2 TWENTY BOXES AND PERMUTATIONS

2.1 THE GAME

Number of players:	10–20
You will need:	1 box for each player; each player has an ID card and a pen (if ID cards are unavailable, each player gets 1 piece of paper)

Let us say we have 20 players. The 20 players stand in a line in alphabetical order. In front of each player is a box with her full name written on it (that is what the pen is for). For simplicity, let us assume there are no players with identical names. The players put their identity cards into their own boxes and are

then asked to leave the room. The audience is now allowed to arbitrarily shuffle the ID cards around between the boxes, making sure there is still an ID card inside each box. Then the players are asked back into the room one at a time, where they find the 20 boxes still lined up on the table, as before. The first player called back in has to try to find her personal ID card but is only allowed to check inside half of the boxes, that is, she can look into $20/2 = 10$ of the boxes to find her ID. She is only allowed to look into the box and is not allowed to move any of the ID cards. After the player has checked 10 boxes, she leaves the room, and the next player enters. Every player is only allowed to look into 10 boxes. The players collectively win the game if and only if every one of them finds her ID card. That is, if one or more players do not find their ID cards, all of them fail. After the game has started, the players are neither allowed to communicate with one another nor to interfere with the way the game is set up. Thus, each player will find the boxes and their contents set up in exactly the same way; the players must not move around either the IDs or the boxes.

What is a good search strategy for the players to agree on? How often will they be able to win the game?

Just to make things simple, let us first assume that the audience rearranges the IDs randomly into the boxes. We will discuss later what happens if the audience tries to fool the players.

Some strategies that the players could use are not very clever. Suppose, for example, that every player looks into the first 10 boxes. Working with this strategy, they will always lose the game, since, regardless of the arrangement of IDs, 10 players will not find their IDs.

If, alternatively, every player looks into 10 (independently and uniformly) randomly chosen boxes, then each player recovers their ID with probability 1/2. The chance that all players

recover their ID is the product of the individual chances, as the success or failure of a player is independent of the success or failure of any other player with this strategy. For the whole group, the probability of winning is then $p_{\text{win,lower}} = 1/2 \cdot 1/2 \cdot \cdots \cdot 1/2 = 1/2^{20}$. This number is rather tiny: It is comparable to the probability that 2 people sitting silently in front of each other for up to 14 days will choose the same second to shout "Boo!" if both choose a second independently and uniformly within these 14 days. We call $p_{\text{win,lower}}$ a lower bound, since an optimal strategy will have at least the same chance of winning (the same chance only if this strategy were indeed an optimal one, which it is not).

The players could also divide the boxes beforehand and say that players 1–10 check boxes 1–10 and players 11–20 check boxes 11–20. This strategy is obviously better than all players using the same set of 10 boxes, as the players now have a nonzero chance of winning. It also improves on the purely random strategy mentioned above.[1] In fact, it is about 5.6 times better than with the purely random strategy. This is still an almost negligible chance of winning (approximately $3/2^{19}$), but why did it increase in relative terms? The reason is that the division of boxes induces a dependence between the success of individual players. If player 1 is successful (that is, his ID is in one of boxes 1–10, then the chance of, say, the eleventh player to succeed has also increased. Before we know that the first player is successful, the probability that the eleventh player finds his ID under boxes 11–20 is one half. After knowing that the first player has succeeded, the chance of player eleven to succeed is now $10/19 > 1/2$, as there are just

1. One can prove that the chance of winning now equals 1 divided by $\binom{n}{n/2}$, where $n = 20$. According to the so-called Stirling approximation this is roughly $\sqrt{\frac{\pi}{2}n}/(2^n)$ and therefore only slightly larger than $1/(2^n)$, the probability of winning with the purely random strategy.

19 boxes left for his ID instead of the original 20 and each of the 19 boxes left is equally likely to contain his ID. This results in just a very modest boost of the overall success probability, but it is a step in the right direction.

The key to a good strategy is to induce a much stronger dependence between the outcomes, so that if the first player is successful, it will make success of the other players very likely or almost certain (and vice versa, if the first player fails, the others will also very likely fail). The reader should pause here and think once more about the questions presented on page 18.

