

Part I

Mind Games

If you can solve these puzzles, you're not management material. You're too good.



Note: This Screaming Head icon denotes an especially difficult puzzle.

COMPETITION

We can't all be winners.

Sweet Tooth

Byzantine Bettors 

A Touch of Luck

Information Gain 

Reach for the Sky!

Pork Politics

Social Games

Escape Management

Flu Math

Sweet Tooth

Two children, perhaps similar to ones you know, love cake and mathematics. For this reason, Jeremy convinces Marie to play the following game on two identical rectangular cakes chef Martine has prepared for them.

Jeremy will cut the first cake into two pieces, perhaps evenly, perhaps not. After seeing the cut, Marie will decide whether she will choose first or allow Jeremy to do so. If she goes first, she will take the larger piece. If she goes second, she can assume that Jeremy will take the larger piece.

Next, Jeremy will cut the second cake into two pieces (remember that one of the pieces can be vanishingly small if he so chooses). If Marie had chosen first for the first cake, then Jeremy gets to take the larger piece of the second cake. If Marie had chosen second for the first cake, then she gets to take the larger piece of the second cake.

Warm-Up

 Assuming each child will strive to get the most total cake possible, what is an optimal strategy for Jeremy?

Hint: Before you look at the solution: Assume that Jeremy divides the first cake into fractions f and $1-f$ where f is at least $1/2$. Then explore the consequences if Marie chooses to take the piece of fraction f or if she goes second, so gets the piece having fraction $1-f$.

Solution to Warm-Up

Following the hint, Marie reasons as follows: If she takes the fraction f piece, then Jeremy will take the entire second cake (he'll cut the smallest crumb from the second cake and then will take the large first piece, leaving Marie only the crumb). So, Marie will get exactly f and Jeremy will get $(1-f) + 1$. If Marie takes the smaller piece of the first cake (fraction $1-f$), Jeremy will do best if he divides the second cake in half. This gives Marie $(1-f) + 1/2$. Jeremy follows this reasoning, so realizes that the best he can do is to make $f = (1-f) + 1/2$. That is $2f = 1 + 1/2$ or $f = 3/4$. If Marie takes the first piece of the first cake, then Jeremy will get $1/4$ of the first cake and all of the second cake. If Marie takes the second piece of the first cake, then Jeremy will get $3/4$ of the first cake and $1/2$ of the second cake. In both cases, Marie gets a total of $3/4$ of a cake and Jeremy $1 + 1/4$. Note that if Jeremy cuts the first cake such that the larger fraction is less than $3/4$, Marie will get more than $1/4$ of the first cake by going second and will still get $1/2$ of the second cake, thus increasing her cake amount beyond $3/4$. By contrast, if Jeremy cuts the first cake such that the larger fraction is more than $3/4$, Marie will just take that larger piece, again increasing her cake amount beyond $3/4$.

I have discussed the warm-up at great length, because a week later a harder challenge has arisen. Chef Martine has made three new identical rectangular cakes. Jeremy and Marie both eye them greedily.

They decide on the following rules. Again, Jeremy will cut. But this time, Marie gets the first choice twice and Jeremy only once. That is, Jeremy cuts the first cake. Marie decides whether she wants to choose first or second for that cake. Next, Jeremy cuts the second cake; again Marie decides whether she wants to choose first or second. Same for the third cake. The only caveat is that Marie must allow Jeremy to choose first at least once.

1. How does Jeremy maximize the amount of cake he gets, given these rules? How much will he get?
2. Suppose there are seven cakes and Marie gets the first choice for six of the seven cakes. Who has the advantage? By how much?
3. Is there any way to ensure that each child receives the same amount of cake, assuming Jeremy always cuts?



Byzantine Bettors



In the popular imagination, Byzantium is famous for the incessant intrigues and plotting that supposedly took place among the courtiers of the palace. Apparently, the Byzantine court in no way deserved this reputation. Modern historical research suggests that Byzantium achieved a remarkable stability and harmony by the standards of the first millennium. Reputations die hard, however, and this puzzle concerns a game show inspired by the imagined intrigues of Byzantium. We call it Byzantine Bettors.

Each bet works as follows. There are some number of “advisors” and you. One advisor will write either 1 or 0 on a piece of paper, show it to the other advisors but not to you, and put that paper face down in front of you. Each advisor will tell you what the value is. They are all very good actors, so no obvious tick or facial expression will reveal to you who is telling the truth. The amount you can wager on each bet is anything from 0 to the total amount you have.

Warm-Up



Suppose there are four advisors, two of whom always tell the truth, though you don't know which ones. You get three even money bets. You start with \$100. How much can you guarantee to win?

Solution to Warm-Up

If all four or three out of four agree on the first bet, then bet the maximum. There must be at least one truth teller in that group. If they are two and two, then bet nothing. After the first bet, you will know who the truth tellers are and will be able to bet the maximum each time. Thus, at least two times you will be able to bet all you have and be sure to win. This will give you a total of \$400 at the end.

1. Now suppose there are only three advisors and only one of them always tells the truth. Again, you can make three even money bets. You start with \$100. How much can you guarantee to win?

The game has just become a little longer but also a lot more nasty. You have four bets, but there is no longer a truth teller, but merely a "partial truth teller." That advisor is not guaranteed to tell the truth all the time, but must do so at least three out of four times. Further, the advisors can actually substitute what's on the paper for a worse result for you once they hear your bet. However, they cannot change the result if so doing would eliminate the possibility that at least one of them is a partial truth-telling advisor.

2. What can you guarantee to win in four bets, with four advisors, three of whom can lie at will, and one who must tell the truth at least three out of four times?
3. Under the reliability conditions of one partial truth teller who tells the truth four out of five times and three liars at will, but assuming you have five bets, can you guarantee to end with at least \$150?



A Touch of Luck

The movie *Intacto* posits a world in which luck is a more or less permanent quality of a person. Lucky people are lucky at the gambling table and they are lucky at traffic intersections. But such people have to avoid physical contact with a certain group of people who know how to steal their luck by touch. Finding lucky people is the quest of one of the movie's protagonists. In one case he tests the luck of possible recruits by having them run blindfolded through a forest. The winner is the one who gets to the destination first. Many others crash into one of the trees.

We will try a physically gentler game. In its general form, there are N people and B chances to bet. Every player knows the value of both N and B . Each person starts with an initial (not necessarily equal) wealth in units.

Each bet is an even money bet depending on the flip of a shared fair coin. So, if you bet an amount x and win, your wealth increases by x more. Otherwise you lose the x wagered.

Before each flip, each person bets an amount of his/her choosing (from 0 to any amount he/she has at that point) on either heads or tails.

The winner is the person having the greatest number of units after all B bets are done. If two people tie, then nobody wins. The units are worthless after the betting is done. So, a player receives a reward if and only if he or she is the absolute winner.

Warm-Up

 Bob and Alice have the same number of units and Bob must state his bet first. Suppose there is just one remaining flip of the shared coin. What is Alice's chance of winning?

Solution to Warm-Up

If Bob bets x on heads, then Alice could bet $x+1$ on heads and Alice will win on heads and will lose on tails. Alternatively, Alice could bet nothing and then Alice will win on tails. Either way, Alice has a probability of $1/2$ to win.

Warm-Up 2

 Alice and Bob are the only players again. Alice has more units than Bob. There are five more coin flips to do. If Bob must state his bet before Alice on every flip, then how can Alice maximize her chance to win?

Solution to Warm-Up 2

Alice can win the game every time. For each flip, Alice simply copies Bob's bet. Suppose Bob bets b on heads. Alice bets b on heads too. Whether the shared coin lands on heads or tails, Alice will end up ahead of Bob.

Now here are some challenges for you.

1. Bob, Carol, and Alice play. Alice has 51 units, whereas Bob and Carol have 50 each. Bob must state his bet first, then Carol, and then Alice. Bob and Carol collude to share the reward if either wins. How can Bob and Carol maximize the probability that at least one of them will win if there is just one coin flip left?
2. Does the result change if Alice must state her bet first?
3. Suppose Bob has 51 units, and Alice 50. There are two coin flips left. Bob bets first for the penultimate flip. Alice bets first for the last flip. Does Bob have more than a $1/2$ chance of winning? If so, how much more?
4. Suppose Bob has 51 units and Alice 50. Again, there are two coin flips left. This time, Alice bets first for the penultimate flip. Bob bets first for the last flip. Does Bob have more than a $1/2$ chance of winning? If so, how much more?
5. Suppose Bob has 51 units and Alice 50. Again, there are two coin flips left. Again, Alice states her bet first in the penultimate round and Bob states his bet first in the final one. This time, Bob announces that he will bet 20 in the penultimate round, though he will wait to see Alice's bet before saying whether he will bet on heads or tails. Can Alice arrange to have more than a $1/2$ chance to win?

Like joining a ballet company, winning the Olympics, or getting ahead in a narrow hierarchy of power, the more competition there is for fewer spots, the more risk one must take. Let's see whether you agree.

6. Bob, Alice, Rino, and Juliana have 100 units each and there are two flips left. Each person is out to win — no coalitions. Bob and Alice have bet 100 on heads. Rino has bet 100 on tails. Juliana must now state her bet, knowing that she will state her bet first on the last flip. What are her chances of winning if she bets 90 on either heads or tails? What should she bet?



Information Gain



The well-known and always well-dressed game show host Jeff Nicholas has approached Jordan and his friends Ariana, Bob, Caroline, David, and Ellen with a contest proposition. Jordan and the five are the leading lights of the omniheurist club, a group of outstanding puzzle solvers.

“Our contest is live. I will blindfold your friends, then put a hat on each of them bearing a number between 1 and 10 (more than one person may have the same number) and lead them into a televised game room. Once they have arrived, I will arrange them in a circle in an order that I choose and then replace each blindfold by very dark but non-reflecting sunglasses to eliminate the possibility of eye signals.

“You and the audience will see them in the game room and the numbers on their heads via the television monitor, but they won’t be able to see you. I will be holding one blue ticket and one red ticket. You can ask me to deliver one ticket to one of the five. That’s all you can do. No knocking on the window or your team will be disqualified.

"The mathematicians may not speak or signal one another in any way, or the entire team will be disqualified. (Obviously, I won't help them in any way other than delivering the ticket.) They will, however, see to whom the ticket is being delivered and the color of the ticket as well as the number on each other person's hat. There is no way they can see the numbers on their own heads.

"At my signal each person will put up a certain number of fingers. If the number of fingers matches the number on that person's hat, then he or she will receive that many thousands of dollars. If all of them win, then you, Jordan, will receive \$5,000. If any lose, then you have to buy me a new Armani suit."

"That's it?" Jordan replied. "The only information they receive from the outside is who gets a ticket and the color of that ticket?"

"That's right," said Nicholas. "Remember, though, that each person sees the numbers on the other people's hats too. You may not be able to solve this puzzle. I do want that suit."

1. Is it possible for Jordan to design a protocol to ensure that each of his friends will raise the correct number of fingers? If so, explain it. If not, does Jordan have a high probability of winning?

Warm-Up

 Here's a much easier problem to give you an idea of the kind of protocol Jordan might design. Suppose Jeff Nicholas were required to put consecutive numbers on the five hats (as in 4, 5, 6, 7, 8). Then what could Jordan do?

Solution to Warm-Up

Jordan could arrange the following protocol with his friends. Let Ariana represent 1, Bob represent 2, Caroline represent 3, David represent 4 and Ellen represent 5 and 6. If Jordan sends a ticket to Ariana, then the consecutive numbers start with 1. If to Bob, then they start with 2. If to Caroline, then 3. If to David, then 4. If a blue ticket to Ellen, then 5. If a red ticket to Ellen, then 6.

Thus, when Jordan sends in a ticket, each mathematician will know the first number in the sequence. And by looking at the numbers on his or her fellow players' hats, that mathematician can work out by process of elimination which number appears on his or her own hat.

In the Jeff Nicholas challenge, however, the numbers are not necessarily consecutive or even all distinct. Do you think it can be done?



Reach for the Sky!

Reach for the Sky! is based on hollow lottery balls. Here is how it is played.

Your adversary is given 100 identical pieces of paper. In a room hidden from your view, he writes a number on each one—he can choose any value in the range 1 to 1 million but without duplicates—then inserts them into small opaque envelopes. Subsequently, these papers are given to an independent third party called the “stuffer.” The stuffer works in full view of you and your adversary at all times. He shows you 100 hollow and empty lottery balls. He then shuffles the envelopes and inserts one envelope into the first lottery ball, another into the second, and so on up to 100. After inserting each envelope, he screws the ball shut. These balls have been examined using drop tests, bounce tests, and resiliometric tests to be sure they all share the same physical properties.

The stuffer then puts the balls into a lottery machine. The lottery machine mixes the balls thoroughly until one comes out. Call that ball number 1. The ball is opened and you are told the value on the paper (what we will call the ‘paper value’ for short). You have the option to ‘capture’ that paper value or not. If you capture it, you put it in your capture pile and you have used up one capture. If you don’t, it goes in the discard pile never to come out, though you may find it useful to remember the discarded paper values. Repeat this procedure for all 100 balls. You are allowed three ‘captures’ altogether. Your goal is to capture the highest paper value written by your adversary. If you do, you win \$100,000. If you don’t, you lose \$100,000. Should you engage in this bet? If so, what is your probability of winning?

This puzzle is a variant on the sultan’s daughters problem. A young suitor may choose any of the 100 daughters of the sultan. They are presented to him in some random order. He has little basis for judgment, so judges only by outward beauty and grace. If he rejects one, he never sees her again. Once he selects one, he must marry her and no other.

Warm-Up

 Can you design a strategy that gives the suitor at least a $1/4$ chance of marrying the most beautiful daughter, assuming the daughters are presented to him in an order that is unlikely to be related to their beauty, e.g., based on the hour and minute of the day in which they were born?

Solution to Warm-Up

Look at but reject the first half of the daughters. Then take the first daughter who is more beautiful than any of those you have seen so far. This does not guarantee that you will marry the most beautiful daughter (or any daughter at all), but it has the virtue of offering an easy rough analysis: you have a $1/2$ chance that the most beautiful daughter is in the second half and a $1/2$ chance that the second most beautiful daughter is in the first half. The two together arise with probability $1/4$, assuming the daughters are presented to you in an order that is independent of their beauty—a situation that likely holds in the hour-minute ordering mentioned above. That case is sufficient to ensure that you marry the

most beautiful daughter. In fact, other cases also lead to this outcome. For example, the third most beautiful daughter could be in the first half and the most beautiful one could precede the second most beautiful one in the second half. A deeper analysis shows that it is better, in fact, to reject the first 37 daughters and then choose the next one that is more beautiful than any you have seen.

1. What would be a good strategy for the \$100,000 Reach for the Sky! game, where you have three captures? Using that strategy, what would be your chance of winning?

Hint: A bit of programming might come in handy.

2. How would your answer change if there were 1,000 lottery balls?



Pork Politics

A grant of a whopping 100 million dollars is available to the five districts making up a state. The districts have representation proportional to their population in the state assembly. District A has 35 representatives, district B has 25, district C has 16, district D has 14, and district E has 10.

If a coalition consisting of at least 51 members forms, then they can force the funds to come to their districts, where the money will be divided proportionally to the representation. For example, if the districts having 35 and 25 get together, they can force all 100 million to go to their two districts and that money will be divided such that $(35/60) \times 100$ million goes to the district A and $(25/60) \times 100$ million goes to B. Each district has only its own self-interest in mind.

Warm-Up

 What would B most prefer?

Solution to Warm-Up

B prefers the coalition {B, C, E} to any other winning coalition, because B would get the fraction $25/55$ of the money, whereas if B allied itself with A, it would get only $25/60$ of the money.

-
1. Which coalitions might form?

The representatives of B have an ancient privilege. They can control the rules of when an allocation can take place. That is, they may require that an allocation require only a simple majority as we've assumed until now, but they can also demand that it require a 67% majority, or a 75% majority.

2. How strong a majority should B prefer (51%, 67%, or 75%) in order for B to receive as much money as possible?



Social Games

Since Adam Smith observed that the “invisible hand” of the free market would force self-interested manufacturers to offer low prices to consumers, governments and politics have never been the same. It took almost two more centuries, however, to achieve a proper mathematical analysis of the consequences of selfish behavior beginning with Morgenstern’s and Von Neumann’s game theory and the work of John Nash. (You may recall Nash as the mathematician who was the subject of the movie *A Beautiful Mind*.)

This puzzle explores game theory as it applies to social goods. Let’s start with the invisible hand.

