a balanced configuration after $i \leq n-1$ flips. $\square$

This deterministic strategy is fast, but still takes the square of the time required by the randomized strategy. In contrast, for equalizing all coins, randomization would help by only a constant factor in expectation. Is there a better deterministic strategy for reaching a balanced configuration? The answer is negative, even for strategies that flip multiple coins in a single move, using results from coding theory:

Theorem 5 *Every deterministic strategy for balancing n coins makes at least $n-1$ moves in the worst case.*

Proof: A general deterministic strategy can be viewed as a list of k bit vectors $s_0, s_1, \ldots, s_{k-1}$ where 0s represent coins in their original state and 1s represent coins flipped from their original state. For example, our $(n-1)$-flip strategy is given by the n vectors

$$s_i = (\underbrace{1, 1, \ldots, 1}_{i}\, \underbrace{0, 0, \ldots, 0}_{n-i}), \qquad 0 \leq i < n.$$

For a strategy to balance any initial configuration given by a bit vector x, where 0s represent heads and 1s represent tails, $x \oplus s_i$ must have exactly $n/2$ 1s for some i, where $\oplus$ denotes bitwise XOR (addition modulo 2). In other words, every bit vector x must differ in exactly $n/2$ bits from some s_i. Alon et al. [ABCO88] proved that the optimal such "balancing" set of vectors $s_0, s_1, \ldots, s_{k-1}$ consists of exactly n vectors, and therefore our $(n-1)$-flip strategy is optimal. $\square$

5 Flipping More Coins At Once

Returning to the magic trick of flipping n coins to become all the same, another generalization is to allow the magician the additional flexibility of flipping more than one coin at once. The number of coins flipped per move might be a constant value (as considered in the next section), or might change from move to move (as in Figure 4, left). In either case, we let k denote the number of coins flipped in a move.

In this section, we consider what happens when the spectator gets to choose how many coins the magician must flip in each move. Obviously, if n is even, then the spectator must choose odd values for k, or else the magician could never get out of the odd parity class. But even then the magician is in trouble. We provide a complete answer to when the magician can still succeed:

Lemma 6 *If $n \geq 5$, the magician is doomed.*

Proof: The spectator uses the following strategy: if the distance between the current configuration and the all-heads or all-tails configuration is 1, then the spectator tells the magician to flip three coins. Otherwise, the spectator tells the magician to flip one coin. Because $n \geq 5$, being at distance 1 from one target configuration means being at distance at least 4 from the other target configuration, and $4 > 3$, so the magician can never hit either target configuration. The spectator always says odd numbers, so this strategy satisfies the constraint when n is even. $\square$

Lemma 7 *If $n = 3$ or $n = 4$, the magician can succeed.*

Proof: As mentioned above, flipping k or $n - k$ coins are dual to each other. For $n = 3$ or $n = 4$, the spectator can only ask to flip 1 or $n - 1$ coins. Thus the magician effectively has the same control as when flipping one coin at a time. More precisely, if the spectator says to flip 1 coin, the magician flips the next coin in the $k = 1$ strategy. If the spectator says to flip $n-1$ coins, the magician flips all coins except the next coin in the $k = 1$ strategy. This transformation has effectively the same behavior because the two targets are bitwise negations of each other. $\square$

Despite this relatively negative news, it would be interesting to characterize the sequences of k values for which the magician can win. Such a characterization would provide the magician with additional flexibility and variability for the equalizing trick. In the next section, we make partial progress toward this goal by showing that the magician can succeed for most fixed values of k.