2.2 HOW WELL CAN A STRATEGY WORK?

Let us now derive an upper bound, so that we can see how often the players are likely to fail, even if they adopt the best-possible strategy. Let us consider the first player to enter the room. She looks into 10 boxes. Her choice of boxes may depend on the agreed-on strategy, of course, but the probability of success does not. No matter which boxes the first person looks into, she will find her ID with probability 1/2. And if the first person fails, the game is certainly lost. (But even if she is lucky, the game is still far from being won.) We therefore have[2]

$$(1/2)^{20} = p_{\text{win,lower}} \leq p_{\text{win,optimal}} \leq p_{\text{win,upper}} = 1/2.$$

Surprisingly, there is a strategy that yields a winning probability that is much larger than the lower bound and that comes close to the upper theoretical limit of the upper bound, so that the group can win almost half of all games.

2. Here, $p_{\text{win,lower}}$ is still based on the purely random strategy. We could, in principle, also use the slightly improved version described above, but then $p_{\text{win,lower}}$ would have a more difficult form.

2.3 SOLUTION

The optimal strategy can be described easily. The player with name j first looks into box j (i.e., the box with her own name on it). If she finds her own name in there, she may stop. If she finds another name k (or rather, the ID of person k), say, she then looks into box k. In that box, she may find another name, ℓ, and continues with that box, and so on. The probability of winning with 20 players is about 1/3 and is thus very close to the bound of one-half that we derived in section 2.2.

At first glance, it might not be clear at all why this procedure is successful. The game allows for a natural formulation using the language of permutations. Next we introduce some basic properties of permutations, which will allow us to provide a simple yet thorough analysis of the strategy.

2.4 SOME MATHEMATICS: PERMUTATIONS AND CYCLES

The solution is best understood when we identify the names with the numbers $1, \ldots, 20$. It does not matter which name corresponds to which number, but for simplicity, let us take alphabetical order, in which (by assumption) the boxes are already ordered. So, let us think about the players having numbers between 1 and 20 as names[3] and we will use the terms "names" and "numbers" interchangeably.

Permutations

A *permutation* π of n numbers is a bijective map

$$\pi : \{1, \ldots, n\} \to \{1, \ldots, n\}.$$

Such a map assigns to each number $1, \ldots, n$ one unique number that is again between 1 and n. The term "bijective" means that

3. In fact, this is the setup that is often used when the game is introduced.

each number is "hit" by π exactly once. It is maybe easiest to explain this within the game. Each distribution of IDs into boxes is a permutation: every box j contains an ID (or number) $\pi(j)$, and each ID appears exactly once. That is, for each ID k, there is a unique box j with ID $\pi(j) = k$ in it. Often, it is convenient to write permutations in the following form:

$$\begin{array}{cc|ccccc} j & \text{(box number)} & 1 & 2 & 3 & 4 & 5 \\ \hline \pi(j) & \text{(ID)} & 4 & 5 & 1 & 3 & 2 \end{array}, \tag{2.1}$$

instead of writing $\pi(1) = 4$, $\pi(2) = 5$, $\pi(3) = 1$, $\pi(4) = 3$, and $\pi(5) = 2$.

The set of permutations of n numbers is often called S_n and referred to as the *symmetric group*. How many permutations are there? Let us consider the boxes again. For the first box, we can choose from 20 ID cards, for the second box there are 19 possibilities left, and so on. In total, there are thus $20 \cdot 19 \cdot \cdots \cdot 2 \cdot 1 = 20!$ different ways of distributing the IDs into the boxes—which is a rather large number:

$$20! = 2432902008176640000.$$

More generally,

$$\#S_n = n!, \text{ where } \#S_n \text{ denotes the number of elements in } S_n.$$

Cycles

There is one more concept that will prove to be helpful when analyzing our strategy. Each permutation can be decomposed into smaller parts, called *cycles*. Suppose that you manually rearrange the ID cards in the boxes. How would you do it? You could take ID card 1 and put it into box 3. You would then take the ID card from that box (ID card 3) and put it into box 4. ID card 4 you would put into the empty box 1. (You would then rearrange the remaining numbers, 2 and 5, in an arbitrary way). You would have thus created a permutation with a cycle that contains the

numbers 1, 3, and 4. More formally, a cycle of length m in a permutation π is a sequence $(x_1, \ldots, x_m)$ such that

$$\text{for } j = 1, \ldots, m-1 : \pi(x_j) = x_{j+1}, \quad \text{and} \quad \pi(x_m) = x_1.$$

You can "walk around a cycle," and when you try to "walk around the permutation," you are always stuck in one of its cycles. For example, the permutation described by equation (2.1) has one cycle of length 3: $(1, 4, 3)$. This is because the permutation satisfies $\pi(1) = 4$, $\pi(4) = 3$, and $\pi(3) = 1$. The other cycle of equation (2.1) has length 2: it is the cycle $(2, 5)$.

