

The Twelve-Coins Puzzle

John Denker

Contents

1	Statement of the Puzzle	1
2	Outline of a Systematic Strategy	2
3	Re-Statement of the Puzzle, with Variations	3
4	Information-Theory Analysis	4
5	Communication-Theory Analysis	6
6	Geometric Interpretation	9
7	For Further Reading	9

Introductory Discussion

1 Statement of the Puzzle

Here is a rough statement of the puzzle:

- You are given twelve coins. All appear to be identical, as far as you can tell by eye, but you are told that one of them is a counterfeit. All genuine coins have the same mass, while a counterfeit is either lighter or heavier than a genuine coin.
- You are provided an equal-arm balance (sometimes called a scale or scales), as shown in [figure 1](#). The balance provides one of three possible indications: the right pan is heavier, or the pans are in balance, or the left pan is heavier. The balance has sensitivity sufficient for the task.
- One objective is to identify the the counterfeit coin, and to ascertain whether it is heavy or light.
- Another objective is to carry out the task using as few weighings as possible.
- More specifically, you need to decide in advance the maximum number of weighings you want, and you must not exceed that. Therefore a strategy that sometimes gets by with a small number of weighings but sometimes requires a large number of weighings is not a good strategy. In other words, we want a *minimax* strategy, i.e. one that does well in the worst case.




Figure 1: Two-Pan Equal-Arm Balance

2 Outline of a Systematic Strategy

It pays to analyze the problem in a systematic way. Here are some hints to help you get started.

1. One of the cardinal rules for scientific thinking is to consider all of the plausible hypotheses. In other words, consider all the plausible scenarios, all the plausible outcomes, all the plausible answers to the question. So the first step is to figure out what are the possibilities.
2. Simple questions you can do in your head, but this question is so intricate that you are likely to get tangled up if you try that. On the other hand, it is simple enough that you can just write down¹ all of the possible outcomes. So, another bit of advice: *Paper is cheap*. So write down the hypotheses. Do it now.
3. When we make a measurement of the coins, it allows us to cross some hypotheses off the list. Some hypotheses are no longer viable, i.e. no longer consistent with the available data.
4. We can measure progress by keeping count of how many hypotheses have been crossed off.
5. Let's do an example: Suppose you weigh one coin against another, with ten bystanders. The measurement has three possible results: Left-heavy, right-heavy, or neutral. In each case, figure out which hypotheses we get to cross off the list.
6. Let's do another example: Suppose you weigh six coins in one pan against six coins in the other pan. What are the possible measurement results? In each case, which hypotheses do we get to cross off the list?
7. Now think about the problem some more. Come up with a strategy for the first weighing that makes the best possible progress – best in a minimax sense – in terms of crossing out hypotheses.
8. After the first weighing, you are left with a new problem, which is basically a smaller, simpler version of the original problem. Figure it out in the same way.

Remark: Note the following contrast:

¹Later we will learn how to handle puzzles with thousands or billions of possible outcomes, for which a more sophisticated representation is necessary, but for this problem, the direct explicit representation works fine.

You may have heard it said that when designing experiments, you should never change more than one variable at a time. As discussed in [reference 1](#), that's reasonable advice in situations where experiments are cheap and thinking about the results is the hard part.

On the other hand, in the real world there are lots of situations where the experiments are very expensive, in which case it pays to minimize the number of experiments and maximize the information you get from each experiment. In this case, you almost always wind up changing more than one variable at a time. You have to be clever about how you analyze the results of each experiment. Sometimes cleverness pays off.

More Advanced Analysis

Let's start over. Let's analyze the same problem, plus several variants of the problem, using some more sophisticated techniques.

3 Re-Statement of the Puzzle, with Variations

Version 1a: You are given twelve coins. All appear to be identical, as far as you can tell by eye, but you are told that one of them is a counterfeit. All genuine coins have the same mass, while a counterfeit is either lighter or heavier than a genuine coin.

You are allowed to weigh coins using a two-pan equal-arm balance (sometimes called a scale or scales), as shown in [figure 1](#). The balance provides one of three possible indications: the right pan is heavier, or the pans are in balance, or the left pan is heavier. The balance has sensitivity sufficient for the task. Note that the result is qualitative not quantitative: there is no indication of how much heavier the heavy pan is.

Your mission, should you decide to accept it, is to identify the counterfeit coin and tell whether it is "light" or "heavy," using *at most three weighings*.

You are not allowed to tamper with or scrutinize the coins, nor gather any information about them except by the three weighings. You may, if you wish, label the coins if that helps you keep track of which is which.

Version 2a: Same as version 1a, but you are told that *at most* one of the coins is a counterfeit, i.e. possibly all 12 are genuine.