6 Flipping Exactly k Coins at Once

In this section, we characterize when the magician can equalize n coins by flipping exactly k coins in each move. Naturally, we must have $0 < k < n$, because both 0-flip and n-flip moves cannot equalize a not-already-equal configuration. Also, as observed in the previous section, we cannot have both n and k even, because then we could never change an odd-parity configuration into the needed even parity of an all-equal configuration. We show that these basic conditions suffice for the magician:

Theorem 8 *The magic trick with k-flip moves can be performed if and only if $0 < k < n$ and either n or k is odd. The optimal solution sequence uses exactly $2^{n-1} - 1$ moves in the worst case.*

A lower bound of $2^{n-1} - 1$ follows in the same way as Section 2. Again we can view the trick on the n-dimensional hypercube, where 0 represents a bit unchanged from its initial configuration and 1 represents a changed bit. The difference is that now moves (edges) connect two configurations that differ in exactly k bits. The lower bound of $2^{n-1} - 1$ follows because we need to visit every bit string or its complement among 2^n possibilities.

Our construction of a $(2^{n-1} - 1)$-move solution is by induction on n. If k is even, we can consider only odd values of n. The base cases are thus when $n = k + 1$ for both even and odd k. The $n = k + 1$ case has $k = n - 1$, so it is effectively equivalent to $k = 1$ from Section 2. We will, however, need to prove some additional properties about this solution.

It seems difficult to work with general solutions for smaller values of n, so we strengthen our induction hypothesis. Given a solution to a trick, we call a configuration *destined for heads* if the solution transforms that configuration into the all-heads configuration (and never all-tails), and *destined for tails* if it transforms into all-tails (and never all-heads). (Because our solutions are always optimal length, they only ever reach one all-heads configuration or one all-tails configuration, never both, even if run in entirety.) We call a transformation *destiny-preserving* if every configuration on n coins has the same destiny before and after applying the transformation. A transformation is *destiny-inverting* if every configuration on n coins has the opposite destiny before and after applying the transformation. Now the stronger inductive statement is the following:

1. for nk even, flipping the first j coins for even $j < k$ preserves destiny, while flipping the first j coins for odd $j < k$ inverts destiny; and

2. for nk odd, flipping the first j coins for even $j < k$ inverts destiny, while flipping the first j coins for odd $j < k$ preserves destiny, and flipping coins $2, 3, \ldots, k$ preserves destiny.

To get this stronger induction hypothesis started, we begin with the base case:

Lemma 9 *For any $k > 0$, the k-flip trick with $n = k + 1$ coins has a solution sequence of length $2^k - 1$ such that flipping the first j coins for even $j < k$ preserves destiny, while flipping the first j coins for odd $j < k$ inverts destiny.*

Proof: The construction follows the Gray code of Section 2. That flip sequence, ignoring the first coin, can be described recursively by

$$G_k \;=\; G_{k-1}, \text{“flip the } (n-k+1)\text{st coin”}, \; G_{k-1}.$$

To flip $n - 1$ coins in each move, we invert this sequence into

$$\bar{G}_k \;=\; \bar{G}_{k-1}, \text{“flip all but the } (n-k+1)\text{st coin”}, \; \bar{G}_{k-1}.$$

In the base case, $G_0 = \bar{G}_0 = \emptyset$. Validity of the solution follows as in Section 2; indeed, for any starting configuration, the number of moves performed before the configuration becomes all-heads or all-tails is the same in sequences G_k and $\bar{G}_k$.

Every move flips the first coin, so the destiny of a configuration is determined by its parity and the parity of n: if n and the number of coins equal to the first coin (say heads) have the same parity, then the configuration is destined is that value (heads); and if n and the (heads) count have opposite parity, then the destiny is the opposite value (tails). To see why this is true, consider the hypercube viewpoint where 0s represent coins matching the initial configuration and 1s represent flipped coins in an execution of G_k (not $\bar{G}_k$). Then, at all times, the number of 1 bits in the configuration has the same parity as the number of steps made so far. At the same time, every move in $\bar{G}_k$ flips the first coin, so the first coin in the current configuration matches its original value precisely when there have been an even number of steps so far. Thus, when we reach a target configuration of all-heads or all-tails, it will match the original first coin precisely if there have been an even number of steps so far, which is equivalent to there being an even number of 1 bits in the G_k view, which means that the initial and target configurations differ in an even number of bits. In this case, the initial and target configurations have the same parity of coins equal to their respective first coins; but, in the target configuration, all coins match the first coin, so in particular n has the same parity as the number of coins equal to the first coin. We have thus shown this property to be equivalent to the target configuration matching the initial first coin.