Bob and Alice are competitive manufacturers. If they fix their prices at a high level, then they will share the market and each will receive a profit of 3. If Alice decides to lower prices while Bob doesn’t, then Alice will enjoy a profit of 4 while Bob gets a profit of 0, because nobody will buy from him. At that point, Bob will lower his price to receive at least a profit of 1. Similarly, Bob receives 4 and Alice 0 if the roles are reversed. Simply following their self-interest, both will lower their prices and their profits will drop to 1 each.

This arrangement can be expressed in the following table, where Bob’s profit is shown on the left in each pair and Alice’s profit is on the right. So, the upper right corner, for example, represents the state in which Bob charges a high price (and receives a profit of 0) and Alice charges a low price (and receives a profit of 4).

	Alice High	Alice Low
Bob High	3, 3	0, 4
Bob Low	4, 0	1, 1

John Nash defined the concept of an *equilibrium state*, since known as a *Nash equilibrium*, in which no party has an interest in deviating from that state. The only Nash equilibrium in this case is the bottom right corner. Neither Bob nor Alice will unilaterally raise prices. If Bob raised prices (thus moving the state to the upper right corner), then his profit would decrease to 0. If Alice raised prices, she would move to the lower left corner, thus decreasing her profits to 0.

This is the invisible hand at work. For the consumers, Bob and Alice’s competition leads to lower prices, a social good. Virtually every modern economy gives evidence of this.

Unfortunately, selfishness does not always lead to socially beneficial consequences. Suppose that instead of representing price choices of competitive manufacturers, the table represented choices about honesty or social responsibility. That is, the “High” row represents states in which Bob acts honestly. By contrast, the “Low” row represents a situation in which Bob cheats (say, steals, pollutes, or bribes lawmakers). The upper right corner portrays a state in which Alice cheats but Bob doesn’t. As you can see, the cheater then benefits. If both cheat, then their benefits go down. Selfishness leads to social loss. Corrupt nations, high crime zones, and brawling families give much evidence of this.

Part I: Mind Games

Game theory is neutral. The same game matrix and the same Nash equilibrium can lead to a good or bad state. Selfishness can yield a social benefit or harm.

Now, let us say that it is your job to design public policy. You are confronted with the matrix representing the selfish benefits of cheating and you want to change it somehow. So you establish a police force and criminal justice system that makes it 10 percent likely that cheaters will be caught and given a value of -5 (e.g., jail time).

In that case, Alice’s benefit in the upper right corner has a 90 percent probability of being 4 and a 10 percent probability of being -5. Thus, her expected benefit is $(4 \times 0.9) + ((-5) \times 0.1)$ or 3.1. In terms of expectations, the matrix becomes:

	Alice Honest	Alice Cheats
Bob Honest	3, 3	0, 3.1
Bob Cheats	3.1, 0	0.4, 0.4

Now, the benefits of cheating when starting from the honest-honest state are much reduced. Further, if the punishment increases or the probability of being caught increases, then honest-honest may become a Nash equilibrium.

Warm-Up

 Suppose you can increase punishments (i.e., make them more negative). Can you do so in a way to make the upper left a Nash equilibrium while keeping the bottom right a Nash equilibrium as well? Assume you cannot change the probability of being caught.

Solution to Warm-Up

Suppose the punishment becomes -8. Alice’s gain in the upper right corner state becomes $(4 \times 0.9) + ((-8) \times 0.1) = 2.8$. The gains at the bottom right become $(1 \times 0.9) + ((-8) \times 0.1) = 0.1$. This yields the matrix:

	Alice Honest	Alice Cheats
Bob Honest	3, 3	0, 2.8
Bob Cheats	2.8, 0	0.1, 0.1

So, honest-honest is a Nash equilibrium, because neither Bob nor Alice has an incentive to cheat to leave this state. On the other hand, the lower right hand corner is also a Nash equilibrium. If both Bob and Alice cheat, then they each still receive a tiny expected gain. If one then becomes honest, that person’s gain is eliminated.

Here are some challenges for you.

1. How would punishment have to change to eliminate cheats-cheats as a Nash equilibrium?

As civilizations have advanced, societies have evolved the notion of punishments that fit the crime. This principle might limit the extent to which overly severe punishments are considered acceptable.

2. If you are a public policy maker and you want to limit the punishment for cheating to -5, how much must you increase the probability of catching criminals to maintain the upper left as the only Nash equilibrium?

In modern societies, inequalities exist so that the benefits of an honest society may benefit one party more than another. Say that Bob gets 5 whereas Alice gets only 2 if both are honest. Assuming for the moment that the probability of being caught is 0, this gives us the following game matrix:

	Alice Honest	Alice Cheats
Bob Honest	5, 2	0, 4
Bob Cheats	4, 0	1, 1

The total social good of the upper left corner has increased at the cost of inequality. One side effect, however, is that incentives to dishonesty now change.

3. Assuming a 10 percent likelihood of being caught, is there a punishment value that would cause the only Nash equilibrium to be the upper right state (Alice commits crime but Bob does not)?
4. Still assuming a 10 percent likelihood of being caught, what must the punishment be to force Alice to stay honest?

Let's close this puzzle on a non-punishment note. Let's return, in fact, to the competing manufacturers Bob and Alice. Recall their game matrix.

	Alice High	Alice Low
Bob High	3, 3	0, 4
Bob Low	4, 0	1, 1

Part I: Mind Games

The invisible hand guides them to lower prices, but neither likes this. They see that if they cooperated, their total profit would be $3 + 3 = 6$ instead of $1 + 1 = 2$ as in the competing scenario. For that reason, they might consider merging their two enterprises into BobAlice and thereby achieving a total profit of 6 instead of 2. That is, when Bob looks at Alice, he doesn't value her company at the profit of 1 that she is currently achieving, but at a profit of 5, the additional profit she could bring to a potential BobAlice relative to Bob in a competing world. Of course, this increased profit to Bob and Alice comes at an approximately equal cost to the consumer.

Next time you hear someone expound about inequality, crime, enlightened self-interest, and antitrust laws, bring out your game matrices. You may find more wisdom in them than in the blah-blah.



Escape Management

A thief has just robbed a bank at the center of a city. He has done the robbery quietly, so he believes he is not being tracked. In fact, though, a silent alarm has alerted the police to the theft and they know the exact position of the thief. Their hope is to catch the thief and his accomplices.

The city is laid out as a 19×19 road grid. The thief (T) starts exactly in the center, that is, nine blocks from any border. Figure 1-1 shows the upper right portion of the whole grid. The thief begins to drive north towards the first intersection in the northern half of the city. At every intersection the car can go right, left, or straight (not backwards). The thief will escape if he arrives at either the northern or southern edges.

The police (P) want to let the thief move to see where T will go with his car, but they want to prevent the car from leaving the grid, at least for a while.

At every intersection, the thief chooses where to go except that the police, by their presence, can force T to go straight or to turn right or left. On the other hand, the police want to leave T as much freedom as possible. They therefore have to decide how to deploy their resources to achieve one of several outcomes.

1. The police want to control T's direction at only five intersections. They know that T is extremely motivated to escape. How long (in terms of single block steps) at most will it take T to arrive at either the northern or southern border?
2. Suppose that the police want to prevent T from reaching the northern or southern borders. Therefore the police plan to control T's direction m out of every n times for some $m < n$. What should m and n be to minimize the fraction of time P forces T to change directions while still preventing T from escaping?




Figure 1-1: The thief T starts where indicated. He is nine blocks from the northern, eastern, southern, and western borders, though only the first two are shown.

3. How does your answer change if T can escape by reaching any border, including the eastern and western ones?



Flu Math



Suppose the government health authorities tell you that you have a 5% chance of dying from a flu vaccine but a 10% chance of contracting and dying of the flu when an epidemic strikes if you don't take the vaccine. Do you take the vaccine assuming it offers perfect protection?

Many people refuse.

This apparent irrationality is commonly attributed to “omission bias” — people tend to prefer inaction to action. But I think there are other reasons. For one thing, the authorities commonly have a bias towards action — if only to justify their own existence. If so, the consequences of the flu may be overstated or the risks of the vaccine may be understated.

But even if these probabilities are to be believed and the bias towards inaction is corrected, there is the observation that the risk of being unvaccinated decreases if most other people take the flu vaccine.

Suppose, in fact, that your likelihood of dying from the flu when unvaccinated goes down according to the following formula: If a fraction f (excluding you) of people take the flu shot, then your probability of dying from the flu is just $(1-f) \times 10\%$. For example, if 65% of the people take it and you are among the 35% who do not, then the probability of your contracting the flu is $0.35 \times 10\% = 3.5\%$.

Warm-Up



Suppose a government official could require 60% of the people to take the vaccine. Then what would be the average risk of death for the entire population due to the flu?

Solution to Warm-Up

Recall that if the official doesn't allow anyone to take the vaccine, then the death toll is 10%. If the official forces everyone to take it, the death toll is 5%. If 60% take the shot, then those people have a 5% chance of dying, but the others have only a 40% chance of dying from the flu, so the death toll among them is 4%. The overall death toll, therefore, is $(0.6 \times 5\%) + (0.4 \times 4\%) = 4.6\%$.

-
1. As long as the government is in a compelling mode, what fraction should be required to take the vaccine in order to minimize the overall average risk?

On the other hand, suppose that the government feels it cannot (or, preferably, should not) compel people to take the vaccine. Instead, the government offers each person the flu shot in turn. If a person refuses, then there is no second chance. Each person knows how many people took the flu shot among those already offered. Each person believes and knows that everyone else believes the government's risk figures (5% if a person takes it; 10% modified by the f formula above if a person doesn't). Each person will take the flu shot if and only if it helps him or her. There is no regard for the greater good.

2. Under these conditions (no compulsion, full belief in government risk figures, selfishness), what percentage of people will take the flu shot?
3. Suppose 25% of the population will not take the vaccine under any circumstances. Nobody else knows beforehand who they are. Then can you determine what percentage of people will in fact take the vaccine under the same assumptions?

The government looks at these results and decides that a little benevolent disinformation is in order. That is, the government will inflate the risk of disease to a number R greater than 10% but with the risk of contracting the flu if one is unvaccinated decreasing based on the f formula, with the 10% replaced by R — that is, $(1-f) \times R$. The exaggerated risk strategy is a carefully guarded secret, so everyone believes the government and knows that everyone else does too.

4. Again, the government will compel nobody. To which percentage should the risk of disease be inflated to achieve the vaccination level you determined in your answer to question 1?

***Disclaimer:** The actual death rates from the flu and flu vaccines are typically far smaller than the numbers used in these examples. Also, the scenario of benevolent disinformation is purely invented. Purely.*



DESIGN

Imagination rules...

Whipping Ice 

Optimal Jargon

Using Your Marbles

Flipping Colors

Scheduling Tradition 

Fractal Biology 

As Easy as Pie 

Whipping Ice



Remember Indiana Jones and his bullwhip? In these problems, Indy must overcome a series of obstacles to reach his goal. Indy stands on a board resembling a skateboard but that cannot tilt. Under the chassis are four metal ball bearings on universal swivels. He travels over a smooth frozen lake that has narrow wooden poles sticking up and out of the ice.

Here are the rules:

- i. If he is going in a given direction and has not wrapped the whip around any pole, then he will keep going in a straight line. He cannot otherwise force the chassis to turn.
- ii. The whip is light, so throwing it in any direction does not change his direction materially.
- iii. He can unwhip a whip from a pole at will.
- iv. He can pull on the whip or just let it rotate him, but pushing the whip has no effect.

Warm-Up

 In Figure 1-2, you see a challenge for Indy to get from one end of a corridor that is partly blocked by two barriers to the other end. Assuming Indy can approach the corridor at any angle, how can he reach the other end without touching any barriers?




Figure 1-2: How can Indy use his whip to navigate from the bottom of the corridor to the top without hitting the barriers or the sides? Assume he can approach the bottom from any angle. The dot represents a pole.

Solution to Warm-Up

Figure 1-3 shows a solution. Indy comes in at an angle and then throws the whip as indicated by the dashed line to rotate around the pole.




Figure 1-3: Indy connects briefly to the pole (the dark circle) to rotate to a line parallel to the walls of the corridor.

Now here are the two problems for you:

1. Does Indy need more poles than the two shown by the dark circles in Figure 1-4 to go from the bottom to the top?




Figure 1-4: Indy wants to go from the bottom to the top without hitting the barriers or the walls. Indy may enter the bottom at any angle. Can Indy do it with the two poles indicated as black circles?

2. There are several L-shaped barriers and several objects (at the Xs in Figure 1-5) that Indy wants to collect in a game called Treasure Skater. What is the minimum number of poles he will need and where should he place them?




Figure 1-5: What is the minimum number of poles Indy will need to skate over all the Xs?



Optimal Jargon

Everywhere you look, you find mini-fads and their accompanying industries. Each comes with a bunch of characteristic phrases that often yield to a kind of popular jargon. To send a text message to someone becomes “to text them.” To find information using the Google search engine becomes “to google it.” Such shortening of phrases is not unique to our time. Movie sets shorten the phrase “we are about to start cameras so actors prepare yourselves and everyone else be quiet” to the single word “camera.” Languages borrow jargon from one another, because of the economy of expression they offer. Even linguistically proud France borrows “milestone” and “IP” from English, whereas it would never think of borrowing, say, “spoon.”

As puzzlists, we can help make better jargon. To show you how, we’ll abstract the phrases into letters and find an encoding in terms of bits. The marketers can find words to fit later. We’ll be very concrete to avoid having to estimate probabilities.

Suppose that every message communicated has 60 characters and each character is one of A, B, C, or D. If you have to encode these in bits and you know nothing else, you might encode them as follows:

A – 00
B – 01
C – 10
D – 11

Because there are 60 characters and each character is encoded in two bits, the whole message requires 120 bits (15 bytes). Whereas this yields a far more succinct expression than the one byte per character encoding of ASCII, for example, some more information might improve things still further.

Warm-Up

 Suppose that we knew that every message of 60 characters consisted of exactly 30 As, 15 Bs, 10 Cs, and 5 Ds. What would be a good encoding in bits of the characters then?

Solution to Warm-Up

Intuition tells us that a good encoding should render A in the fewest bits, then B, then C, and then D. It is that intuition that guided Samuel Morse in his invention of the code that bears his name. But Information Theory could conceivably help us do better.

As originally conceived by Claude Shannon, Information Theory was closely inspired by a statistical mechanics interpretation of thermodynamics. Shannon defined a notion of “entropy” as a formula involving the probability of each character (e.g., A has probability $30/60$ or $1/2$; whereas D has probability $5/60$ or $1/12$), and logs of probabilities. The entropy describes the weighted average length of an optimal encoding in bits of a single

Part I: Mind Games

character. The entropy formula in this case yields: $(1/2 \log(2)) + (1/4 \log(4)) + (1/6 \log(6)) + (1/12 \log(12))$ where all the logs are to the base 2. This comes out to about 1.73 bits per character. Of course, a designer is free to choose the encoding he or she wishes, and we choose a whole number of bits for each character, as follows:

A – 1
B – 01
C – 000
D – 001

In our encoding, sending A thirty times would cost 30, sending B 15 times would cost 30 as well, and so on. The total length would be $30 + (15 \times 2) + (10 \times 3) + (5 \times 3) = 105$ bits. This comes out to $105/60$ or 1.75 bits per character. So the intuitive design is quite good.

So far, we've considered only the single character frequencies. We may in fact know more. For example, in English, "q" is rare and so is "u," but "u" almost always follows "q." So we might consider encoding "qu" as a single character. Could this help us?

1. Suppose that in addition to the frequencies given in the warm-up, you know that in every message B is always followed by A and D is always followed by C. Could you arrive at a shorter encoding of the message?

To be really useful, jargon must reduce long phrases to a word or two. So, let's see if we can simulate that case.

2. Suppose that in addition to the frequencies given in the warm-up, you know that in every message, D is always followed by C, C is always followed by B, and B is always followed by A. Could you arrive at an even shorter encoding of a 60-character message? Which phrases should be rendered as jargon?

English is famous for its exceptions to almost all rules (e.g., the country "Iraq" has a "q" that is not followed by a "u"). Such exceptions could force an optimal encoding to include a codeword for "q" as well as one for "qu." But if there were many exceptions, then the extra codewords might not be worth it. Let's see how this plays out in our mini-language.

3. Suppose that in addition to the frequencies given in the warm-up, you know that in every message, D is always followed by C, C is always followed by B, and B is followed by A 13 out of 15 times? What is the shortest encoding you could find of a 60-character message?



Using Your Marbles

Your friend Carol loves to prepare gifts of marbles in bags for children. On Monday, she has 5 red marbles, 6 blue marbles, and 7 white ones.

Warm-Up



How many different bags can she pack so each bag will have a different collection of marbles? Two bags are different if for at least one color, one bag has a different number of marbles of that color from the other. For example, a bag with two red marbles and one blue is different from a bag having one red and one blue or a bag having one red and two blue.