2.5 UNDERSTANDING THE SOLUTION

Using the notation introduced above, let us describe the distribution of numbers into 20 boxes by a permutation $\pi \in S_{20}$: Box 1 contains number (or ID) $\pi(1)$, box 2 contains number $\pi(2)$, and so forth:

box number	1	2	3	$\cdots$	20
ID	$\pi(1)$	$\pi(2)$	$\pi(3)$	$\cdots$	$\pi(20)$

Remember that permutations are bijective (i.e., all the boxes contain different IDs). For now, we model the distribution of the IDs by randomly choosing a permutation π from a uniform distribution on S_n. That is, when the game is set up, each permutation can be chosen with probability

$$\frac{1}{20!}.$$

How can we interpret the proposed strategy in terms of permutations? Player j looks into box number j and finds the ID with number $\pi(j)$. She then looks into box number $\pi(j)$, finds ID $\pi(\pi(j))$, and so on. The key step is to understand when the game is lost. Let us therefore consider a small example with 10 players, where each player is allowed to look into 5 boxes, say:

box number	1	2	3	4	5	6	7	8	9	10
ID	5	4	8	6	1	9	3	7	2	10

The distribution of IDs can also be visualised as a graph, where each node corresponds to a box. Each box has a directed edge to another box. The value of the ID hidden under the box determines which box the arrow points to, and it is the next box the player would check if following the strategy.

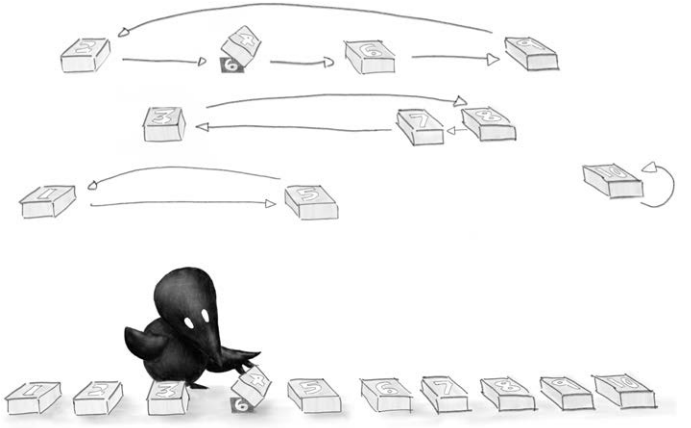

Let us consider the game specifically from player 4's point of view (see the above figure). She starts with box 4 and finds ID number 6 underneath. She then checks box number 6 and finds number 9. In box 9, she finds number 2, and in box 2, finally, she finds her own ID, the number 4. Mathematically, she has walked around a cycle of the permutation!

The strategy guarantees that each player starts in "her own" cycle, the cycle that contains her ID. Moreover, her ID will always be under the last box in the cycle, as this box is "pointing" toward the box that the player started from. Player 4 starts at box 4, and box 2 is pointing toward her starting box, as it contains her ID with the number 4. In this way, each player

walks around her own cycle and continues until she either finds her number, or the maximum number of attempts is reached. In the example above, not only does player 4 find her number, but so do all other players. This is not always the case. In the following example, the strategy fails:

box number	1	2	3	4	5	6	7	8	9	10
ID	3	4	8	6	9	10	2	7	5	1

Visualized as a graph, the situation is now as shown in the following figure.

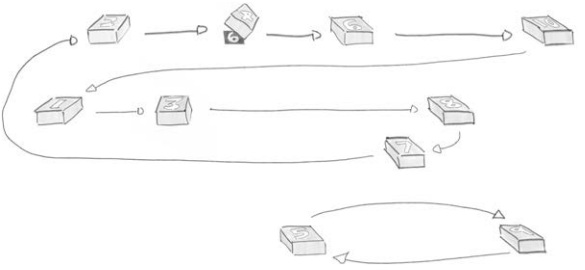

In this example, player 4 fails to find her number; she has to stop after 5 attempts, that is, after looking into box 3. What is the crucial difference between the two examples? The second permutation contains a cycle of length 8: $(4, 6, 10, 1, 3, 8, 7, 2)$. The fact that it is longer than the allowed number of attempts, 5, leads to a disastrous loss of the group: Not only does player 4 remain out of luck, but so do players 2, 6, 10, 1, 3, 8, and 7. These are all the players with a number in the same cycle, and none of them is able to find the correct ID. Players 5 and 9, in contrast, do find their own numbers after looking under just two boxes. If, as in the first example, all cycles are of short length, then all players are guaranteed to find their number.