It remains to verify the flipping claims. Flipping the first j coins for even $j < k$ preserves the parity of the number of heads as well as the number of tails, but inverts the first coin, so inverts the destiny. Flipping the first j coins for odd $j < k$ changes the parity of the number of heads as well as the number of tails, and inverts the first coin, which together preserve the destiny. □

With this base case in hand, we complete the induction to conclude Theorem 8. In the nonbase case, $n > k + 2$. There are three cases to consider:

Case 1: Both n and k are odd. By induction, we obtain a solution sequence σ' of length $2^{n-2} - 1$ for $n' = n - 1$ satisfying the destiny claims. We view σ' as acting on only the last $n - 1$ of our n coins. Then we construct a solution σ for n as follows:

$$\sigma \ = \ \sigma', \text{ “flip the first } k \text{ coins”}, \ \sigma'.$$

This solution has length $|\sigma| = 2\,|\sigma'| + 1 = 2^{n-1} - 1$.

Next we prove that sequence σ solves the trick. Consider any configuration on n coins, and assume by symmetry that the last $n - 1$ of its coins are destined for heads in σ'. If the first coin is also heads, then the magician arrives at the all-heads configuration within the first σ' prefix of σ. If the first coin is tails, then the σ' prefix will not complete the trick, at which point the magician flips the first k coins. This move has the effect of flipping the first coin to heads as well as flipping the first $k - 1$ of the σ' subproblem, which is destiny-preserving because $k-1$ is even and $(n-1)k$ is even. Therefore, during the second σ', the magician will arrive at the all-heads configuration.

Now we verify the destiny claims. Note that the destiny of a configuration in σ equals the destiny of the last $n - 1$ coins in σ', so we can apply induction almost directly. Flipping the first j coins for even $j < k$ flips the first $j - 1$ of the last $n - 1$ coins, which inverts destiny by induction because $j - 1$ is even and $(n - 1)k$ is even. Similarly, flipping the first j coins for odd $j < k$ preserves destiny by induction. Finally, flipping coins $2, 3, \ldots, k$ flips the first $k - 1$ coins of the last $n - 1$ coins, which preserves destiny by induction because $k - 1$ is even.

Case 2: For n even and k odd, by induction we again obtain a solution sequence σ' of length $2^{n-2} - 1$ for $n' = n - 1$, viewed as acting on only the last $n - 1$ coins. We construct a solution σ for n as follows:

$$\sigma \ = \ \sigma', \text{ “flip coins } 1, 3, 4, \ldots, k + 1\text{”}, \ \sigma'.$$

Again σ has length $2^{n-1} - 1$. Flipping coins $1, 3, 4, \ldots, k+1$ has the effect of flipping the first $2, 3, \ldots, k$ of the last $n - 1$ coins, which by induction is destiny-preserving because $(n - 1)k$ is odd. Thus, if the destiny of σ' does not match the first coin, then it will match the newly flipped first coin during the second σ'. As before, destiny in σ matches destiny in σ' on the last $n - 1$ coins. Flipping the first j coins for even $j < k$ flips the first $j - 1$ of the last $n - 1$ coins, which preserves destiny by induction because $j - 1$ is odd and $(n - 1)k$ is odd. Similarly, flipping the first j coins for odd $j < k$ inverts destiny by induction.