Solution to Warm-Up

For brevity, the colors will be represented by their first letters: R (red), B (blue), and W (white). Carol can pack ten different bags: R, B, W, RR, BB, WW, RB, BW, RW, BWW.

That was just the warm-up.

1. On Tuesday, she again starts with 5 red marbles, 6 blue marbles, and 7 white ones. In order to make the bags more varied, she wants to pack the bags so that each bag differs by at least two “in-or-outs” from any other bag. An in-or-out consists of inserting a marble into a bag or removing a marble from a bag. So, for example, R and RR differ by only one in-or-out (insert R). On the other hand, RW and RB differ by two (remove W and insert B). How many different bags of marbles can Carol create that differ by at least two in-or-outs?


On Wednesday, Carol is given a bag that she cannot open.

2. She knows she has 18 marbles. She also knows that there is at least one red, at least one blue, and at least one white. Knowing only this information, but before seeing the marbles, Carol receives a phone call. “Can you guarantee to be able to give each child a bag with a different collection of marbles (i.e., at least one in-or-out apart), if there are eight children at a party? If not, then can you make this guarantee if there are seven children? If so, then how about nine?” How should Carol answer?
3. On Thursday, she again knows only that she has 18 marbles and at least one of each color, and is in the same situation as Tuesday as far as wanting the children to enjoy variety. So she wants each gift to differ from every other by at least two in-or-outs. How many children can she guarantee to prepare bags for in that case?

4. On Friday, her friend Diane packs the bags. Diane assures Carol that (i) every bag contains a different (some difference of number in at least one color) collection of marbles, (ii) that there are 18 marbles to start with, (iii) there are more white than blue and more blue than red, and (iv) that the maximum different ones she could pack was seven. Assuming that Diane is an extremely capable packer, what is the maximum number of red marbles there could have been at the beginning?



Flipping Colors



The well-known knight's tour problem is the challenge of placing a knight at a starting position on a chessboard and having it visit every square on the board exactly once. You can look up the knight's tour problem on the Web to get inspiration for the much harder puzzle I'm going to ask you to do. In fact, you may decide this problem is not even possible.

This puzzle also concerns a knight, but we are going to suppose that each knight move consists of traveling two squares vertically and one square horizontally or two horizontally and then one vertically. Each move of the knight flips the colors of the three squares along the L-shaped path chosen, (not including the starting square, but including the landing square). The problem is to place the knight at an initial square of your choice, then move the knight through a series of L-shaped moves in such a way that the color of each square on the board is the reverse of the color before the knight began its journey. Assume that when you place the knight on a square to start, you flip the color of that square first.

Note that the knight may flip the color of a square more than once but must ultimately flip it an odd number of times.

1. Can the knight flip the color of every square an odd number of times? If so, in how few moves?



Scheduling Tradition



There are 12 school teams, unimaginatively named A, B, C, D, E, F, G, H, I, J, K, and L. They must play one another on 11 consecutive days on six fields. Every team must play every other team exactly once. Each team plays one game per day.

Warm-Up



Suppose there were four teams A, B, C, D and each team has to play every other in three days on two fields. How can you do it?

Solution to Warm-Up

We'll represent the solution in two columns corresponding to the two playing fields. Thus, in the first day, A plays B on field 1 and C plays D on field 2.

AB CD
AC DB
AD BC

Not only does the real problem involve 12 teams instead of merely four, but there are certain constraints due to traditional team rivalries: A must play B on day 1, G on day 3, and H on day 6. F must play I on day 2 and J on day 5. K must play H on day 9 and E on day 11. L must play E on day 8 and B on day 9. H must play I on day 10 and L on day 11. There are no constraints on C or D because these are new teams.

1. Can you form an 11-day schedule for these teams that satisfies the constraints?

It may seem difficult, but look again at the warm-up. Look in particular at the non-A columns. They are related to one another. If you understand how, you can solve the harder problem.



Fractal Biology



Biologists study not only genes, but also genomes, proteomes, interactomes, basically anything *omic*, that is, anything having to do with entire species. It's a perspective that might hurt many delicate human egos.

For example, we have roughly the same number of genes as mice and most are very similar in form and function. We think of ourselves as a higher form of life, but the genome itself doesn't bear that out. The difference must be in the interactions, the networks of binding and repulsion that give rise to the unique capabilities of each species.

A puzzle that obsesses biologists is to figure out why protein interaction networks have the topologies they do. You see, most proteins have very few connections, and very few have many connections. We call members of the latter groups the *hubs*. This gives rise to what are sometimes called *scale-free* networks reminiscent of Mandelbrot's fractals or in fact the Web. You can build a scale-free network as follows: Make 70 percent of all nodes have one edge and the remaining 30 percent have two or more. Of the 30 percent, 70 percent have two edges exactly and 30 percent have three edges or more. Of these 30 percent, 70 percent have three edges exactly and 30 percent have four edges or more, and so on.

One theory is that scale-free networks have good failure properties. The thinking goes like this: Because mutations strike randomly at the genome, most mutations will wound proteins that interact with very few other proteins, thus are presumably less important. Fatal mutations of hub proteins can occur too, but they are rare, because hubs are rare.

A typical problem is to determine how many interaction links are needed to achieve high fault tolerance while ensuring that no unwounded protein is more than two interaction links away from another one.

If there were just four protein nodes, then a simple square of interactions would achieve this goal (see Figure 1-6). Removing any node allows any remaining pair of nodes to communicate over at most two links.



Figure 1-6: A four-node network in which there is a one- or two-link path from any node to any other even if a node is deleted.

Warm-Up



Can you achieve this distance two condition for six nodes even if one node is wounded (call that the *wounded distance two condition*), using ten links or fewer and assuming no node can have more than three interaction links?

Solution to Warm-up

Figure 1-7 shows a nine-link solution.




Figure 1-7: A six-node network in which there is a one- or two-link path from any node to any other even if a node is deleted.

-
1. Can you achieve the wounded distance two condition for eight nodes using 16 links, assuming no node can have more than four interaction links?
 2. What is the fewest number of links needed to achieve the wounded distance two condition for twelve nodes and at most five interaction links for any node?
 3. What is the fewest number of links needed for 12 nodes, but without any limit on the number of interaction links any particular node can have?
 4. We have a particular network having 108 proteins. We don't yet know all the interactions. What is the fewest number of links needed to achieve the wounded distance two condition for any pair among these 108 nodes if there is a limit of 60 interactions that any single node can have? Try for a solution that uses fewer than 600 links.



As Easy as Pie



If you've ever asked a young child to cut a pie for several people, the child will likely cut a piece for himself or herself and then pass you the knife. Perpendicular cuts across the length of the pie come naturally to adults for the sake of overall efficiency. Young children simply find this to be too altruistic. This puzzle is an attempt to help you recapture that inner child.

You will design a series of cuts to partition a square pie into equal pieces to share among several people. The piece finally offered to each person is called a "final piece." Here are the rules:

- i. All cuts must be straight, vertical, and parallel to one of the original sides.
- ii. All final pieces must contain the same amount of pie volume (and, because of the first rule) the same amount of the top of the pie.
- iii. (Child cut rule) Every cut should yield one final piece except the last one, which yields two final pieces.

The goal, which admittedly might occur only to a lazy, clever, or lazy and clever child, is to minimize the sum of the perimeters of all the final pieces.

Warm-Up



We start with the square pie in Figure 1-8.





Figure 1-8: Two possible ways of cutting the original pie (above) into three pieces of equal size. The bottom left shows two parallel cuts. The bottom right shows two perpendicular cuts.

We could cut that square pie using two parallel cuts (A) or using two perpendicular cuts. Both satisfy the three rules: equal size final pieces, vertical cuts, and child cut rule. Which yields a smaller total perimeter?

Solution to Warm-Up

In the parallel cut case (A), each final piece has a perimeter of length $1 + 1/3 + 1 + 1/3 = 8/3$. So the three final pieces altogether have a total perimeter of 8. Two perpendicular cuts (B) yield a first final piece having perimeter $8/3$ and the remaining two final pieces each having perimeter: $2/3 + 1/2 + 2/3 + 1/2 = 7/3$. In this case, the total perimeter is $7 \frac{1}{3}$.

Now it's your turn.

1. Find a cut design obeying the above three rules that minimizes the total length of the perimeters for five pieces.

For the remaining questions, let's pretend we're working with adults. So, we drop the child cut rule, but keep the first two.

2. How much better can you do for five pieces, after dropping the child cut rule?
3. What about nine pieces?

Hint: There is a Russian doll feeling about this one.

4. Can you get a smaller perimeter for five pieces, if, in addition to dropping the child cut rule, you drop the rule that cuts must be parallel to the original sides?

Try this puzzle on a smart child sometime. Enjoy the reaction you get. In my experience, someone always ends up laughing.



CHANCE

Getting on the right side of luck

Lucky Roulette

Legal Logic

The Box Chip Game 

Feedback Dividends

Lucky Roulette

Interviews at hi-tech companies often include puzzles. Here is one that has made the rounds in spite of its violent setting. We will spice it up a bit.

I put two bullets in two adjacent chambers of a six shot revolver. I point it at your head and pull the trigger. Click. You are still alive. The chamber has advanced by one. I am prepared to try again.

Warm-Up



Is it better for you if I spin again or not before pulling the trigger a second time?

Solution to Warm-Up

It's better for you to keep the chamber where it is, i.e., not to spin. Why? Well, let E represent an empty chamber and B represent a chamber with a bullet.

Putting the chamber first used at the beginning on the left, the gun could have been in any of the following configurations with equal probability:

BBEEEE
EBBEEE
EEBBEE
EEEEBE
EEEEBB
BEEEEB

Because you didn't die, the equi-probable configurations left are:

EBBEEE
EEBBEE
EEEEBE
EEEEBB

If I don't spin, then you die in only 1 in 4 cases on the second trigger pull. If I spin, you have a 1 out of 3 chance to die. So, you don't want me to spin.

Now suppose I am going to pull the trigger two more times for a total of four. If you survive the second trigger pull, then you must be in one of these cases:

EEBBEE
EEEEBE
EEEEBB

1. Do you ask for a spin of the chamber before the third trigger pull? How about before the fourth trigger pull? Using your best strategy, what are your chances of survival given that you've survived the first trigger pull?



Legal Logic

In my last year of college, I briefly flirted with the idea of going to law school. The reasoning seemed kind of fun. Lots of people (some might say too many) were doing it. But I chose technology, deciding that for the most part lawyers divide an existing wealth pie and take a slice for themselves: a zero sum game for society. A career in technology promised the possibility of increasing the size of the pie. Of course, my analysis was simplistic, but I still think fundamentally correct. Technology opens new worlds. Legal reasoning retreads paths that the Mesopotamians blazed more than 4,000 years ago. Still, my career has since brought me in close contact with very smart lawyers. I have seen how they argue and how their arguments affect society. I offer you a small axiomatization of what I think is going on and then propose a modest reform.

Suppose there is a new device (or drug) D aimed at sickness S . Legal logic plays with the following predicates:

$\text{mayuse}(x, D)$ — Person x may use device D (i.e., D is available).

$\text{hurtby}(x, D)$ — Person x has been hurt by device D .

$\text{newdevice}(D)$ — Device D is new.

$\text{sick}(x, S)$ — Person x is sick with sickness S .

$\text{sueandwin}(x)$ — Person x sues someone with a good probability of winning.

This is first order predicate logic so:

" $\text{ThereExists } x \text{ such that } \dots$ " means at least one person x has the property represented by the ellipses.

" $\text{ForAll } x \dots$ " means that all people have the property represented by the ellipses.

$A \rightarrow B$ means that if A holds, then B must hold. If A doesn't hold, then B need not either.

❑ Rule 1) $\text{ForAll } x \text{ hurtby}(x, D) \rightarrow \text{sueandwin}(x)$

In words: *If any person x is hurt by D , then x sues and will probably win.*

❑ Rule 2) $\text{ThereExists } x \text{ such that } \text{hurtby}(x, D) \ \& \ \text{sueandwin}(x) \rightarrow \text{ForAll } y \text{ not mayuse}(y, D)$

In words: *If at least one person x is hurt by D and sues successfully, then nobody is able to use device D . Admittedly, this is extreme, but in the United States, at least, we have almost arrived at that point.*

❑ Rule 3) $\text{ForAll } x \text{ sick}(x, S) \ \& \ \text{not mayuse}(x, D) \rightarrow \text{die}(x)$

In words: *If x has sickness S and x doesn't have device D , then x will die.* So, we're assuming the device D is important.

Warm-Up



Let's see the consequences of some starting assumptions. Suppose we assume that D is new and there is an axiom:

$$\text{ForAll } Y \text{ newdevice}(Y) \rightarrow \text{ThereExists } x \text{ hurtby}(x, Y)$$

What are the consequences?

Solution to Warm-Up

Given the axiom, the invention of D will inevitably lead to lawsuits and no savings in lives based on the following rules.

```
ThereExists x hurtby(x,D) → sueandwin(x)
ThereExists x such that hurtby(x,D) & sueandwin(x) →
ForAll y not mayuse(y,D)
ForAll x sick(x,S) & not mayuse(x,D) → die(x)
```

But this state of affairs is clearly undesirable. If D hurts very few people and saves many more, then D is overall good for society, at least in some utilitarian sense. Is there any way to take this into account in the lawsuits?

That is, suppose the judge at a trial could inform juries of the global cost-benefit history of the device. If the device has helped many more people than have been hurt, then the jury might take this into account. This could reduce the probability of success in lawsuits.

So, we enhance our logic with new predicates:

"MoreThanFraction(f) $x \dots$ " means that more than a fraction f , including x , of people have the property represented by the ellipses.

This then allows us to define a new rule 1 based on a fraction of, say, .02 (2.0%) for some device that is risky but overall very helpful:

□ Rule 1) $\text{MoreThanFraction}(0.02) \ x \text{ hurtby}(x, D) \rightarrow \text{sueandwin}(x)$

In words: *If more than 2% of the people treated by device D are hurt, then those people can sue and win. For someone to sue and win, then that person must have been hurt and more than 2% must have been hurt.*

Using this rule and assuming only one person in a thousand is hurt, you can see that there are no lawsuits.

This may be the right idea, but a few problems remain. First, even if the effective probability of being hurt is in fact one quarter the threshold (say 0.005 in this case), the early history of the device could exceed the threshold.

Part I: Mind Games

1. If each patient has a 0.005 probability of being hurt, how likely is it to get one failure within the first 50 patients?

Hint: You can solve this using probability theory or by writing a small program that simulates this probability.

But overall success might cause us to miss a real lemon. After all, failures aren't necessarily independent. In the hands of some hospitals, for example, device D might be far more dangerous than in others. Suppose only 10 people out of 2,000 have been hurt overall, but all 10 were hurt in the same hospital, which had treated only 100 patients. What should be done? This is where your advice as a mathematician comes in. But first, suppose that the threshold for successful lawsuits is reduced to 1% per hospital.

2. Suppose that 18 people out of 4,000 have been hurt by the device. There are eight hospitals, each of which has treated 500 patients. What is the likelihood that at least one hospital exceeds a 1% test, i.e., has 6 patients who were hurt, even if the underlying chance of failure were in fact independent of hospital and is 0.005 per patient?
3. Suppose the distribution of patients were much more skewed, with seven hospitals having treated only 200 patients each and the remaining one having treated all the rest. How likely is it that at least one of those smaller hospitals had hurt three people or more under the same conditions (i.e., 18 people who were hurt overall and the underlying chance of failure is 0.005 per patient)?

So far, these cases have made the lawsuits look frivolous. After all, even if the ground probability of being hurt is 0.005 and the hospitals are all equally careful, hospitals are quite likely to be vulnerable to lawsuits based on the threshold test. But suppose we fix this problem by increasing the threshold back to 2%. If the higher threshold test is met or exceeded in some hospital, then the likelihood that the 0.005 model is correct becomes so small that a statistically informed jury would infer that the hospital had not done its job.

4. Suppose we say that if a single hospital has more than 2% of bad outcomes, then that is a bad hospital. Assume that each hospital treats 500 patients. What is the likelihood for at least one hospital to hurt 10 or more people if the 0.005 hurt probability held and there were no bias in any hospital?
5. How would you answer the last question in the case that seven hospitals each treat only 200 patients and one treats 3,600?

The bottom line: If a hospital wants to offer the use of a new device, it should use it a lot.



The Box Chip Game



Sometimes a problem comes along for which a seemingly insignificant twist completely changes its character. In this puzzle, we design the odds for a new game similar to Reach for the Sky!