This finding brings us much closer to understanding the solution. Using the proposed strategy, the game is lost if, and only if, the permutation representing the distribution of ID cards

contains a cycle that is longer than the maximal number of attempts. If all cycles are small (in our example with 20 players, that is, length 10 or smaller) the players will eventually find their own ID in 10 attempts or fewer. If there is a large cycle (length 11 or larger), there will be players (at least 11 of them) who will not find their number, because they can only check 10 boxes and therefore cannot complete the walk around their own, long cycle. To compute the probability of losing, we therefore need to answer the question:

What is the probability that the permutation has a long cycle?

We will come back to this question soon, but first, we point out a similarity between the box game and the hat game presented in chapter 1: The strategy makes the players' outcomes dependent on one another.

Relation to the Hamming Code Solution

This game is similar to the game described in chapter 1, in that the players win and lose together. There, we discussed a solution that collects all wrong answers in single instances of the game and distributes the correct answers over as many games as possible. In the box game discussed here, we observe a similar effect: If there is a long cycle that prevents one player from finding her ID card, many other players will be unsuccessful, too. In this sense, the outcome of the individual players are no longer independent. In many games that are won and lost jointly, the key idea is to introduce a dependence between the individual successes. The challenge is usually finding out how to create that dependence.

Computing the Probability of Winning

To compute the probability that a permutation has a long cycle, let us count the number of permutations in S_{20} that give rise to

a cycle of length k with $11 \le k \le 20$. There are $\binom{20}{k}$ possibilities of choosing the k numbers that are involved in the cycle. Here,

$$\binom{20}{k} := \frac{20!}{k!(20-k)!}$$

is called the "binomial coefficient" and describes the number of different possibilities to choose a set of k numbers out of a set of 20 numbers; see "What Is … a Binomial Coefficient?" (appendix B.4). These k numbers can be sorted in $(k-1)!$ different ways to form a cycle of length k (because we can always start with the lowest number). The cycle leaves us with $20-k$ elements for the rest of the permutation, which can be arranged in $(20-k)!$ ways. Thus, there are $\binom{20}{k}(k-1)!(20-k)!$ permutations with a cycle of length k. To summarize, the number of permutations with a cycle of length 11 or larger equals

$$\sum_{k=11}^{20} \binom{20}{k}(k-1)!(20-k)! = \sum_{k=11}^{20} \frac{20!}{k}.$$

Importantly, we did not count any permutation twice: If we form a cycle of length 15, for example, it is impossible to obtain, in the remaining numbers, another cycle of length 11 that we might have counted before. The probability that a randomly chosen permutation has a cycle of length 11 or larger therefore becomes

$$\frac{\sum_{k=11}^{20} \frac{20!}{k}}{20!} = \sum_{k=11}^{20} \frac{1}{k} \approx 0.67.$$

The probability of winning equals the probability that there is no large cycle, and therefore we obtain

$$p_{\text{win, optimal}} \approx 1 - 0.67 = 0.33,$$

which is surprisingly close to the upper bound of 0.5. Strictly speaking, we do not know yet that this strategy is really optimal, but we will come back to that question later.

More Players

If we keep the rule that players are allowed to look into half of the boxes, what happens to our strategy if we increase the number of players (while keeping it even, say)? Surprisingly, the probability of winning does not change much! The analysis is completely analogous to the one already presented. The probability of losing the game for n players who are allowed to look into $n/2$ boxes equals

$$\frac{\sum_{k=n/2+1}^{n} \frac{n!}{k}}{n!} = \sum_{k=n/2+1}^{n} \frac{1}{k}, \tag{2.2}$$

still, assuming that n is even. The number on the right side of equation (2.2) is not so easy to compute directly, but we can use integrals to approximate it. (If you have not seen integrals before, just skip to the final result of this paragraph.) By drawing the graph of the function $x \mapsto 1/x$, it is not difficult to see that we have

$$\int_{n/2+1}^{n+1} \frac{1}{x}\, dx \leq \sum_{k=n/2+1}^{n} \frac{1}{k} \leq \int_{n/2}^{n} \frac{1}{x}\, dx.$$