Case 3: For n odd and k even, by induction we obtain a solution sequence σ' of length $2^{n-3} - 1$ for $n' = n - 2$, which we view as acting on only the last $n - 2$ coins. Then we construct a solution σ for n as follows:

$$\sigma \ = \ \sigma', \text{ “flip the first } k \text{ coins”}, \ \sigma', \text{ “flip coins } 1, 3, 4, \ldots, k + 1\text{”}, \ \sigma', \text{ “flip the first } k \text{ coins”}, \ \sigma'.$$

This solution has length $|\sigma| = 4\,|\sigma'| + 3 = 2^{n-1} - 1$. Restricting attention to the first two coins, σ first flips both coins, then flips the first coin only, then flips both coins again. Together these enumerate all possibilities for the first two coins. Restricting to the last $n - 2$ coins, these moves correspond to flipping the first $k - 1$ coins, coins $2, 3, \ldots, k$, and again the first $k - 1$ coins. By induction, all three of these operations preserve destiny because $k - 1$ and $(n - 2)k$ are odd. Therefore all three executions of σ' produce the same target configuration (all-heads or all-tails) which will eventually match one of the combinations of the first two coins. Flipping the first j coins for even $j < k$ flips the first $j - 2$ of the last $n - 2$ coins, which preserves destiny by induction because $j - 2$ is even and $(n - 1)k$ is odd. Similarly, flipping the first j coins for odd $j < k$ inverts destiny by induction.

This concludes the inductive proof of Theorem 8.

7 Permuting Coins Between Flips

Our final variation of the coin-flipping magic trick is parameterized by a group G of permutations on $\{1, 2, \ldots, n\}$. We start with n coins, labeled 1, 2, …, n, in initial orientations decided by the spectator. At each step, the blindfolded magician can choose an arbitrary collection of coins to flip. Prior to flipping the coins, the spectator chooses an arbitrary permutation from the permutation group G, and re-arranges the coins according to that permutation. The spectator then flips the coins at the locations specified by the magician. The magician then asks "are the coins all the same?" and the trick ends if the answer is positive.

Whether the magician has a winning strategy depends on the permutation group G. In this section, we will characterize exactly which permutation groups allow the magician to perform such a trick. Our characterization also applies to the superficially easier version of the trick where the spectator flips the coins specified by the magician before permuting the coins, because we consider deterministic strategies.

Our characterization of valid groups turns out to match the existing notion of "2-groups". A group G is a 2-*group* if the number $|G|$ of its group elements is an exact power of 2. The simplest example of such a group is the cyclic group C_{2^k} of order 2^k, that is, a rotating table of coins with 2^k coins. Another simple example is the dihedral group D_{2^k} of symmetries of a regular 2^k-gon (acting as a permutation group on the vertices), that is, allowing the spectator to confuse "left" (counterclockwise on the table) from "right" (clockwise on the table). A more sophisticated example is the iterated wreath product of k copies of the group S_2 of permutations on two elements. This group can be viewed as permutations on the 2^k leaves of a perfect binary tree, generated by reversal of the leaves beneath any internal node in the tree. Of course, we can also obtain a 2-group by taking direct products of 2-groups.

Theorem 10 *The magician can successfully perform the n-coin trick with permutation group G if and only if G is a 2-group. In this case, the worst-case optimal solution sequence makes exactly $2^{n-1} - 1$ moves.*

To prove this theorem, it is convenient to speak in the language of group representations. For a group G and a field $\mathbb{F}$, an *n-dimensional $\mathbb{F}$-representation of G* is an n-dimensional $\mathbb{F}$-vector space V together with a left action of G on V such that $g(v + \lambda w) = gv + \lambda gw$ for all $g \in G$ and $v, w \in V$ and $\lambda \in \mathbb{F}$. (A *left action* is a function $G \times V \to V$ such that $(gh)v = g(hv)$ for all $g, h \in G$ and $v \in V$.)