You have in front of you some even number n of chips each having a different color and you have also n identical, opaque latched boxes. As you watch, your adversary puts one chip in each box. He then puts all the boxes in a mixing bowl and tumbles them extensively. He then removes them and lays them out in front of you. The net effect is that you have no idea which box contains which color.

Here is what you bet on. For each color, you point to $n/2$ boxes where that color might be. You make all guesses about all colors before any box is opened. If you are correct about every color, you win. Otherwise you lose.

It might seem that your chances of winning are $(1/2)^n$, but you can do better.

Warm-Up



Suppose there are two chips — blue and red — and two boxes. So you must identify the box for red and the box for blue. Which guesses can you make so your chances of being right in both guesses is $1/2$?

Solution to Warm-Up

Guess red in, say, the left box and blue in the right box. You'll either be both correct or both wrong. You might as well be consistent about the two colors. It's no worse for you for both to be wrong than if one is wrong.

Let's see how that generalizes. Say there are four chips ($n=4$) whose colors are Black, White, Red, and Green. You are to guess two boxes for each chip. Number the boxes 1, 2, 3, and 4. Your guesses will be a simple list of box numbers for each color. For example,

Black guess: 1, 2

White guess: 2, 4

Red guess: 1, 3

Green guess: 3, 4

1. For four chips and two guesses per color, can you figure out a way to win at least $1/6$ of the time, assuming all rearrangements are equally likely? (The example just given might not achieve this.)

Part I: Mind Games

2. What if there are six chips and three guesses per color?
3. How does this probability generalize for n chips (n even) and $n/2$ guesses per color?

Now here comes the twist. It's small, but it changes everything. Under the old rules, for each color you provide $n/2$ boxes as guesses. Under the new rules, for each color you provide an initial guess and a procedure for determining each next guess based on the actual color found by the previous guess.

Here is an example for Black in the case of four colors:

```
Start with box 1;  
case result of first step is  
  Black, then win  
  Red, then try box 2  
  White, then try box 3  
  Green, then box 4
```

For convenience, we will abbreviate this as follows:

Black guess: 1; Black→win, Red→2, White→3, Green→4

You must provide all initial guesses and procedures before any checking occurs and the procedures may not share information. If every color is found in $n/2$ or fewer guesses, you win.

We might conceptualize the situation as follows: Each chip color is represented by an agent who follows a procedure like the one above, though different agents may follow different procedures. A given agent enters the room with the boxes, looks inside the boxes as dictated by his procedure, but doesn't otherwise disturb the boxes. He or she then leaves the room without being able to communicate with any other agent, unless he loses, in which case the game is over.

Continuing with this scenario: A judge determines whether that "procedural agent" has won or not. If all agents win, then you (the person betting) win. Otherwise, the house wins. It may seem surprising, but you can do a lot better under these new rules.

4. How well can you do in the case of four chips and two guesses under the procedural agent assumption? (*Hint: You can win over 40 percent of the time.*)
5. How well can you do if you have six chips and three guesses?



Feedback Dividends

Have you ever admired your own skill at navigating a narrow road at high speed? If not, imagine the following alternative method of travel: Study a detailed map of the same road, figure out how much the wheel should turn and the accelerator should be pressed at every time point, and then drive down the road blindfolded. Even without obstacles, this is beyond the memory and trigonometric capacity of most of us.

In fact, we're hardly conscious of the intellectual effort of driving. Perhaps the reason is that the act of driving consists of very short-term plans (few seconds at most) followed by adaptation based on eyesight. The driver has an overall goal—get to some destination—but the plan is incremental and adaptive. This requires less brainpower and is far more robust to changes in the environment.

Any person on the street understands this argument, but we require quantification. So to make this concrete, examine the following game. Consider a standard checkerboard having 8 rows and 8 columns as in Figure 1-9.

You want to go from row 1, column 4 (the black square above the S) to row 8, column 5 (the black square below the E). Each move goes from black square to black square and proceeds up a row and either to the left or right diagonally adjacent square. If you fall off the checkerboard or reach the top row without reaching the correct square, you lose.

At each move, you as player get to aim to go either right or left. You will achieve that step's aim with probability P_{good} , whose values we will discuss in a minute. There are two kinds of strategies: FeedYes and FeedNo.

A FeedYes strategy can decide where to aim on the i^{th} move after seeing the results of the first $i-1$ moves. A FeedNo strategy must decide where to aim at step i from the very beginning.




Figure 1-9: The goal is to go from the S to the E, always staying on black squares. Don't fall off the board.

Here is an example to show you the difference. Suppose you want to go from row 1, column 4 to row 3, column 4. Suppose that P_{good} is 0.9. Then in the FeedYes strategy you might aim right the first move. If you in fact go right (probability 0.9), then you would aim left the second move. But if you go left on the first move (probability 0.1), you will aim right the second move. The net result is that you have a probability of 0.9 to hit your destination. In the FeedNo strategy, you might say something like “aim right the first move and aim left the second.” There are two cases in which you would win with that strategy: You in fact move right in move 1 and left in move 2 (probability $0.9 \times 0.9 = 0.81$) or you move left in move 1 and right in move 2 (probability $0.1 \times 0.1 = 0.01$). So FeedNo has a probability of 0.82 of hitting the destination. Call the *feedback dividend* the probability of hitting the destination with the optimal FeedYes strategy divided by the probability of hitting it with the optimal FeedNo strategy. (Optimal means that you do as well as you can based on the probabilities.) In the example here, the feedback dividend is $0.9/0.81$.

Warm-Up



Are there any values of P_{good} for which the feedback dividend is 1 regardless of source and destination?

Solution to Warm-Up

If P_{good} were 0.5 or 1, the feedback dividend would be only one. In the first case, it doesn't matter where you aim. In the second, you don't need feedback. For all other P_{good} values, the dividend will exceed one.

Now here is the full problem. You start at row 1, column 4 and you want to hit row 8, column 5.

1. If P_{good} is 0.9, what is the probability of hitting row 8, column 5 using the FeedYes strategy and using the FeedNo strategy?
2. For which value of P_{good} does the feedback dividend reach its greatest value? What is the feedback dividend in that case?

You may be surprised by the result of this second question. It's not intuitive at all.

3. If we cut off the three rightmost columns and the two leftmost columns, then which value of P_{good} would give the highest feedback dividend? Assume that falling off the board gives a certain loss.



INFERENCE

What are you thinking?

Number Clues

Mind Games 

Refuse and Reveal

A Biting Maze

Mad Mix

Number Clues

A certain number theorist has recently disappeared. His safe may contain papers that would lead police to his whereabouts. The trouble is that his safe will self-destruct if it is forced open. So we must infer the combination.

He has left a few hints. They involve pairs of numbers. Each pair, p and q , has a product ($p \times q$), a greatest common divisor [$\gcd(p,q)$], and a least common multiple [$\text{lcm}(p,q)$].

Consider, for example, $p = 10$ and $q = 25$. The greatest common divisor $\gcd(p,q) = 5$, because both 10 and 25 are divisible by 5, but not by any number greater than 5. ("Divisible by" always implies a zero remainder.) The least common multiple $\text{lcm}(p,q) = 50$, because 50 is divisible by 10 and by 25, but no positive number less than 50 is divisible by both 10 and 25. So the greatest common divisor has all prime factors in common to p and q , in this case 5 alone. The least common multiple has the prime factors contained by either, in this case 2, 5, and 5.

The number theorist has left hints obeying the following rules:

- i. These hints concern pairs of positive whole numbers p and q . For every pair, $p \leq q$.
- ii. Each of p and q has two digits (i.e., between 10 and 99 inclusive).

Warm-Up

 Suppose I tell you that $q = 18$, p is unknown, but $\gcd(p,q) = 6$ and $\text{lcm}(p,q) = 36$. What is p ?

Solution to Warm-Up

The factors of 36 are 2, 2, 3, and 3. The factors of 18 are 2, 3, and 3. The factors of 6 are 2, 3. So, we know that p has at least 2 and 3, because $\gcd(p,q) = 6$ so 6 must divide p . Further, since $\text{lcm}(p,q) = 36$, either p or q has to have two 2s. Because q doesn't, p must, so p has factors 2, 2, and 3. Therefore $p = 12$.

Here are the first three hints, given stipulations (i) and (ii) above.

1. For the first p, q pair, $\text{lcm}(p, q) = 60$, $p \times q = 240$. What are p and q ?
2. The product of p and q is 140. $\text{gcd}(p, q) = 2$. What are p and q ?
3. There is a pair p and q whose $\text{lcm}(p, q) = 35$. What are p and q ?

The number theorist left another hint that the police have just decrypted: "If you've gotten this far, then you are doing well. Here now is one last hint."

4. The sequence of numbers that opens the safe is a sequence of five numbers from the hints. The sequence enjoys the following property: The greatest common divisors between neighbors of the sequence strictly increase as one proceeds from left to right in the sequence, but all greatest common divisors are single digits. What is the sequence that opens the safe?



Mind Games



You and I play a game that bears a strong family resemblance to the board game *Mastermind*. I think of a five binary digit secret number, e.g., 10101. You are allowed to pose questions about bit sequences of length five. Such a question is called a “bit question.” My responses will report how many bits in your guess have their correct value in the correct place.

For example, if your bit question is 10110 and my secret is 10101, then three of your bits (the first three, though I won’t tell you that) are in the correct place with their correct value. On the other hand, I could have other secret numbers that are consistent with the three-are-correct answer such as 00100.

Warm-Up



In fact, how many would be consistent with the answer of three correct for a bit question of 10110?

Solution to Warm-Up

Ten answers would be consistent. They would be: 10101, 10011, 10000, 11111, 11100, 11010, 00111, 00100, 00010, and 01110.

Warm-Up 2

 If you guessed 10110 and were told that all five were incorrect, then which possibilities would there be?

Solution to Warm-Up 2

Only one sequence is consistent with that answer: 01001. Just reverse all the bits.

Warm-Up 3

 If I give you the following answers, what is the secret number?

00000 — 2 correct

01000 — 1 correct, so second bit must be 0

00100 — 3 correct, so you know x01xx

01110 — 3 correct, so x011x

11100 — 1 correct, so 0011x

00001 — 3 correct, so last bit is 1 (compare with first response)

Solution to Warm-Up 3

Therefore, the secret number is 00111. These questions don't follow any particular systematic procedure, however. We should be able to do better.

1. What is the smallest number of bit questions sufficient to determine a five-bit sequence, no matter what it is? Remember that none of your bit questions need be correct. You just need to know how to crack the secret at the end.

Hint: A five question solution is not too hard to find, so try that first. You can do better, however.

2. Suppose we change the game to make it a little more difficult: You ask some number of questions and then I answer them all after you have asked all your questions. In that case, how many bit questions are sufficient?

Part I: Mind Games

I've become decidedly less friendly. You pose an initial bit question, to which I give no response at all — ever. As with question 2, I answer all your questions after you have asked the last one. After that, I give you only “low info” answers, consisting of one of three responses:

- (i) You have the same number correct as in the last bit question (S).
- (ii) You have more correct than in the last bit question (+).
- (iii) You have fewer correct than in the last bit question (-).

That's all I say. The term “correct” here still means correct bit in the correct position.

- 3. Given that the low info answers come at the end and there is no answer at all to the first bit question, is it possible to find the secret number, no matter what it is? If so, how many bit questions (counting the first one) do you need to pose in order to guarantee finding the secret number?
- 4. How many bit questions do you need, if you hear the low info answers immediately after posing them?
- 5. How does this generalize if the code is N digits long, but the digits could be base b for some $b > 2$? For example, b could be 10 and all digits 0 through 9 would be allowed. I have a solution when all answers come at the end, but that requires $1 + (b-1) \times N$ questions. I don't think it's optimal.



Refuse and Reveal

One of the approaches to understanding a large amount of data is to characterize it using a few numbers. Statistics such as minimum, maximum, and the various kinds of averages tell you global properties. Sometimes they are enough to reveal information about individuals. This is why databases that give only statistical information are an issue for privacy advocates: Enough statistical questions can reveal personal data.

Consider a simple game between a questioner Quentin and a responder Rosalba. Quentin can ask only about global properties of a group of numbers, e.g., are they all whole numbers, are they distinct, and the statistics — mean, median, minimum, and maximum.

Rosalba always tells the truth. She may refuse to answer, however, if she knows the answer may divulge all the numbers. As we'll see, the refusal to answer may itself divulge all the numbers. Sometimes, she will volunteer information just for the fun of it.

Warm Up

Rosalba: "I have five integers, all distinct."

Quentin: "What is the minimum?"

Rosalba: "15."

Quentin: "What is the maximum?"

Rosalba: "I won't tell, because you would know everything."



What are the numbers?

Solution to Warm-Up

Because the numbers are all distinct, the maximum would tell everything only if it were 19. Then the collection consists of 15, 16, 17, 18, and 19. Okay, this was easy, but the deductions will get more interesting.

Before we go on, though, let me remind you of the definition of the mean and median. The *mean* of a collection of numbers is their sum divided by the number in the collection. For example the mean of 20, 22, 22, 40, and 101 is $205/5 = 41$. The *median* is the middle number in the sorted order, so 22 for this example. That is, the median is the middle in a sorted ordering of the values (our examples will always have an odd number of values).

1. Rosalba: "I have five integers that may or may not be distinct."

Quentin: "What is the minimum?"

Rosalba: "20."

Quentin: "Which of these would not allow me to infer all their values: number distinct, mean, maximum, or median?"

Rosalba: "Only the median."

Quentin: "Great. I know the numbers."

What are they?

2. Rosalba: "I have seven integers that may or may not be distinct."

Quentin: "What is the minimum?"

Rosalba: "20."

Quentin: "Which of these are you willing to tell me (i.e., would not allow me to infer all their values): mean, median, and maximum?"

Rosalba: "All of them."

Quentin: "Ok, what is the maximum?"

Rosalba: "21."

Quentin: "I know which of mean and median you're willing to tell me now."

Which? Why?

3. Rosalba: "Can you find some situation in which I would be happier to tell you the mean rather than the median?"

Quentin: "Could you give me a hint?"

Rosalba: "In an example I can think of, there are three numbers, two of them distinct." Give it a try.

4. Rosalba: "Can you find some situation in which all of minimum, maximum, mean, and median are necessary and sufficient to find the identities of five numbers that are all integers?"

Rosalba: "So far we've been playing games with just a few numbers. I've given you hints and you've been able to infer them all. But just a handful of numbers are not interesting. Let's try for more.

"Before we do that, let me define one new global property: the *total distance* to a point. Let's say we have the five numbers 10, 15, 20, 30, 60. The total distance to a point — say, 22 — is the sum of (22-10), (22-15), (22-20), (30-22), and (60-22). Mathematically, the total distance to x from a set of numbers is the sum of the absolute values of the differences between each number and x .

5. Rosalba: "Now we are ready. There are 17 numbers that are not all distinct. Their minimum is 30, mean is 34, and median is 35."

Quentin: "What is their total distance to 35?"

Rosalba: "I won't tell you, but the total distance to 35 is five less than the total distance to 38. Whoops! I shouldn't have told you that."

Quentin laughing: "You're right. Now I know the numbers."

What are they?

6. How would your answer to this question change if there were 1701 numbers, but otherwise the same information as in the previous question?



A Biting Maze

A conventional maze is a set of walls or hedges enclosing roughly a rectangular plot of ground. One enters from the outside of the rectangle and wants to exit somewhere else to the outside of the rectangle. A well-known strategy to guarantee to exit such a maze is to keep your left hand on the wall and walk forward. Eventually you will reach the exit.

For example, consider the simple maze in Figure 1-10.




Figure 1-10: You are to travel this maze from entrance to exit by keeping your left hand on the wall and walking forward. Will you make it?

Keep your left hand on the wall at the entrance at the bottom and walk forward. This will lead you under the T shaped wall (though you will never touch it), up and above the entrance to the right, eventually to the bottom left corner (labeled A), then B, and then the dead end at C. But the left hand strategy will lead you out of this dead end until you will finally approach the exit from the left. This is certainly not the most efficient way to go through the maze, but it is guaranteed to work.

Warm-Up



 Can you see why this strategy will always work on any maze whose entrance and exit are gaps in the outside wall?

Solution to Warm-Up

Imagine that the exterior wall surfaces of the maze are painted black and the interior wall surfaces are painted white. Imagine that as you walk, you draw a line with your left hand using a crayon. Even though your path may not seem to be efficient, it ensures that you never go over the same section of white wall twice until you get to the exterior of the maze

Part I: Mind Games

(and hit a black surface). There is clearly only a finite amount of white wall surface between the entrance and the first exterior point of the maze. Therefore, you will eventually get out. (By symmetry, a right hand on the right wall will also work.)