Version 1b: Same as version 1a, with the added restriction that we want a *parallel* algorithm. That is, you must decide *in advance* (before any weighings are carried out) which coins are to be weighed in each of the three weighings. Most people are tempted to use the results of the first weighing to help design the subsequent weighings, but that is disallowed in this version.

Having a parallel algorithm means, among other things, that the three weighings can be performed in any order.

Version 2b: Same as version 2a, with a parallel algorithm.

Extended versions: There are four more versions x1a through x2b, the same as the above, except that there are *thirteen* unknown coins, and you are allowed to supply one known-good coin of your own and use it in the weighings however you like.

Beginners are encouraged to start with version 1a, which is the easiest. In this note we will solve version 2a and immediately extend it to version x2b, which is the hardest version. All the other versions are then seen as corollaries.

4 Information-Theory Analysis

This is an example of what scientists call a “Design of Experiment” (DoE) question. Entire books have been written about DoE.

Let’s start by reconnoitering the situation using the methods of information theory. For a nice, accessible introduction to information theory, see [reference 2](#).

The first step is to calculate the entropy,² i.e. how much we don’t know about the problem. In version 1, there are 24 possible outcomes:³ perhaps coin #1 is heavy, perhaps coin #2 is heavy, and so forth for twelve possibilities, plus another twelve possibilities if the counterfeit coin is light. Invoking the minimax⁴ principle, we assume that these 24 outcomes are equally likely. That gives us a total entropy of

$$\begin{aligned} S &= \log_2(24) \text{ bits} = 4.585 \text{ bits} \\ S &= \log_3(24) \text{ trits} = 2.89 \text{ trits} \end{aligned} \tag{1}$$

We get to compare this with the information we hope to get. The balance provides only three possible results per weighing. If we are lucky, we might hope that the results are equally likely, or nearly so, in which case each weighing gives us one trit of information. If the results are not equally likely, the information will be less. Therefore we have the following upper bound on the amount of information available:

$$\begin{aligned} I &\leq \log_2(27) \text{ bits} = 4.755 \text{ bits} \\ I &\leq \log_3(27) \text{ trits} = 3 \text{ trits} \end{aligned} \tag{2}$$

The upper bound is high enough, but there is not much to spare. We need an approach that is very nearly optimal.

As an example of bad design, consider for a moment what happens if the first weighing involves three coins in the left pan, three coins in the right pan, and six bystanders. There is a distinct possibility that the pans will balance. That tells us that the counterfeit is among the six bystanders. So we are left with a reduced version of the original problem: six coins with two weighings remaining. Information theory tells us that this reduced problem is, alas, provably impossible. There are twelve possible outcomes (which coin, light or heavy) but there are only nine possible results of the remaining weighings. That is, we have 2.26 trits of entropy remaining but at most 2 trits of information available. This is a hopeless situation.

As an example of much better design, consider what happens if the first weighing involves four coins in each pan plus four bystanders. This has the nice property that the three possible outcomes (left heavy, balance, right heavy) are equiprobable, so the first weighing returns the maximum possible information, i.e. one full trit.

Suppose the pans balance. We are left with a reduced version of the problem: four coins with two weighings remaining. This problem is not hopeless: we have eight possible outcomes (1.89 trits of entropy) and we have a reasonable hope of getting sufficient information from the remaining weighings.

Now suppose the pans don’t balance. This leaves us with eight coins and two weighings remaining. However, there are only eight possible outcomes (not sixteen), since once we identify which of the eight coins is the counterfeit, we know (based on the results of the first measurement) whether it is heavy or light. So once again we have eight possible outcomes (1.89 trits of entropy) with two weighings remaining, so there is hope.

Additional examples of this sort of hypothesis testing can be found in [reference 3](#), [reference 4](#), and [reference 5](#).

²For a general discussion of what entropy is, see [reference 3](#).

³The analysis of version 2 is left as an exercise for the reader.

⁴By way of analogy: when you are playing chess, you should plan your moves on the assumption that your opponents will make their best moves; you shouldn’t assume they are going to make things easy on you. In this case, if the possibilities were not equiprobable, the entropy would be less and the puzzle would be easier.

5 Communication-Theory Analysis

Let's treat the balance as a sort of communication channel. That is, suppose it is trying to transmit to you a code telling you the answer to the puzzle. The channel has an alphabet with only three symbols (L, B, and R) denoting respectively left-pan heavy, balance, and right-pan heavy. The channel can give you only a single codeword consisting of three such symbols. So the overall task reduces to this: using some subset of the 27 possible codewords, we want to assign a codeword to each coin. The codeword must be unique to that coin, and the set of codewords must obey certain constraints as discussed below.

First, let's just look at all 27 codewords. It is easy to construct this table; it's just counting in base 3.