In the context of our magic trick, we have a permutation group G on the coins; call the coins $1, 2, \ldots, n$ for simplicity. The vector space $V = (\mathbb{F}_2)^n$ represents all possible configurations of the n coins. We consider the $\mathbb{F}_2$-representation of G on V defined by $g(v_1, v_2, \ldots, v_n) = (v_{g(1)}, v_{g(2)}, \ldots, v_{g(n)})$. In other words, a group action g simply permutes the coins. In this algebraic language, we can view one move in the magic trick as follows. Suppose the current configuration of coins is $v = (v_1, v_2, \ldots, v_n) \in V$, where v_i is 0 if the ith coin is heads and 1 if it is tails. The blindfolded magician specifies a vector $w = (w_1, w_2, \ldots, w_n) \in V$, where w_i is 1 if the magician specifies to flip the ith coin and 0 otherwise. The spectator then picks a permutation $g \in G$, applies that permutation to v, and applies the flips specified by w to $g(v)$. Hence the resulting configuration is $g(v) + w$. If $g(v) + w = (0, 0, \ldots, 0) = \vec{0} \in V$ (all heads) or $g(v) + w = (1, 1, \ldots, 1) = \vec{1} \in V$ (all tails), then the magician has succeeded.

Our proof of Theorem 10 consists of three lemmas. The first lemma shows that, if G is not a 2-group, then the magician cannot guarantee a successful performance of the trick. Next we define the notion of a "G-invariant flag". The second lemma shows that the existence of G-invariant flag on V implies a winning strategy for the magician. The third lemma establishes that V has a G-invariant flag if G is a 2-group. Together, these three lemmas prove the theorem.

Lemma 11 *If G is not a 2-group, then the magician is doomed.*

Proof: Suppose G is not a 2-group, i.e., $|G|$ is not a power of 2. Thus there is an odd prime p that divides $|G|$. By Cauchy's group theorem, there is a permutation $g \in G$ of order p, i.e., for which g^p is the smallest power of g that equals the identity permutation. The order of a permutation is the least common multiple of its cycle lengths in its disjoint-cycle decomposition, and p is prime, so there must in fact be a cycle of length p, i.e., a coin $i \in \{1, 2, \ldots, n\}$ such that $i, g(i), g^2(i), \ldots, g^{p-1}(i)$ are all distinct, while $g^p(i) = i$. We can assume by renaming some of the coins that this cycle appears among the first p coins: $i = 1$, $g(i) = 2$, $g^2(i) = 3, \ldots, g^{p-1}(i) = p$.

We define the set X of "trouble configurations" to consist of configurations in which the first three coins are not all equal, i.e., $X = \{x \in V : (x_1, x_2, x_3) \notin \{(0, 0, 0), (1, 1, 1)\}\}$. The spectator chooses a configuration in X as the initial configuration. We next give a strategy for the spectator that guarantees staying within X, no matter how the magician moves. This strategy then foils the magician, because not all the coins can be equal if the first three coins are never equal.

Consider any trouble configuration $x \in X$ and magician move $w \in V$. We need to show that the spectator has a move $h \in G$ resulting in configuration $h(x) + w \in X$. Look at the magician moves for the first three coins: w_1, w_2, w_3. There are eight possibilities for these three bits. We can factor out a symmetry by letting $a \in \{0, 1\}$ be arbitrary and letting $b = 1 - a$. Then the three bits have four possible patterns: aaa, aab, aba, and abb. The aaa pattern flips none or all of the first three coins, which means they remain not all equal, and thus the configuration remains in X if the spectator chooses the identity permutation (i.e., does not permute the coins). Three patterns remain: aab, aba, and abb.