This next puzzle, however, escapes the limitations of the physical world because the maze is on the Web. So there is no analogue to putting your left hand on the wall — even if you always go left, you could go in a circle.

The first Web page is at <http://cs.nyu.edu/cs/faculty/shasha/papers/mazebook.d/f43.html>. You want to get to a Web page that tells you (albeit in code) that you have arrived. There are hints along the way.

1. The challenge is to find a route to the final page, and to decrypt the words and phrases on the way. There are hints in the encrypted words (they form a sentence about natural history) and in other parts of the Web page. Give it a try.



Mad Mix

Imagine we have two large reservoirs of fluids. One contains the liquid 'supertox' and one contains purified water. We want to get certain proportional mixtures of the two.

There is also a drain for waste. Supertox is very dangerous, however, so we'd like to pour as little down the drain as possible, even in dilute form. Supertox is polar, so supertox and water mix perfectly and instantly whenever one is poured into the other.

In addition, we have two opaque measuring vessels, one of 10 liters and one of 7. We call these the 10-vessel and the 7-vessel for short. There are no measurement lines on the vessels. When we fill a vessel from a reservoir, we fill it from a faucet, as our client is nervous about contamination from handling.

So we can fill a vessel to completion from a reservoir or we can pour liquid from one vessel to the other, stopping either when the receiving vessel is full or the source vessel is empty. We are allowed to pour pure water back in the water reservoir and to pour pure supertox back into the supertox reservoir.

Warm-Up

 We want to create a 10-liter solution having 40% of supertox without pouring anything down the drain. How do we do it?

Solution to Warm-Up

Fill the 7-vessel with supertox. Pour into the 10-vessel. Fill the 7-vessel with supertox again. Pour into the 10-vessel until the 10-vessel is full. So we have 4 left in the 7-vessel. Pour the contents of the 10-vessel back into the supertox container. Pour the contents of the 7-vessel into the 10-vessel, and then fill the rest of the 10-vessel with water.

-
1. We want a mixture having 2.7 liters of supertox in a 3-liter mixture. We don't want to pour more than 10 liters of any liquid laced with supertox down the drain. How can we achieve this?
 2. Can we get a mixture that is $\frac{2}{9}$ supertox and the rest water without pouring any supertox down the drain?

Now, chemists require extremely precise mixtures, so let me explain our terminology with an example. A 25% concentration of supertox in the 10-liter container would correspond to 2.5 liters of supertox and 7.5 liters of water.

Also, we have a big empty bucket at our disposal that can hold over 100 liters, though we don't know exactly how big it is.

3. Can we get a 26% concentration of supertox in the 7-liter container?

4. Suppose that we want to get mixtures which are all supertox, $1/2$ supertox, $1/3$ supertox, $1/4$ supertox, ... $1/50$ supertox. You are allowed to specify the size of the two vessels. You may even include fractional sizes as long as the smallest vessel contains at least one liter and the largest contains at most 30 liters. In addition, you have the big bucket again with capacity over 100 liters. Which two vessel sizes would you choose in order to obtain each of these mixtures? Given those vessel sizes and given a target mixture, how would you obtain it?



OPTIMIZATION

Doing more with less

Dig That! 

Preferential Romance

No Change for the Holidays

Quiet in the Depths

Dig That!



Imagine a road grid with seven rows as in Figure 1-11.




Figure 1-11: Somewhere between Start and End in this road network is an eight-block tunnel some bad guys have built.

Some bad guys have placed a tunnel from the bottom, beginning at the intersection marked Start in the south, and ending at End in the north. The tunnel follows the path of the roads somehow but may wind around. It is also a simple path (no dead ends and no loops along the way). You want to probe a minimum number of times and yet be able to find the exact route of the tunnel.

A probe device takes an hour to set up. We can set them up at intersections or in the middle of the street. If you establish a probe on a street, you can tell whether the tunnel follows the street. If you set the probe up at an intersection, you can tell whether the tunnel passes through the intersection, and, if so, according to the engineering team, which adjacent streets it goes to.

1. If the tunnel is at most eight blocks long and begins at Start and ends at End, then what is the minimum number of probe devices you would need to guarantee to determine the precise route of the tunnel in one hour?
2. If you had only one probe, what is the minimum time it would take to guarantee to determine the outcome, assuming you could move the probe every hour?

Suppose all a probe could tell was whether a tunnel ran under it or not, but could not tell the direction. In the engineering jargon, suppose we had point probes rather than directional probes.

3. If the tunnel is at most eight blocks long and begins at Start and ends at End, then what is the minimum number of point probe devices you would need to place to guarantee to determine the precise route of the tunnel in one hour?

Hint: Remember that you can place probes at intersections or in the middle of streets.

Suppose you have two hours, but you are not allowed to reuse any probe?

4. How few point probes could you use over two hours to guarantee to find the tunnel?



Preferential Romance

A certain marriage counseling service called Marriage Success, or MS for short, advises couples on how to get along better. MS's idea is simple: Each spouse writes down his/her preferences about various criteria of common interest. "Our criteria go beyond those elements of physical appearance and passion that guide early romance and tend to blind judgment. We want to understand your values as you live day to day. The happy couple is the one whose preferences are compatible or can be made compatible."

Here are some of the positive qualities each person might wish a spouse to have: biker (B), cultured (C), enthusiastic (E), foodie (F), hiker (H), juggler (J), kayaker (K), movies (M), organized (O), puzzles (P), rich (R), theatre (T), and windsurfer (W).

Suppose X and Y are qualities of people. $X \rightarrow Y$ implies a preference for X to Y . Preferences are transitive, so $X \rightarrow Y$ and $Y \rightarrow Z$ implies $X \rightarrow Z$. Two people are *fully compatible* if their preferences are consistent. The test for consistency is that there is some list of qualities that reflects both of their preferences. If two people aren't fully compatible, then perhaps at least they are *passably compatible*, which means they would be fully compatible if each spouse dropped at most one preference.

Warm-Up

Suppose Bob's preferences are

$$K \rightarrow M \rightarrow P, \text{ and } R \rightarrow T$$

and Alice's preferences are

$$O \rightarrow P \rightarrow R \text{ and } W \rightarrow T.$$

 Then would they be fully compatible?

Solution to Warm-Up

Yes. Ignoring the qualities not mentioned (because they can be put anywhere), here is a consistent ordering for these preferences: KOMPRWT. This is consistent in the sense that for every preference $X \rightarrow Y$, X comes before Y in the list. On the other hand, if Alice added the preference $R \rightarrow M$, then the couple would not be compatible because $P \rightarrow R$ and $R \rightarrow M$ implies $P \rightarrow M$, but Bob has $M \rightarrow P$. If Bob drops his $M \rightarrow P$ preference, then we would be left with:

$$\text{Bob: } K \rightarrow P, \text{ and } R \rightarrow T$$

$$\text{Alice: } O \rightarrow P \rightarrow R \rightarrow M \text{ and } W \rightarrow T.$$

This could yield the following consistent ordering, among others: KOPRMWT. So Bob and Alice would then be passably compatible.

MS requires that each person be consistent in that he or she has no cyclic preferences of the form X is better than Y which is better than Z which is better than X. The bad news is that the real Bob and Alice are both very opinionated, so they have each listed many preferences, so many, in fact, that I have drawn Figures 1-12 and 1-13 to illustrate.




Figure 1-12: Bob's preferences in life. Note that he has no preferences for certain qualities, e.g., he has not said whether he prefers Alice to be cultured or to be a windsurfer.





Figure 1-13: Alice's preferences in life. Note that she has no preference about whether Bob is rich or a biker.

- 1. Are Bob and Alice compatible, passably compatible, or neither? After dropping the smallest number of edges in zero, one, or both spouse's preferences, try to find a consistent ordering.
- 2. [For the entrepreneurial at heart] Can you describe an algorithm for the counseling company to use to help marriages in distress? That is, try to find a method so spouses have to drop as few preferences as possible.



No Change for the Holidays



Have you ever noticed that if a teenager is given \$50 to buy a \$20 item, no money comes back? Some other item was just *so* essential...

In this puzzle, we provide the teenager with a good excuse. *Artiste* and street vendor Claude sells beautiful handmade items for \$100 or less but refuses to give change.

You have no cash, but you have three checks. You will make them out to Claude in whole dollar amounts. Your teenager is to give Claude the combination of checks that is the minimum amount more than necessary for the purchase price. For example, if you give your child a \$50 check, a \$30 check, and a \$20 check and the item costs \$53, your child will give Claude the \$50 and \$20 checks and Claude will keep the \$17 in change in addition to the purchase price.

You like Claude's stuff, but you very much begrudge him his "I keep the change" attitude. So you'd like to minimize the amount he keeps beyond the purchase price.

Warm-Up



If you knew that the item in question cost either \$20, \$40, \$50, or \$60, which combination of three check amounts could you give in order to leave no change for Claude?

Solution to Warm-Up

\$20, \$40, and \$50 is one among many solutions. Claude will not be able to keep anything in change.

1. If you did not know how much the purchase price is except that it is a whole number amount between \$1 and \$100 inclusive, which amounts would you write on your three checks in order to minimize Claude's change?
2. Suppose Claude publishes his four whole-number prices in an advertisement that you see. Can you show how he can guarantee to do so in such a way that at least one item will yield him non-zero change no matter which check amounts you write?



Quiet in the Depths

The need for near-silence in a submarine leads to strict procedures. As a submarine captain explained, “One of our drills, in case we are under conditions of extreme radio silence inside the submarine, is to exchange information among our five key departments. Naturally, we can’t tell you which those are.

“Each department consists of five sailors. Our basic protocol works as follows: As captain, I request certain information by broadcasting the request to all departments; each department responds with a number that is sent only to the captain (no broadcast); then the captain may broadcast a number; and so on. Remember, the captain broadcasts; the others don’t.

“To demonstrate, the departments will exchange very few messages with me — as few as possible — in order for me to determine some information. Naturally, we can’t reveal to you the nature of such information, but my first officer has come up with a sanitized version.”

Warm-Up

 Suppose the captain broadcasts the request for the arithmetic mean of the blood pressures of all sailors in those departments. (“Blood pressure,” technically, systolic blood pressure, is the sanitization of a three digit quantity of interest.) Can the departments respond in such a way that the captain can then compute the mean?

Solution to Warm-Up

Each department calculates the mean of its sailors’ blood pressures and reports that number. The captain computes the sum of these means and divides by five. This works because each department has the same number of sailors (otherwise the captain would have had to weight each reported mean by the number of sailors in that department).

Okay, that was pretty easy. It turns out they weren’t so interested in the mean as in the median. (The mean is too sensitive to outliers.) But the median itself wasn’t necessary, simply a “blood pressure” among the 10 middle values of the 25.

1. Assuming blood pressures are three digits long, can each department report one three digit number so that the captain can infer a blood pressure among the 10 middle blood pressures of the 25 sailors?
2. The captain wants to find some blood pressure value that is sure to be less than or equal to the median and some blood pressure value that is sure to be greater than or equal to the median. Using the assumptions of the first question, can the captain infer such a value after one three-digit report from each department?



Solutions

Solution to Sweet Tooth

1. How does Jeremy maximize the amount of cake he gets, given these rules? How much will he get?

Jeremy cuts the first cake in the fractions f and $1-f$, where f is at least $1/2$. Then if Marie chooses first, Jeremy will get $1\frac{1}{4}$ cakes from the last two by the argument of the warm-up. If Marie chooses second, Jeremy will get f and then will divide the last two cakes in half. So, $f + 1/2 + 1/2 = (1 - f) + 1\frac{1}{4}$. That is, $f + 1 = 2\frac{1}{4} - f$. $2f = 1\frac{1}{4}$. That is $f = 5/8$. Therefore, Jeremy will receive $1\frac{5}{8}$ cakes and Marie will receive $1\frac{3}{8}$ cakes.

2. Suppose there were seven cakes and Marie got the first choice for six of the seven cakes. Who has the advantage? By how much?

There is an inductive pattern. If for k cakes where Marie gets the first choice in all but one, Jeremy still has an advantage A (that is, Jeremy would receive A more cake than Marie among those k), then Jeremy will still have an advantage for $k+1$ cakes. Here is why: Jeremy, knowing that he will have an advantage of A if Marie chooses first for the first cake, cuts the first piece into two pieces $1/2 + A/4$ and $1/2 - A/4$. If Marie chooses the piece having $1/2 + A/4$, then Jeremy will suffer a disadvantage of $A/2$ for that first cake, but will gain an advantage of A for the remaining k cakes, so have a net advantage of $A/2$. By contrast, if Marie chooses to go second for the first cake, then Jeremy will cut the remaining cakes in half. Again he will have an overall advantage of $A/2$. Thus, Jeremy's advantages from $1/2$ for two cakes, to $1/4$ for three cakes, to $1/8$ for four cakes, to ... $1/64$ for seven cakes.

3. Is there any way to ensure that each child receives the same amount of cake, assuming Jeremy always cuts?

Yes, simply make it so Marie gets the first choice every time. That will force Jeremy to cut evenly every time.

Solution to Byzantine Bettors

1. Now suppose there are only three advisors and only one of them always tells the truth. Again, you can make three even money bets. You start with \$100. How much can you guarantee to win?

If the vote is two against one, then bet x on the majority. If you lose, the advisor in the minority must be the truth-teller so you can bet everything on the next two bets by following the advice of the truth-teller. This gives you $4 \times (100 - x)$. If you win by following the majority the first time, you should disregard the advisor in the minority and you'll have $100+x$ after the first bet. In the next bet, if the remaining two agree, you bet everything because one must always tell the truth. But if not, you'll find out who is telling the truth by betting nothing. On the last bet, you follow the truth teller and end up with at least $2 \times (100 + x)$.

Setting $4 \times (100 - x) = 2 \times (100 + x)$, you conclude that $6x = 200$ or $x = 33 \frac{1}{3}$. Plugging that into either equation corresponds to handling both cases when the vote is two against one. Thus, in either case, you end up with $2 \frac{2}{3}$ of what you start with—a total of \$266.66.

2. What can you guarantee to win in four bets, with four advisors, three of whom can lie at will, and one of whom must tell the truth at least three out of four times?

In this case, unfortunately, you cannot guarantee anything. The reason is that your “advisors” can arrange it so that you never have a reason to choose one outcome over the other. Call the advisors A, B, C, and D.

Bet 1: AB advise 0; CD advise 1.

Outcome: 1, so A and B have each lied once.

Bet 2: AC advise 0; BD advise 1.

Outcome: 1, so A has lied twice so is a definite liar.

B and C have each lied once.

D has not yet lied.

Bet 3: BD advise 0; AC advise 1.

Outcome: 1, so B is a definite liar. C and D have each lied once.

Bet 4: AC advise 0 and BD advise 1.

Notice that, in bet 2, if you had bet on 1 but the outcome had been 0 (because the advisors had changed the written result), then B would be out, A and C would have lied once, and C would have never lied. This is symmetrical to what happened with just a rearrangement of labels. So, you would have learned no more if the outcome had been 0. Similarly, in the third bet, if the outcome were 0, then D and B would both be possible partial truth tellers in bet 4. Thus, no series of outcomes ever teaches you whom to trust. The advisors can arrange it so that every bet you make is a loser.

3. Under the reliability conditions of one partial truth teller who tells the truth four out of five times and three liars at will, but assuming you have five bets, can you guarantee to end with at least \$150?

Even with five bets, the advisors can keep you from winning much. Their strategy starts similarly to their strategy when you had only four bets.

Bet 1: AB advise 0; CD advise 1.

Outcome: 1, so A and B have each lied once.

Bet 2: AC advise 0; BD advise 1.

Outcome: 1, so A has lied twice so is a definite liar.

B and C have each lied once.

D has not yet lied.

Bet 3: BC advise 0; AD advise 1.

Part I: Mind Games

Now, we must do a case analysis.

Case 1: Suppose you bet nothing on bet 3. Suppose the outcome is 1. Now you know that D is the partial truth teller because everyone else has lied twice. However, D may still lie once. If D advises 1 in bet 4 and you bet x on 1, then if $x < 50$ and D tells the truth in bet 4, you cannot bet anything on bet 5 because D could lie then. If $x \geq 50$, then D could lie in bet 4 and you will never do better than regain your original capital. So, if you bet nothing, you cannot attain \$150.

Case 2: Suppose you bet something on 0 in bet 3. Then your situation is as in case 1, only worse.

Case 3: Suppose you bet x on 1 in bet 3. If $x < \$12.50$, then suppose further the outcome is in fact 1 and you bet y in bet 4 on the outcome offered by D and $x + y < 50$. Then, if you win in bet 4, you cannot bet in bet 5 (because D may lie). If $x + y > 50$, then $y > \$37.50$, so if you lose bet 4, you will have less than \$75, so cannot attain \$150 in bet 5. Therefore $x \geq \$12.50$.