RRR	13	
RRB	12	
RRL	11	
RBR	10	
RBB	9	
RBL	8	
RLR	7	
RLB	6	
RLL	5	
BRR	4	
BRB	3	
BRL	2	
BBR	1	
BBB	0	(3)
BBL	-1	
BLR	-2	
BLB	-3	
BLL	-4	
LRR	-5	
LRB	-6	
LRL	-6	
LBR	-8	
LBB	-9	
LBL	-10	
LLR	-11	
LLB	-12	
LLL	-13	

Now, suppose you pick one of the codewords such as RRB and paint it on one of the coins. You can think of the codeword in two ways:

- The codeword tells how the coin participates in the weighings: the RRB coin is in the right pan in the first weighing, is again in the right pan in the second weighing, and is a bystander in the third weighing.
- The codeword tells you how to interpret what the balance is telling you: if the first weighing is right-pan heavy, the second weighing is again right-pan heavy, and the third weighing is balanced, then the coin marked RRB is counterfeit and is heavier than a genuine coin.

At this point the alert reader will be wondering what happens if the coin we marked RRB is *lighter* than a genuine coin. Then the balance will report LLB, that is, left-pan heavy, left-pan heavy, neutral balance. From this we conclude that if we have marked one of the coins RRB, we must not mark any of the other coins LLB, lest there be ambiguity. More generally, every time we use a codeword from the list in [equation 3](#), we must “burn” its inverse code. (The inverse is formed by replacing L with R and vice versa.)

Consequently, even though it appears we have 27 codewords, this approach has no chance of handling more than 13 coins. Since we are only required to handle 12 coins, things are still on track.

There is one other constraint: The balance is only capable handling an equal number of coins in the two pans. So we must choose a set of codewords (a subset of [equation 3](#)) with the property that in each of the three columns the number of Rs equals the number of Ls.

The design process proceeds as follows: We immediately discard the BBB code as being unsuitable for marking on any coin ... since a coin that participates in none of the weighings can never be properly characterized.⁵ Now take the top half of [equation 3](#), i.e. codes 1 through 13 as a tentative set of codewords. We get to discard one, because we only need 12. The obvious choice is to discard the RRR code, because by doing so we create an *even* number of Rs and Ls in the remaining set. Finally, we replace four of the codes by their inverses. (Remember, we get to use any given code *or* its inverse.) We need to do some inverting because the tentative set started out too rich in Rs in the first column, and we need to turn four of them into Ls. To make a long story short, it suffices to invert 12, 11, 9, and 7. The result is:

LLB	-12	
LLR	-11	
RBR	10	
LBB	-9	
RBL	8	
LRL	-7	
RLB	6	(4)
RLL	5	
BRR	4	
BRB	3	
BRL	2	
BBR	1	

If you scrutinize [equation 4](#), you will discover the following nice properties:

- There are twelve rows of three columns.
- Each of the columns has exactly four Rs, four Ls, and four Bs.
- For each codeword in the list, its inverse does not appear.

The codewords in [equation 4](#) solve version 1 of the problem directly. They also solve version 2, if we adjoin the observation that if the balance generates the BBB codeword it means no counterfeit is present.

To solve the extended versions of the puzzle, adjoin the RRR codeword to [equation 4](#) and put the known-good coin in the left pan.

6 Geometric Interpretation

In the language of Design of Experiment, these twelve codewords are the *test vectors*.

If you are good at visualizing things in higher dimensions (most people aren't) it may be amusing to visualize these test vectors as shown in [figure 2](#).

If you scrutinize the figure, you can discover quite a number of interesting properties.

- For each vector that is present, the inverse (inverted through the midpoint of the cube) is absent. This property is necessary in order to prevent ambiguity, as discussed in [section 5](#). You can add one (and only one) more vector while preserving this property.

⁵The BBB code will resurface below, in connection with version 2 of the puzzle, in the case where none of the coins is counterfeit.




Figure 2: Test Vectors for the Twelve-Coins Puzzle

- There are nine planes through the cube (three perpendicular to each axis), and each plane contains exactly four vectors. Nine planes is quite a bit more structure than we noticed from the symbolic representation of the codewords in equation 4.

7 For Further Reading

1. John Denker,
“Changing One Variable at a Time ... or Not”
www.av8n.com/physics/one-variable-at-a-time.htm
2. John R. Pierce,
Symbols, Signals, and Noise
3. For a general discussion of what entropy is, see:
www.av8n.com/physics/thermo/entropy.html
4. John Denker,
“Twenty Questions”
www.av8n.com/physics/twenty-questions.htm
5. John Denker,
“Learning, Remembering, and Thinking”
www.av8n.com/physics/thinking.htm