The cyclic sequence $x_1, x_2, \ldots, x_p$ of p bits forming the p-cycle in g consists of an odd number of bits. Because $x \in X$, at least one of these bits is 0 and at least one is 1, Thus both the patterns ccd and eff must occur in the cyclic sequence, where $d = 1 - c$ and $f = 1 - e$. Now, if w_1, w_2, w_3 has pattern aba or abb, we use the ccd pattern; and if w_1, w_2, w_3 has pattern aab, we use the eff pattern. In either case, say the latter pattern appears as $(x_{k+1}, x_{k+2}, x_{k+3})$, where $k \in \{0, 1, \ldots, p - 1\}$. The spectator then chooses $h = g^k$, so that $h(x)$ puts the pattern in positions $1, 2, 3$. Thus $h(x) + w$ sums the two patterns, resulting in $aba + ccd = ghh$, $abb + ccd = ghg$, or $aab + eff = iji$. In all cases, $h(x) + w \in X$. $\square$

The next two lemmas use the notion of "G-invariant flag". A subspace $W \subseteq V$ is *G-invariant* if $gv \in W$ for all $v \in W$ and all $g \in G$. A *flag* on V is a chain of subspaces $\{\vec{0}\} = W_0 \subset W_1 \subset \ldots \subset W_{n-1} \subset W_n = V$ where $\dim(W_i) = i$ for $i = 0, 1, \ldots, n$. A flag $W_0 \subset W_1 \subset \ldots \subset W_{n-1} \subset W_n$ is *G-invariant* if each W_i is G-invariant.

Next we describe the known connection between G-invariant flags and 2-groups:

Lemma 12 [Miy71] *If G is a 2-group and W is any $\mathbb{F}_2$-representation of G, then there a G-invariant flag on W.*

Finally we show the connection between G-invariant flags and the magic trick. For simplicity, we show here how to perform a more specific version of the trick: make the coins all heads. This version requires $2^n - 1$ moves. A slight modification allows the all-tails case and reduces the number of moves to $2^{n-1} - 1$. The characterization of valid groups G remains unaffected. These move bounds are optimal because in particular we are solving the regular n-coin game (with the trivial group $G = \{\vec{0}\}$) from Section 2.

Lemma 13 *If V has a G-invariant flag, then the magician can make all coins heads in $2^n - 1$ moves.*

Proof: Suppose $W_0 \subset W_1 \subset \ldots \subset W_{n-1} \subset W_n$ is a G-invariant flag on V. Choose any element $w^{(i)} \in W_i \setminus W_{i-1}$ for each $i = 1, 2, \ldots, n$. Define the move sequences $\sigma_1, \sigma_2, \ldots, \sigma_n$ recursively by $\sigma_0 = \emptyset$ and $\sigma_i = \sigma_{i-1}, w^{(i)}, \sigma_{i-1}$ for $i = 1, 2, \ldots, n$. By a simple induction, σ_i consists of $2^i - 1$ moves. The magician's strategy is σ_n with $2^n - 1$ moves.

We prove by induction on i that σ_i brings any initial configuration $v \in W_i$ to the all-heads configuration $\vec{0}$. Then, in particular, σ_n brings any $v \in W_n = V$ to $\vec{0}$. In the base case, $i = 0$ and $v \in W_0 = \{\vec{0}\}$, so the magician has already won. In the induction step, $i > 0$, and there are two cases: $v \in W_{i-1}$ and $v \in W_i \setminus W_{i-1}$. If $v \in W_{i-1}$, then by induction the prefix σ_{i-1} of σ_i brings v to $\vec{0}$. Otherwise, we analyze the three parts of σ_i separately. In the prefix σ_{i-1} of σ_i, we transform configuration v' into $g(v') + w^{(j)}$ where $1 \leq j < i$. Because W_i is G-invariant, $v' \in W_i$ implies $g(v') \in W_i$. Because W_{i-1} is G-invariant, $v' \notin W_{i-1}$ implies $g(v') \notin W_{i-1}$. Because $w^{(j)} \in W_{i-1}$ for $j < i$, $v' \in W_i \setminus W_{i-1}$ implies $g(v') + w^{(j)} \in W_i \setminus W_{i-1}$. (In contrapositive, $g(v') + w^{(j)} \in W_{i-1}$ implies $(g(v') + w^{(j)}) - w^{(j)} = g(v') \in W_{i-1}$ and by G-invariance $v' \in W_{i-1}$.) Therefore the configuration v' remains in $W_i \setminus W_{i-1}$ throughout the prefix σ_{i-1} of σ_i. Next σ_i takes the resulting configuration v'' and applies $w^{(i)} \in W_i \setminus W_{i-1}$, so the resulting configuration $v'' + w^{(i)}$ drops to W_{i-1}. (A simple counting argument shows that $v'' = x - w^{(i)}$ for some $x \in W_{i-1}$, and hence $v'' + w^{(i)} = x \in W_{i-1}$.) Finally, by induction, the second copy of σ_{i-1} brings the configuration to $\vec{0}$. $\square$