Now suppose that you bet x on 1 in bet 3, but the outcome is 0. Then you have $\$100 - x$ and you know that B, C, and D have all lied once. In bet 4, two of them must advise one outcome and one must advise the other. Let's say you bet z on the outcome suggested by the majority. Then $x + z \leq 25$; otherwise losing in bets 3 and 4 will leave you less than \$75, making it impossible to attain \$150 in bet 5. But this implies $z \leq \$12.50$ and therefore if you win in bet 4, you will have at most \$100 and D can still lie in bet 5. If you bet with the minority in bet 4, the majority could be telling the truth and then there would remain two advisors who have lied only once.

Therefore, you cannot guarantee to attain \$150. In fact, a refinement of this argument shows you cannot even attain \$134. Ivan Rezanka has argued that \$133.33 can be guaranteed, but a short elegant proof is still missing.

Solution to A Touch of Luck

- 1. Bob, Carol, and Alice play. Alice has 51 units, whereas Bob and Carol have 50 each. Bob must state his bet first, then Carol, and then Alice. Bob and Carol conclude to share the reward if either wins. How can Bob and Carol maximize the probability that at least one of them wins if there is just one coin flip left?**

Bob bets 50 on heads and Carol bets 50 on tails. If Alice bets 49 or less, then she can't possibly win because even if she guesses the correct coin orientation, her 100 will equal the 100 of one of Bob or Carol. If Alice bets 50 or more on heads, then she wins on heads and Carol wins on tails. So, Bob and Carol can guarantee to have a $1/2$ chance to win.

- 2. Does the result change if Alice must state her bet first?**

If Alice bets 48 or less, then Bob and Carol bet as above. One of them will surely win. But if Alice bets at least 50, Bob and Carol can do no better than before. So, the result stays the same.

- 3. Suppose Bob has 51 units and Alice 50. There are two coin flips left. Bob bets first for the penultimate flip. Alice bets first for the last flip. Does Bob have more than a $1/2$ chance of winning? If so, how much more?**

Bob will win with probability $3/4$. In the penultimate round, he bets nothing. No matter what Alice bets, she cannot exceed 100 units. If Alice loses in the penultimate round, then since she must state her bet first in the last round, Bob just bets as she does in that round and is guaranteed to win. If Alice wins in the penultimate round, then, in the last round, Bob simply bets all his units on the coin face opposite to the one Alice chooses. (If Alice doesn't bet in the last round, Bob can bet everything on, say, heads.) So Bob will win $3/4$ of the time.

- 4. Suppose Bob has 51 units and Alice 50. Again, there are two coin flips left. This time, Alice bets first for the penultimate flip. Bob bets first for the last flip. Does Bob have more than a $1/2$ chance of winning? If so, how much more?**

Suppose Alice bets 2 on heads in the penultimate round. If Bob bets nothing or bets on tails and the penultimate flip lands on heads, then Alice wins for sure by using the strategy from the warm-up. If Bob bets all 51 on heads and the result is tails, then again Alice is sure to win on the last flip. If Bob bets anything on heads and the result is heads, then Bob has at most 102 and Alice has 52. In that case, on the last flip, Alice bets everything on the opposite face to Bob's and she will win $1/2$ the time regardless of what Bob bets. So, Alice wins at least $1/2$ the time and will win more often if Bob plays foolishly. A good move for Bob would be simply to bet the same amount as Alice on the same face (so 2 on heads if Alice bets 2 on heads) in the penultimate flip. He can then bet nothing on the last flip and be sure to win with probability $1/2$.

- 5. Suppose Bob has 51 units and Alice 50. Again, there are two coin flips left. Again, Alice states her bet first in the penultimate round and Bob states his bet first in the final one. This time, Bob announces that he will bet 20 in the penultimate round, though he will wait to see Alice's bet before saying whether he will bet on heads or tails. Can Alice arrange to have more than a $1/2$ chance to win?**

Yes. Alice bets nothing in the penultimate round. If Bob loses, Alice wins for sure in the final round using the strategy of the second warm-up. If Bob wins, Alice bets everything on the coin face opposite to the one Bob chooses in the final round. She will win with probability $3/4$.

- 6. Bob, Alice, Rino, and Juliana have 100 units each and there are two flips left. Each person is out to win — no coalitions. Bob and Alice have bet 100 on heads. Rino has bet 100 on tails. Juliana must now state her bet, knowing that she will state her bet first on the last flip. What are her chances of winning if she bets 90 on either heads or tails? What should she bet?**

If Juliana bets 90, she has no hope of winning. If she bets 90 on heads and the outcome is heads, then Bob and Alice will each have 200 compared to her 180 after the first flip. If she then bets x on, say heads, on the last flip, she knows that (i) Bob and Alice must split their bets on that last flip (because ties do them no good) and (ii) one of them can arrange to bet x or more on heads as well. Covered in this way, she has no chance of winning. If Juliana bets 90 on tails in the first flip and the outcome is tails, then Rino will cover her on the last flip.

Part I: Mind Games

So, Juliana should bet 100 on tails in the penultimate flip. If the outcome is tails, then she and Rino will each have a $1/2$ chance to win in the last flip. Rino won't bet as she does, again because ties do him no good. So, overall Juliana can arrange to have a $1/4$ chance of winning, but she must bet high.

John Trono and Tom Rokicki lent their always-helpful insight to these solutions.

***Note:** All questions relating to this Touch of Luck puzzle mentioned so far have paper-and-pencil solutions. But the general situation having multiple players, multiple rounds, and different initial wealths is still elusive. The person who can solve it would have to mix, I think, expertise in game theory with algorithmic insight.*

Solution to Information Gain

- 1. Is it possible for Jordan to design a protocol to ensure that each of his friends will raise the correct number of fingers? If so, explain it. If not, does Jordan have a high probability of winning?**

Jordan's team can win for sure. This time, by prior agreement, all the mathematicians will know: Ariana represents 0 if she receives a blue ticket and 5 if she receives red; Bob represents 1 for blue and 6 for red; Caroline represents 2 for blue and 7 for red; David represents 3 for blue and 8 for red; and Ellen represents 4 for blue and 9 for red.

Jordan adds up all the numbers on all the hats. He divides them by 10 and takes the remainder. For example, suppose the numbers are 3, 2, 7, 9, and 5. He adds them up and gets 26. The remainder after dividing by 10 is 6.

He sends a ticket reflecting the remainder value to the appropriate person. In this example he would send a red ticket to Bob to represent 6.

Now each mathematician adds up the other numbers and then derives his or her own number as the unique value to make the remainder work out.

In this example, Caroline would have a 7 on her hat. She would see 3, 2, 9, and 5. Adding those up, she would get 19. She knows the final remainder has to be 6, so she would hold up 7 fingers.

Solution to Reach for the Sky!

- 1. What would be a good strategy for the \$100,000 Reach for the Sky! game, where you have three captures? Using that strategy, what would be your chance of winning?**

Following the idea of the sultan's daughters problem, we will express our protocols in terms of three numbers in increasing order: n_1 , n_2 , and n_3 . The idea is to reject the paper values in the first n_1 balls, then choose the first ball whose paper value is greater than the highest of those in the first n_1 balls. Call the ball number of the first capture b_1 (so $b_1 > n_1$).

If $b_1 > n_2$, then make the second capture at a ball whose paper value is a greater number than any paper value seen in the first b_1 balls. Otherwise (when $b_1 \leq n_2$), reject all paper values until the first n_2 balls (altogether) have been seen and then make the second capture occur at a ball whose paper value is greater than any seen so far in the first n_2 . Call the position of the second capture b_2 .

If $b_2 > n_3$, then choose, for the third capture, the paper value greater than any seen before position b_2 . Otherwise, reject until position n_3 and then pick the next one better than any seen before the first n_3 . Call the position of the third capture b_3 .

Of course, if the highest value happens to be in the first n_1 balls, then you will never choose anything, but it wouldn't help to do so anyway. That suggests that it is useful to make n_1 rather small.

For 100 balls in total and 3 captures, you win roughly 68 percent of the time if you set n_1 , n_2 , and n_3 to be 14, 32, and 64, respectively.

Consider, for example, the following sequence of 100 values where each dollar amount written by your adversary has been replaced by its rank. That is, the highest dollar figure is replaced by 99 and the lowest by 0.

```
78 3 80 90 25 95 51 27 57 40 65 48 55 72 26 73 54 31 15 2 89 61 97 98 8
50 38 18 88 52 4 42 68 16 62 9 94 99 20 28 56 58 76 93 10 96 63 35 81 91
66 11 30 5 0 24 82 29 41 12 47 71 44 92 43 32 85 84 7 59 60 86 69 21 83
79 64 67 74 37 1 46 22 19 33 39 87 45 36 13 23 75 34 70 53 49 77 17 6 14
```

The b_1 position value in this case would be 23 where the value 97 is found, because 97 is the first value larger than 95, which is the largest of the first 14 numbers. The b_2 value would be 38 where 99 is found and the third capture then is irrelevant.

The bottom line: Take the bet.

There may be an elegant analytical way to solve this problem as there is for the original sultan's daughters problem, but the puzzle is so amenable to programming, it didn't seem worth it. Simply choose ascending triples and then, for each triple, generate 10,000 random permutations of the numbers between 0 and 99 (the ranks) and see how often you win. Such a strategy is not limited to triples, obviously.

2. How would your answer change if there were 1,000 lottery balls?

For 1,000 balls with three captures, you win roughly 66 percent of the time if you set n_1 , n_2 , and n_3 to be 141, 317, and 650 respectively. So, you should still take the bet. Isn't it amazing that the odds stay so good?

Solution to Pork Politics

1. Which coalitions might form?

A and C will form a coalition together. C prefers A to any other coalition, because C will get the most money in that case. A also prefers C.

2. How strong a majority should B prefer (51%, 67%, or 75%) in order for B to receive as much money as possible?

67% would be best for B. Then B would be part of the A, B, E coalition and B would get the fraction 25/70 of the money. A also would get more with this coalition than with any other winning coalition of which A would be a part. If the cutoff were 75%, A would team up with C, D, and E. B would get nothing.

Solution to Social Games

1. How would punishment have to change to eliminate cheats-cheats as a Nash equilibrium?

If punishment were to increase to -9, then the upper right corner would give Alice $(4 \times 0.9) + ((-9) \times 0.1)$. The bottom right corner would give an expected gain of 0. That yields the matrix:

	Alice Honest	Alice Cheats
Bob Honest	3, 3	0, 2.7
Bob Cheats	2.7, 0	0, 0

In this case, if Alice cheats, Bob receives the same gain whether he cheats or not. So, the lower right and the upper right are the same from Bob’s point of view. If Bob decides not to cheat, then the game state will move to the upper right and from there to the upper left. To encourage this, you could make the punishment even slightly more negative than -9 (e.g., -9.01).

2. If you are a public policy maker and you want to limit the punishment for cheating to -5, how much must you increase the probability of catching criminals to maintain the upper left as the only Nash equilibrium?

Suppose punishment is reduced to -5 and the probability of punishment is raised to 17 percent. Then the expected gain for Alice from cheating is $(0.83 \times 1) + (0.17 \times -5) = -0.02$ in the bottom right state. In the upper right corner, the expected gain will be $(0.83 \times 4) + (0.17 \times -5) = 2.47$.

	Alice Honest	Alice Cheats
Bob Honest	3, 3	0, 2.47
Bob Cheats	2.47, 0	-0.02, -0.02

So the only Nash equilibrium is the upper left.

3. Assuming a 10 percent likelihood of being caught, is there a punishment value that would cause the only Nash equilibrium to be the upper right state (Alice commits crime but Bob does not)?

Under the assumption of a 10 percent likelihood of being caught, let's say the punishment is -9.01. Then the lower right has an effective value of $(-9.01 \times 0.1) + (1 \times 0.9) = -0.001$. The upper right has (for Alice) a value of $(-9.01 \times 0.1) + (4 \times 0.9) = 2.699$, which we'll approximate to 2.7.

This yields a game matrix of:

	Alice Honest	Alice Thief
Bob Honest	5, 2	0, 2.7
Bob Thief	2.7, 0	-0.001, -0.001

So, Alice will take the matrix to the upper right corner. Bob can rail about immorality, but Alice sees it all as a game.

4. Still assuming a 10 percent likelihood of being caught, what must the punishment be to force Alice to stay honest?

Recall the game matrix relevant to this question:

	Alice Honest	Alice Thief
Bob Honest	5, 2	0, 4
Bob Thief	4, 0	1, 1

If punishment increases to, say, -17 (which might represent cutting off the hand of a thief who steals bread), then the game matrix becomes:

	Alice Honest	Alice Thief
Bob Honest	5, 2	0, 1.9
Bob Thief	1.9, 0	-0.8, -0.8

The only Nash equilibrium is the upper left.

Solution to Escape Management

1. The police want to control T's direction at only five intersections. They know that T is extremely motivated to escape. How long (in terms of single block steps) at most will it take T to arrive at either the northern or southern border?

Figure 1-14 shows most of the northern half. The situation is symmetric for the south. The police are represented by dashed arrows and the thief by solid arrows. The earliest T can escape is after traveling 22 blocks (which we call steps). This is the best possible for P

Part I: Mind Games

because each southerly step that P forces requires for compensation at most one east-or-west plus one northerly step by T. Here is the calculation. T needs nine steps even without obstruction. Add to that one by P to change directions to east-or-west plus four southerly steps forced by P, each of which requires two compensatory steps by T. So the best possible strategy for P is $9 + 1 + (4 \times 3) = 22$.




Figure 1-14: The northern half of the city T wants to escape from. The police are represented by dashed arrows and the thief by solid arrows. T leaves after 22 block steps.

- 2. Suppose that the police want to prevent T from reaching the northern or southern borders. Therefore the police plan to control T's direction m out of every n times for some $m < n$. What should m and n be to minimize the fraction of time P forces T to change directions while still preventing T from escaping?

The police can use a strategy more or less in the spirit of the solution to question 1. The idea is to let the thief go eight steps (to one step away from the northern border) and then to force the thief to go east (or west) and then south just as in the figure above. The thief then moves east or west and then the police force a move south again. The police eventually force the thief all the way down to row 1, but facing away from the northern border. This requires eight steps by the police and 15 steps by the thief. So the police intervene 8 times out of every 23.

3. How does your answer change if T can escape by reaching any border, including the eastern and western ones?

If the police intervene two times out of every four, then they can force the thief to stay in a six by six square around the center. Each time the thief goes, say, east and then north, the police force the thief to go west and then south. If the thief goes in a single direction, then the police force the thief to stay within the central six by six square. This plan clearly lacks subtlety, but it works. I see no better general solution, but maybe a clever reader will find one. If so, please let me know.

These solutions are due to Ivan Rezanka.

Solution to Flu Math

1. As long as the government is in a compelling mode, what fraction should be required to take the vaccine in order to minimize the overall average risk?

If a fraction f take the flu shot, then the death toll is $5f + 10(1-f)(1-f) = 5f + 10f^2 - 20f + 10$. That is, $10f^2 - 15f + 10$. This is minimized when $20f = 15$ or $f = 0.75$. At that point, we'd have a death toll of only $(5 \times 0.75) + (2.5 \times 0.25) = 4.375\%$. Notice, however, that those who don't take the vaccine have a probability of only 2.5% of dying whereas the ones who do take the flu shot have a 5% chance of dying.

2. Under these conditions (no compulsion, full belief in government risk figures, selfishness), what percentage of people will take the flu shot?

Exactly half (or maybe one less than half, given the bias towards inaction). Here is why. Take the i th person. If half have already taken the flu shot, then there is no reason to take it. Person i will die with a 5% probability if he or she takes the shot and will die with a 5% probability or less if he or she doesn't take the shot. If less than half have taken it and there are just enough people left for half to take it, then person i will take it for sure, because if he or she doesn't, the chances of living will be less than if he or she took the shot. We call this the "must-vaccinate" point.

3. Suppose 25% of the population will not take the vaccine under any circumstances. Nobody else knows beforehand who they are. Then can you determine what percentage of people will in fact take the vaccine under the same assumptions?

In this case, at most 50% will take the vaccine. Once 50% have taken the vaccine, nobody (neither a refuser nor a normal citizen) will take it. If fewer than 50% have taken the vaccine and we are at the must-vaccinate point, then a non-refuser will take the vaccine, but a refuser of course will not.

4. Again, the government will compel nobody. To which percentage should the risk of disease be inflated to achieve the vaccination level you determined in your answer to question 1?