But these inequalities imply that the probability of losing is (slightly smaller than)

$$\log(n) - \log(n/2) = \log(2) \approx 0.6931472$$

and converges for $n \to \infty$ to $\log(2)$. (Here, we consider the "natural" logarithm to the base of the Eulerian number; see "What Is … an Exponential Function?" in appendix B.3.) The probability of winning therefore converges from above to

$$1 - \log(2) \approx 0.3068528.$$

Optimality

We mentioned that the above strategy is optimal. Let us now see why this is indeed the case. Remember that we are assuming that the IDs are distributed randomly. First, let us alter the

game: The players enter the room according to their number, starting with 1, then 2, and so forth. Each player is allowed to look into as many boxes as she likes until she finds her ID. The game is won if each player requires $n/2$ attempts or fewer. In terms of winning and losing, this game is equivalent to the original version, so we call this slightly adapted version the "original game" below. Now, let us change the game even further, so that we create a new game: All players are present all the time, and after the first player has found her ID, all players whose ID has been revealed by the first player are out of the game, together with the corresponding boxes. Then the player with the smallest number among those whose ID has not yet been found starts to check some of the remaining boxes until she finds her number, and so forth. As before, the game is won if none of the players used more than $n/2$ attempts. There are two key observations:

1. Any strategy in the original game can be applied to the new game and will yield (at least) the same probability of winning in the new game.

Let us now assume that, for the new game, we construct a protocol by writing down the order in which the IDs have been revealed (ignoring the information regarding which player revealed it). This protocol order lets us reconstruct what happened during the game. For example, the protocol

$$(19, 8, 1, 10, 3, 9, 20, 5, 18, 11, 12, 7, 6, 4, 15, 2, 17, 16, 13, 14)$$

tells us that the first player used 3 attempts, the second player opened 13 boxes, player number 13 opened 3 boxes, and, finally, player 14 opened the last box. This order from the protocol can again be interpreted as a permutation: one that contains the cycles $(19, 8, 1)$, $(10, 3, 9, 20, 5, 18, 11, 12, 7, 6, 4, 15, 2)$, $(17, 16, 13)$, and (14). We see that the new game is won if and only if that permutation does not contain a cycle that is longer

than 10 (if it did, one of the players would have looked into more than $n/2$ boxes).

2. The protocol describes a permutation that is chosen (uniformly) at random.

Why is this? The protocol order clearly depends on the team's strategy, but since we started with a random permutation in the first place, the protocol order will be a random permutation of the numbers $1, \ldots, 20$, independently of what strategy the players have used. The permutation induced by the protocol is therefore also chosen uniformly at random (this is because there is a one-to-one correspondence between protocol orders and protocol permutations).

In summary, the probability of winning the new game is independent of the teams' strategy and equals the probability that a random permutation has no cycle that is longer than $n/2$. Finally, observation 1 tells us that this probability is larger than that of success for any strategy in the original game. This proves that the above strategy is indeed optimal for the original game.

Playing against an Adversary

If the audience knows about the strategy, they could try to set a trap for the players. Instead of choosing a random permutation, they could deliberately choose one with a long cycle. If the players follow their own strategy, they will definitely lose the game. However, the players can protect themselves from an audience that is working against them: They can draw a random permutation τ before starting to play, and then they can reinterpret each ID number k, including their name and the numbers that they find as $\tau(k)$ (but not the box labels). Again, they will end up with the above-computed winning probability.

2.6 SHORT HISTORY

The game was introduced by computer scientists Anna Gál and Peter Bro Miltersen in 2003 [Gál and Miltersen, 2003]. The game was used to illustrate the proof of one of their theorems, and they won a best-paper award for their conference paper. In their original version of the game, each player's ID card is either red or blue (which the players do not know), and every player has to guess the color of their ID card while turning over only half of the boxes. The probability of winning is marginally higher in this game, as players can of course just guess when they have not found their card. Some further variations of the game can be found in a paper by Navin Goyal and Michael Saks [Goyal and Saks, 2005], while a nicely written overview of the history can be found in an article by Eugene Curtin and Max Warshauer in [Curtin and Warshauer, 2006].

2.7 PRACTICAL ADVICE

When playing the game, one can consider increasing the proportion of boxes that the players are allowed to check. This yields a larger probability of success, which might increase the players' motivation. For example, when allowing the players to look into 3/4 of the boxes, the probability of losing equals

$$\sum_{k=3n/4+1}^{n} \frac{1}{k},$$

and therefore the probability of winning converges to

$$1 - (\log(n) - \log(3n/4)) = 1 - \log(4/3) \approx 0.7123179.$$

This game is more likely to yield a successful run, especially if the game is repeated a few times.