Acknowledgments. We thank Patricia Cahn, Joseph O'Rourke, and Gail Parsloe for helpful initial discussions about these problems. We thank the participants of *Gathering for Gardner 8* for pointing us to [LR81]; and Noga Alon, Simon Litsyn, and Madhu Sudan for pointing us to [ABCO88].

References

[ABCO88] N. Alon, E. E. Bergmann, D. Coppersmith, and A. M. Odlyzko. Balancing sets of vectors. *IEEE Transactions on Information Theory*, 34(1), January 1988.

[ES95] Richard Ehrenborg and Chris M. Skinner. The blind bartender's problem. *Journal of Combinatorial Theory, Series A*, 70(2):249–266, May 1995.

[Etz97] Tuvi Etzion. The depth distribution—a new characterization for linear codes. *IEEE Transactions on Information Theory*, 43(4):1361–1363, July 1997.

[Ful80] Karl Fulves. *The Children's Magic Kit: 16 Easy-to-do Tricks Complete with Cardboard Punchouts*. Dover Publications, Inc., New York, 1980.

[Gar91] Martin Gardner. The rotating table and other problems. In *Fractal Music, Hypercards and More...: Mathematical Recreations from Scientific American Magazine*. W. H. Freeman & Co., October 1991. Based on "Mathematical Games: About rectangling rectangles, parodying Poe and many another pleasing problem", *Scientific American*, 240(2):16–24, February 1979, and the answers in "Mathematical Games: On altering the past, delaying the future and other ways of tampering with time", *Scientific American*, 240(3):21–30, March 1979.

[Gra53] F. Gray. Pulse code communication. U.S. Patent 2,632,058, March 1953. http://patft.uspto.gov/netacgi/nph-Parser?patentnumber=2632058.

[Gua98] Dah-Jyh Guan. Generalized gray codes with applications. *Proc. Natl. Sci. Counc.*, 22:841–848, 1998.

[Joh98] Richard Johnsonbaugh. A discrete intermediate value theorem. *The College Mathematics Journal*, 29(1):42, January 1998.

[LFW00] Yuan Luo, Fang-Wei Fu, and Victor K.-W. Wei. On the depth distribution of linear codes. *IEEE Transactions on Information Theory*, 46(6):2197–2203, September 2000.

[LR81] William T. Laaser and Lyle Ramshaw. Probing the rotating table. In D. A. Klarner, editor, *The Mathematical Gardner*. Wadsworth, Belmont, California, 1981. Republished in 1998 by Dover in *Mathematical Recreations: A Collection in Honor of Martin Gardner*, pages 285–307.

[LW80] Ted Lewis and Stephen Willard. The rotating table. *Mathematics Magazine*, 53(3):174–179, May 1980.

[Miy71] Takehito Miyata. Invariants of certain groups I. *Nagoya Mathematical Journal*, 41:69–73, 1971.

[Roo02] Eric Roode. Coin puzzle. Posting to Philadelphia Perl Mongers mailing list, May 7 2002. http://lists.netisland.net/archives/phlpm/phlpm-2002/msg00137.html.

[YEM93] Reuven Bar Yehuda, Tuvi Etzion, and Shlomo Moran. Rotating-table games and derivatives of words. *Theoretical Computer Science*, 108(2):311–329, February 1993.