The government wants $f = 0.75$ from the answer to 1. So if the initial perceived risk is R , each person will take the vaccine until $R(1-f) = R \cdot 0.25 = 5\%$. So $R = 20\%$. That is, if everyone believes the initial unvaccinated risk is 20%, then 75% of the people will take the vaccine.

Solution to Whipping Ice

1. Does Indy need more poles than the two shown by the dark circles in Figure 1-4 to go from the bottom to the top?

This solution requires two connections and disconnections to the same pole, as shown in Figure 1-15.




Figure 1-15: Indy uses the first pole both to straighten out and then eventually to curve in.

2. There are several L-shaped barriers and several objects (at the Xs) that Indy wants to collect in a game called Treasure Skater. What is the minimum number of poles he will need and where should he place them?

Four poles are enough. Figure 1-16 shows the basic solution. Notice that poles can be used many times.




Figure 1-16: This is the route Indy should take starting from the bottom to win at Treasure Skater. The numbers indicate the order in which Indy visits the Xs. Notice that he does a sharp turn using the lower left pole and then uses that pole again to curve up to X1. He then uses the lower left pole and the lower right pole twice to swing to 2 and 3.

The tricky part is near the X labeled 4 (see Figure 1-17).




Figure 1-17: The difficult part of the solution is to change directions as Indy descends between X2 and X4. He has to use a distant pole.

These solutions are due to TJ Takei.

Solution to Optimal Jargon

- 1. Suppose that in addition to the frequencies given in the warm-up, you know that in every message, B is always followed by A, and D is always followed by C. Could you arrive at a shorter encoding of the message?

In that case, we can form a code like this:

A-1
BA-01
C-000
DC-001

There will be 15 As by themselves with a cost of 15. There will be 15 BAs with a cost of 30. There will be 5 DCs with a cost of 15. There will be 5 Cs with a cost of 15.

So the total cost is only 75.

- 2. Suppose that in addition to the frequencies given in the warm-up, you know that in every message, D is always followed by C, C is always followed by B, and B is always followed by A. Could you arrive at an even shorter encoding of a 60-character message? Which phrases should be rendered as jargon?

This enables us to form the encoding:

```
A-1
BA-01
CBA-000
DCBA-001
```

In this case, there are 5 DCBAs with a cost of 15; 5 additional CBAs (i.e., not included in the DCBAs) with a cost of 15; 5 additional BAs with a cost of 10; and 15 additional As with a cost of 15. This gives a total of only 55 bits, less than one bit per character. The jargon terms should correspond to the phrases DCBA, CBA, BA, and A.

- 3. Suppose that in addition to the frequencies given in the warm-up, you know that in every message, D is always followed by C, C is always followed by B, and B is followed by A 13 out of 15 times? What is the shortest encoding you could find of a 60-character message?**

We'll try two encodings. The first encoding includes the same phrases as in the solution to question 2, and then adds long codewords for the exceptional cases.

```
A-1
BA-01
CBA-000
DCBA-00101
DCB-00100
CB-00110
B-00111
```

For this first encoding, let's consider two cases. In the first case, assume that there will be 5 DCBAs with a cost of 25; 5 additional CBAs (i.e., not included in the DCBAs) with a cost of 15; at least three additional BAs with a cost of 6; two Bs with a cost of 10; and 15 additional As with a cost of 15. In the second case, assume instead that, say, two of the CBAs were CBs instead, but that there were five BAs. The cost of the phrases beginning with C would then increase from 15 to 19, but the cost of the phrases beginning with B would decrease from 16 to 10. If two of the DCBAs were CDBs instead, then the cost of phrases beginning with D would not increase at all. So the worst kind of message from the point of this encoding is that two of the phrases that begin with B have no following A. The total cost of such a message is $25 + 15 + 6 + 10 + 15 = 71$.

The second encoding uses only jargon terms having no exceptions:

```
A-1
B-01
CB-000
DCB-001
```

There will 5 DCBs with a cost of 15; 5 additional CBs (i.e., not included in the DCBs) with a cost of 15; five additional Bs with a cost of 10; and 30 As with a cost of 30. This gives a total of $15 + 15 + 10 + 30 = 70$. Exceptions make jargon less attractive.

Solution to Using Your Marbles

- 1.** On Tuesday, she again starts with 5 red marbles, 6 blue marbles, and 7 white ones. In order to make the bags more varied, she wants to pack the bags so that each bag differs by at least two “in-or-outs” from any other bag. An in-or-out consists of inserting a marble into a bag or removing a marble from a bag. So, for example, R and RR differ by only one in-or-out (insert R). On the other hand, RW and RB differ by two (remove W and insert B). How many different bags of marbles can Carol create that differ by at least two in-or-outs?

R, B, W, RBB, BWW, WWW, RBW, BRR; 8 altogether. The method here was to build up from small numbers. There could be no bags with two marbles because then they would differ from the singleton bags by just one marble.

- 2.** On Wednesday, Carol is given a bag that she cannot open. She knows she has 18 marbles. She also knows that there is at least one red, at least one blue, and at least one white. Knowing only this information but before seeing the marbles, Carol receives a phone call. “Can you guarantee to be able to give each child a bag with a different collection of marbles (i.e., at least one in-or-out apart), if there are eight children at a party? If not, then can you make this guarantee if there are seven children? If so, then how about nine?” How should Carol answer?

If there is one R, one B and 16 Ws, then the best that Carol can do is to pack seven bags: W, R, B, WW, WWW, WWWW, WWWW. The other W can be put either in a bag with a R or a B. Eight is not possible in this case because once there are only white marbles left, bags may differ only by number. If there are more than one red (or blue or both) marbles, but still fewer of these than white ones, then those red or blue marbles can be substituted for white ones in the larger bags. So Carol can guarantee at least seven bags, but no more.

- 3.** On Thursday, she again knows only that she has 18 marbles and at least one of each color, and is in the same situation as Tuesday as far as wanting the children enjoy variety. So she wants each gift to differ from every other by at least two in-or-outs. How many children can she guarantee to prepare bags for in that case?

Again, if there is just one R, one B, and 16 whites, Carol can guarantee to prepare just six bags: W, R, B, WWW, WWWW, WWWW. She can’t do better, because, once she gets to white marbles alone, the bags must differ by two marbles each.

- 4.** On Friday, her friend Diane packs the bags. Diane assures Carol that (i) every bag contains a different (some difference of number in at least one color) collection of marbles, (ii) that there are 18 marbles to start with, (iii) there are more white than blue and more blue than red, and (iv) that the maximum different ones she could pack was seven. Assuming that Diane is an extremely capable packer, what is the maximum number of red marbles there could have been at the beginning?

If there is one red and two blues, this works: W, B, R WW, BW, WWW, WWWW.

There are not enough marbles left for a bag containing five white marbles. If there are two reds and three blues, then it is easy to form eight bags even without using all the marbles: W, B, R, WW, WB, WR, WWW, WBW.

If there are two reds and at least five blues (the maximum number of blues is seven because there must be more whites than blues), we could have at least eight bags, e.g., W,

B, R, WW, WB, WR, BB, WBW. Even more are possible, but by not specifying them, we handle more cases.

Similarly, if there are three reds and at least four blues, we could have at least eight bags, e.g., W, B, R, WW, WB, WR, BB, WRW.

Similar reasoning holds if there are four or five red marbles. Five is the upper limit of red marbles. There cannot be six or more red marbles, because then there would have to be at least seven blue marbles, leaving only five white ones.

Solution to Flipping Colors

1. Can the knight flip the color of every square an odd number of times? If so, in how few moves?

Steve Schaefer has found a 25 move solution which he describes as follows: “If we take the upper left corner to be square (0,0), then the first twenty moves flip the outer three rows/columns of the board, returning to the starting point at square (2,3). The four squares in the center of the board have not been flipped, and the starting square has been flipped twice. The remaining squares (2,3), (3,3), (3,4), (4,3), and (4,4) can be flipped with the following five moves”:

(3,3), (4,3), (4,2)
 (4,3), (4,4), (3,4)
 (3,3), (3,2), (2,2)
 (3,2), (4,2), (4,3)
 (3,3), (2,3), (2,2)

With that introduction, we can now present Schaefer’s 25-move solution, starting at (2,3):

1: (1,3), (0,3), (0,2)
 2: (0,1), (0,0), (1,0)
 3: (1,1), (1,2), (2,2)
 4: (2,1), (2,0), (3,0)
 5: (3,1), (3,2), (4,2)
 6: (4,1), (4,0), (5,0)
 7: (6,0), (7,0), (7,1)
 8: (6,1), (5,1), (5,2)
 9: (6,2), (7,2), (7,3)
 10: (6,3), (5,3), (5,4)
 11: (6,4), (7,4), (7,5)
 12: (7,6), (7,7), (6,7)
 13: (6,6), (6,5), (5,5)
 14: (5,6), (5,7), (4,7)
 15: (4,6), (4,5), (3,5)
 16: (3,6), (3,7), (2,7)
 17: (1,7), (0,7), (0,6)
 18: (1,6), (2,6), (2,5)
 19: (1,5), (0,5), (0,4)
 20: (1,4), (2,4), (2,3)
 21: (3,3), (4,3), (4,2)
 22: (4,3), (4,4), (3,4)
 23: (3,3), (3,2), (2,2)
 24: (3,2), (4,2), (4,3)
 25: (3,3), (2,3), (2,2)

Solution to Scheduling Tradition

1. Can you form an 11-day schedule for these teams that satisfies the constraints?

This solution is from the classic puzzle book of Dudeney written a century ago. The pairs for each day are separated by spaces. You'll see that except for the A column, all other columns are different rotations of BEGFKHICDLJ, a sequence of letters in which all the non-A letters appear exactly once. Consider some letter, say E. By construction, E is in each one of the non-A columns exactly once. Therefore E plays against A. Also, by construction, E is in a given day's row exactly once (because all the rotations differ).

day 1: AB CD EF GH IJ KL

day 2: AE DL GK FI CB HJ

day 3: AG LJ FH KC DE IB

day 4: AF JB KI HD LG CE

day 5: AK BE HC IL JF DG

day 6: AH EG ID CJ BK LF

day 7: AI GF CL DB EH JK

day 8: AC FK DJ LE GI BH

day 9: AD KH LB JG FC EI

day 10: AL HI JE BF KD GC

day 11: AJ IC BG EK HL FD

How do we know that E encounters every other letter? In every pair of columns, except the leftmost, E encounters two teams. We just have to ensure that the teams in question are different every time. This is the case because if we list the pairs of columns by the offset from the first letter in the sequence BEGFKHICDLJ, then we see that the column paired with A begins at B, so is at an offset of 0. The next pair begins with C and D respectively at offsets 7 and 8, so their offsets differ by 1. We'll denote that fact by the sentence "CD are at offsets 7 and 8, so the offsets differ by 1."

CD are at offsets 7 and 8, so the offsets differ by 1

EF are at offsets 1 and 3, so the offsets differ by 2

GH are at offsets 2 and 5, so the offsets differ by 3

IJ are at offsets 6 and 10, so the offsets differ by 4

KL are at offsets 4 and 9, so the offsets differ by 5

Because the various column pairs differ by different offsets, a given pair XY will meet in only one column pair. For example, E and B will meet in the column pair with offset 1, as will E and G. E and F as well as E and J will meet in the column pair with offset 2 and so on. Whereas we have used E as an example, the same reasoning holds for all letters B to L. The case of A is clear by inspection of the first two columns.

Solution to Fractal Biology

1. Can you achieve the wounded distance two condition for eight nodes using 16 links, assuming no node can have more than four interaction links?

With eight nodes, one needs only 16 links, as demonstrated in Figure 1-18.




Figure 1-18: An eight-node network in which there is a two-link path from any node to any other even if a node is deleted. Within one foursome, there is a single link from any node to any other. For a trip from a node n in one foursome to a node m in the other foursome, there are two paths of length at most two. These go through different intermediate nodes.

2. What is the fewest number of links needed to achieve the wounded distance two condition for 12 nodes and at most five interaction links for any node?

With 12 nodes, one could design three sets of four nodes that are completely connected (requiring $3 \times 6 = 18$ nodes) and then add in four links between every pair of four, leading to an additional 12 links for a total of 30 altogether.

- 3. What is the fewest number of links needed for 12 nodes, but without any limit on the number of interaction links any particular node can have?**

If there is no limit on the number of links per protein node, try the design of Figure 1-19. You would need only 21 links. This is called the two-fan design, because each hub creates a fan. One needs two fans in case one hub is wounded.




Figure 1-19: A twelve node, wounded distance two network using only 21 links.

- 4. We have a particular network having 108 proteins. We don't yet know all the interactions. What is the fewest number of links needed to achieve the wounded distance two condition for any pair among these 108 nodes if there is a limit of 60 interactions that any single node can have? Try for a solution that uses fewer than 600 links.**

Divide the nodes into 96 nodes (the base nodes) and 12 other nodes that will perform a connection function (the switchboard nodes). Number the 96 base nodes from 1 to 96. Call 1 through 24 the A nodes, 25 through 48 the B nodes, 49 through 72 the C nodes, and 73 through 96 the D nodes.

A good way to conceptualize the problem is that each base node group forms the supervertex of a tetrahedron (large black circles). Along each of the six edges linking supervertices there are two “switchboard” nodes, as you can see in Figure 1-20.




Figure 1-20: Switchboard nodes are clear circles. Supervertices, each consisting of 24 nodes, are big dark circles. The basic design represents 576 links. In addition, we have shown two links (the dashed curves) between non-neighboring switchboard nodes. Altogether there are 12 such links.

Each switchboard node is connected to every node in the two supervertices that form its edge (2×24 links). Thus, the total number of links involving base nodes is $2 \times 2 \times 24 \times 6 = 576$. Two switchboard nodes are “neighbors” if they share a supervertex. Each switchboard node needs a direct link to non-neighbors only. Each switchboard is the neighbor to all but two of the other switchboards. Linking to those non-neighbors adds a total of $12 \times 2 = 24$ to the collective degree of the switchboard nodes. This translates to half that number of edges, for a total of 12. This yields a total of $588 = 576 + 12$ links.

This solution is due to TJ Takei.

Note: The network problems here fit a general formulation: Find the minimum number of links necessary for N nodes, maximum distance D between any pair of unwounded nodes where the maximum degree (number of interactions per node) K , and there may be up to X wounded nodes. Such questions might bear distinct relevance to biology. Here are two mathematical biology papers on similar topics that I like very much:

“Regulation of metabolic networks: understanding metabolic complexity in the systems biology era,” Lee Sweetlove and Alisdair Fernie, *New Phytologist* 2005, 168: 9-24.

“Evolutionary capacitance as a general feature of complex gene networks,” Aviv Bergman and Mark Siegal, *Nature* 2003 Jul 31; 424(6948): 549-52.

Solution to As Easy as Pie

1. Find a cut design obeying the above three rules that minimizes the total length of the perimeters for five pieces.

Make the first cut with a width of $1/5$ of the total, and the second cut perpendicularly with a width of $1/4$. Repeat the process with ever-wider widths ($4/15$ and $3/8$) to keep the same area until the last cut halves the remaining piece (see Figure 1-21). This gives a total perimeter of $10\frac{1}{6}$ or about 10.17.




Figure 1-21: Recursive perpendicular cuts satisfying the child cut rule.

2. How much better can you do for five pieces, after dropping the child cut rule?

The best design for five pieces using cuts that are parallel to the edges of the original square is shown in Figure 1-22. Rectangles 1, 2, 3, and 4 all have dimensions b by $1-b$ (having an area of $1/5$). This gives a total perimeter of $4(2 + 1/\sqrt{5})$, which is about 9.8.




Figure 1-22: Without the child cut rule, one can obtain a smaller total perimeter for five pieces.

3. What about nine pieces?

For nine pieces without the child-cut rule, we get the similar design in Figure 1-23. This gives a perimeter of about 15.3.




Figure 1-23: Without the child cut rule, here is the best design I know of for nine pieces.

4. Can you get a smaller perimeter for five pieces, if, in addition to dropping the child cut rule, you drop the rule that cuts must be parallel to the original sides?

A beautifully economical design for five pieces without the child cut rule that permits cuts at an angle to the original edges is shown in Figure 1-24. The slant is about 19 degrees from the horizontal. This gives a total perimeter of under 9.4.




Figure 1-24: Without the child cut rule or the requirement that all cuts must be parallel to an original side, one can reduce the perimeter quite noticeably.

TJ Takei provided these solutions, the best I know of for these problems. For the tilting calculation of question 4, he used Gnu Octave.

Solution to Lucky Roulette

- 1. Do you ask for a spin of the chamber before the third trigger pull? How about before the fourth trigger pull? Using your best strategy, what are your chances of survival ignoring the first trigger pull?**

If you don't have me spin again now, you have a 1 in 3 chance of dying in the third pull of the trigger, which is the same as if you ask me to spin. But the advantage of spinning is that if I spin now and you survive the third trigger pull, then you will return to the 1 in 4 chance of dying for the last trigger pull.

Using that strategy (don't spin before the second trigger pull, spin before the third, and don't spin before the fourth), your chances of survival beyond the first trigger pull are $3/4 \times 2/3 \times 3/4 = 3/8$. If you had asked me to spin every time, your chances of survival would be only $8/27$, a little less than $1/3$. Sorry for the violent setting, but I hope it helped focus your mind.

Solution to Legal Logic

- 1. If each patient has a 0.005 probability of being hurt, how likely is it to get one failure within the first 50 patients?**

As preparation for our later work, write a procedure using a random number generator that produces a 1 with a probability of 0.005 and a 0 otherwise for 50 times. Count the number of 1s. Call this procedure 10,000 times. You can see that there will be one failure in the first 50 patients approximately 22% of the time. This result can also be obtained analytically: probability of at least one failure = $1 - \text{probability of no failures in 50} = 1 - (1 - \text{probability of failure each time})^{50} = 1 - (0.995^{50}) = 0.222$.

- 2. Suppose that 18 people out of 4,000 have been hurt by the device. There are eight hospitals, each of which has treated 500 patients. What is the likelihood that at least one hospital exceeds a 1% test, i.e., has 6 patients who were hurt, even if the underlying chance of failure were in fact independent of hospital and is 0.005 per patient?**

To solve this question empirically, create groups corresponding to 1 to 500 inclusive, 501 to 1,000 inclusive, ..., 3,501 to 4,000 inclusive. Do the following 10,000 times: Take the 18 numbers uniformly at random without replacement (i.e., never take the same one again) over the range of 1 to 4,000. See if any group has more than 1% failures. You will see that between 14% and 15% of the time, at least one group does. This again is a false guilty rate.

3. Suppose the distribution of patients were much more skewed, with seven hospitals having treated only 200 patients each and the remaining one having treated all the rest. How likely is it that at least one of those smaller hospitals had hurt three people or more under the same conditions (i.e., 18 people who were hurt overall and the underlying chance of failure is 0.005 per patient)?

In this case, the tort lawyers will win around 35% of the time. The method is similar to the method we used to answer question 2 except the groups correspond to 1 to 200 inclusive, 201 to 400 inclusive, ..., 1,201 to 1,400 inclusive, and then 1,401 to 4,000.

4. Suppose we say that if a single hospital has more than 2% of bad outcomes, then that is a bad hospital. Assume that each hospital treats 500 patients. What is the likelihood for at least one hospital to hurt 10 or more people if the 0.005 hurt probability held and there were no bias in any hospital?

This would happen only about 5 times out of 10,000. This is so unlikely that there is a good chance the hospital really is bad.

5. How would you answer the last question in the case that seven hospitals each treat only 200 patients and one treats 3,600?

In this case, the probability increases to about 2.5% (250 times out of 10,000). I'd still avoid that hospital.

Solution to The Box Chip Game

1. For four chips and two guesses per color, can you figure out a way to win at least 1/6 of the time, assuming all rearrangements are equally likely?

Black guess: 1, 2

White guess: 1, 2

Red guess: 3, 4

Green guess: 3, 4

You win 4 times out of the 24 (= 4!) possible permutations. Why 4? Because black and white can be in either order among the first two boxes; red and green can be in either order among the last two boxes. Because 4 out of 24 is 1/6, you win 1/6 of the time. I think that's the best one can do. If you can show me wrong, then please let me know.

2. What if there are six chips and three guesses per color?

For six chips with three guesses per color, the best is 1/20. As puzzlist Alan Dragoo puts it: "Guess the same half of the boxes for one half of the chip colors and the other half of the boxes for the other half of the chips. This gives a probability of $(n/2)! \times (n/2)! / n!$ of winning." For 6, this comes out to be $3! \times 3! / 6! = 36 / 720 = 1/20$.

3. How does this probability generalize for n chips (n even) and $n/2$ guesses per color?

Following Alan Dragoo's suggestion is the best approach I know of. It works for all even numbers, which is fine because the puzzle allows only even numbers. Unfortunately, $(n/2)! \times (n/2)! / n!$ rapidly approaches zero. For 8, it's only 1.4%. For 12, it's 0.1%.

4. How well can you do in the case of four chips and two guesses under the procedural agent assumption? (Hint: You can win over 40 percent of the time.)

Identify the colors with numbers. This is an arbitrary association, but fixing it consistently will help.

Black — 1

White — 2

Red — 3

Green — 4

Here are the programs. The boxes are numbered 1, 2, 3, and 4 in left-to-right order.

Black: 1; Black→win, White→2, Red→3, Green→4

White: 2; White→win, Black→1, Red→3, Green→4

Red: 3; Red→win, Black→1, White→2, Green→4

Green: 4; Green→win, Black→1, White→2, Red→3

The notation means the following: Black (the agent associated with the black chip) starts at box 1; if box 1 has a black chip, then Black wins; if box 1 has a white chip, then Black next opens box 2; if box 1 has a red chip, then Black next opens box 3; and so on. Suppose box 1 has a white chip, so Black opens box 2. If box 2 then has a green chip, then Black next opens box 4.

This strategy wins in 10 out of the 24 possible permutations. To see why, consider any ordering of the colors in boxes.

1	2	3	4	win/lose
Black	White	Red	Green	win
Black	White	Green	Red	win
Black	Red	White	Green	win
Black	Red	Green	White	lose
Black	Green	White	Red	lose
Black	Green	Red	White	win
White	Black	Red	Green	win
White	Black	Green	Red	win
White	Red	Black	Green	lose
White	Red	Green	Black	lose
White	Green	Black	Red	lose
White	Green	Red	Black	lose
Red	Black	White	Green	lose
Red	Black	Green	White	lose
Red	White	Black	Green	win
Red	White	Green	Black	lose
Red	Green	Black	White	win
Red	Green	White	Black	lose
Green	Black	White	Red	lose
Green	Black	Red	White	lose
Green	White	Black	Red	lose
Green	White	Red	Black	win
Green	Red	Black	White	lose
Green	Red	White	Black	win

Notice that there are 10 wins here.

5. How well can you do if you have six chips and three guesses?

Out of the 720 possible permutations of the 6 elements, you can win 276 times (or more than 38 percent of the time). The agents all agree in advance on the following arbitrary association A of boxes to colors. These indicate which box next to pick based on the color found in the previous box.

Black→box 1

White→box 2

Red→box 3

Green→box 4

Blue→box 5

Orange→box 6

Here is the program for the agent corresponding to color X. Start at the box corresponding to color X in association A. If you find X, then you've won, so you can stop. Otherwise, look up the color you find, say Y, in A and look at that box. If you find X, then you've won, so you can stop. Otherwise, look up the color you find, say Z, in A and look at that box. If you find X, then you've won. Otherwise, you lose.

For example, White does the following:

White: 2; Black→1, White→win, Red→3, Green→4, Blue→5, Orange→6;

***Note:** I first heard this puzzle from Peter Winkler. Michael Rabin explained the procedural agent idea to me. What is truly fascinating about this problem is the use of arbitrary assignments. In this, as well as in several other puzzles described or invented by Winkler, the idea is to commit everything to one approach. You either win all the time or lose often — kind of a start-up mentality.*

By the way, if you happen to be an algebraist, characterizing when the contestant wins can be expressed succinctly: You win whenever the colors arranged in the boxes guarantee that the longest cycle in the arbitrary permutation is of length $n/2$ or less. Let me explain. First, notice that the mapping from colors to numbers is a permutation (a rearrangement) of those colors. For example,

Black→box 1

White→box 2

says that Black should be put first in the list, White should be put second, and so on. We don't move the chips, of course, but it is in this way that we can think of these assignments as a permutation. Now, a cycle occurs when following the box colors happens to return us to a box we've already visited.

Suppose, for example, that the six boxes have the following colors:

- 1: Blue
- 2: Black
- 3: Red
- 4: Orange
- 5: White
- 6: Green

and we use the arbitrary permutation:

- Black→box 1
- White→box 2
- Red→box 3
- Green→box 4
- Blue→box 5
- Orange→box 6

Suppose we are looking for Red. The permutation says to go to box 3. As it happens, we find Red in box 3. So the length of the Red cycle is one. Suppose we are looking for Black.

- The permutation says to look in box 1.
- We find Blue.
- So the permutation says to look at box 5.
- We find White.
- So the permutation says to look at box 2.
- We find Black.

At this point we win, but note that the next block to look at would be box 1. That is, the color of the chip that sends us back to the first box we looked at must be the color we are looking for. (For Black, we first look at box 1 and the color of the chip in box 2 would send us back to box 1.) In this case, the cycle in the permutation is of length three $1 \rightarrow 5 \rightarrow 2 \rightarrow 1$. In fact, for this permutation, every cycle is of length three or less.

On the other hand, suppose that the arrangement of the colors in the boxes were:

- 1: Blue
- 2: Black
- 3: Red
- 4: White
- 5: Orange
- 6: Green

Part I: Mind Games

Now consider starting with White using the same permutation as above:

Start in box 2.

Find Black.

Go to box 1.

Find Blue.

Go to box 5.

Find Orange.

Go to box 6.

Find Green.

Go to box 4.

Find White.

In this case, we have $2 \rightarrow 1 \rightarrow 5 \rightarrow 6 \rightarrow 4 \rightarrow 2$. This is a cycle of length five.

Moral of the story: A succinct statement of the solution may not lead to a succinct explanation of the statement.

Solution to Feedback Dividends

- 1. If P_{good} is 0.9, what is the probability of hitting row 8, column 5 using the FeedYes strategy and using the FeedNo strategy?**

The probability of winning under FeedYes is about 0.89 when P_{good} is 0.9. It cannot be better than 0.9 because if you are on the second from top row and adjacent to the winning square, there is some chance of losing. The strategy for FeedYes consists of trying to reduce the horizontal distance to the goal to zero (or to one on rows that are an odd distance away). Calculating the probability involves the following key recurrence ideas: If you are n steps ($n > 0$) away from the top row and 0 distance away, then your probability of hitting is the same as the probability of hitting if you are $n-1$ steps away and one away from column 5. Otherwise, with probability P_{good} you will be one column closer to your goal and with probability $1-P_{\text{good}}$ you will be one farther from the goal, both with $n-1$ steps.

Under the FeedNo strategy, the probability is about 0.55 when P_{good} is 0.9. To compute the probability for FeedNo, assume a strategy that will take four aims to the right and three aims to the left since this gives the best likelihood of hitting the destination. Given this, there are four ways to win: All seven aims are true; three of the right aims are true and two of the left are; two of the right are true and one of the left; and one right and zero left.

So, the feedback dividend is about 1.6.

- 2. For which value of P_{good} does the feedback dividend reach its greatest value? What is the feedback dividend in that case?**

The value of P_{good} for which the dividend ratio is highest is not 0.75 as one might think, but something around 0.772. For that value, the dividend ratio reaches 1.99.

- 3. If we cut off the three rightmost columns and the two leftmost columns, then which value of Pgood would give the highest feedback dividend? Assume that falling off the board gives a certain loss.**

In both the FeedNo and FeedYes cases, the analysis combines short trips: One wants to go from row 1, column 4 to row 3, column 4 in the first two moves, then to row 5, column 4 in the second two moves, then row 7, column 4, and finally row 8, column 5. For FeedNo, this gives a formula like the following:

$$(((P_{\text{good}}^2) + ((1 - P_{\text{good}})^2)^3) \times P_{\text{good}}.$$

For FeedYes, this gives a formula that is P_{good}^4 .

This gives a maximum feedback dividend of 1.75 when Pgood is 0.71.

Alan Dragoo contributed greatly to these solutions.

Note: I have a friend who likes to swim backstroke in swimming pools that have no lane markers. He admits to colliding with other people a lot.

Solution to Number Clues

- 1. For the first p, q pair, $\text{lcm}(p, q) = 60$, $p \times q = 240$. What are p and q?**

Here are the possible ways for two numbers to divide 240 and satisfy the lcm constraint.

- 4, 60 — No good because p and q must both be two digit numbers.
12, 20 — This could work.

So the solution must be 12, 20.

- 2. The product of p and q is 140. $\text{gcd}(p, q) = 2$. What are p and q?**

The factors of 140 are 2, 2, 5, and 7. Both p and q have 2 as a factor. So, the possible values of p and q are:

- 2, 70 — No good because p and q must both be two digit numbers.
4, 35 — No good for same reason.
10, 14 — Possible.

- 3. There is a pair p and q whose $\text{lcm}(p, q) = 35$. What are p and q?**

Because $\text{lcm}(p, q) = 35$, at least one of p and q must have one 5 and at least one must have one 7. To satisfy the two-digit requirement, they both must. So $p = q = 35$.

- 4. The sequence of numbers that opens the safe is a permutation of five numbers from the hints. The sequence enjoys the following property: The greatest common divisors between neighbors of the sequence strictly increase as one proceeds from left to right in the sequence, but all greatest common divisors are single digits. What is the sequence that opens the safe?**

The numbers that open the safe are: 10 12 20 35 14. The greatest common divisor between 10 and 12 is 2; the greatest common divisor between 12 and 20 is 4; between 20 and 35 is 5; and between 35 and 14 is 7.

Solution to Mind Games

- 1. What is the smallest number of bit questions sufficient to determine a five-bit sequence no matter what it is? Remember that none of your bit questions need be correct. You just need to know how to crack the secret at the end.**

Five are sufficient, even if all answers come at the end. Here are the guesses:

10000

11000

11100

11110

11111

Here's why this is enough. Number the positions from left to right at 1, 2, 3, 4, and 5. Consider position i . Suppose that there is a 1 in position i , for $i > 1$. Then the number correct will increase by one in the reported answers of the i^{th} bit question compared to the previous 1. For example, suppose the fourth position has a 1. That is, the secret code has the form $xxx1x$. Then 11110 will have a number count that is one higher than 11100. Conversely, if the fourth position has a 0, then 11110 will have a number count that is one less than 11100.

So far, we know about all bit positions except $i = 1$. We can determine the value of that bit position by examining the number correct for 10000. For example, if the last four digits are 0110, then the secret code is either 00110 or 10110. If the secret code were 00110, then the first bit question (10000) would yield a number correct of 2. If the secret code were 10110, then the number correct would be 3.

It turns out that four questions are enough also. Make the first two questions be 00000 and 11100. In most cases, this narrows down the bit sequences quite substantially. For example, if the answers to these first two questions are 4 and 1, respectively, then the remaining possibilities are: 00001 and 00010. So we will focus on the two difficult situations where there are six possibilities left. Those situations occur when the answers to the first two questions are either 2 and 3 respectively or 3 and 2 respectively. Suppose the answers to the first two questions are 2 and 3 respectively. Then the remaining possibilities are:

01101

01110

10101

10110

11001

11010

Let's say we ask 01110, a codeword that differs by at least two bits from both of the first questions. This will give us the following possible answers:

0 — impossible

1 — remaining possibilities: 10101 and 11001

2 — impossible

3 — remaining possibilities 01101, 10110, and 11010

4 — impossible

5 — remaining possibility 01110

The only difficult case is the one where the remaining possibilities are 01101, 10110, and 11010. In this case, we guess 00101. There are only three possible answers, each of which leads to a unique conclusion.

0 — 11010

2 — 10110

4 — 01101

Philip Daniel Rogers suggested this approach to a solution. There is no elegant closed form proof that this always works and the program is a bit complex. However, the program for the case when there is no feedback is quite simple, as you'll see next.

- 2. Suppose we change the game to make it a little more difficult: You ask some number of questions and then I answer them all after you have asked all your questions. In that case, how many bit questions are sufficient?**

Surprisingly, the very same guesses used in question 1 give a unique "signature" for each possible secret. That is, you make the following guesses:

00000

11100

01110

00101

Part I: Mind Games

Each possible sequence of five bits will give four answers and those answers differ for any pair of sequences. For example, suppose 11001 is the secret. That secret yields an answer of 2 correct to the bit question 00000, 3 to the bit question 11100, 1 to the bit question 01110, and 2 to the bit question 00101. No other bit sequence produces this same signature of answers: 2 3 1 2. Thus, you make the above four guesses, look at the signature, and determine which is the secret bit sequence. Writing a program to compute the signature given a series of guesses is, of course, very easy.

- 3. Given that the low info answers come at the end and there is no answer at all to the first bit question, is it possible to find the secret number, no matter what it is? If so, how many bit questions (counting the first one) do you need to pose in order to guarantee finding the secret number?**

If, after the first guess, you are told only whether the number of correct ones has increased or decreased and all these answers come at the end, then six bit guesses suffice to know the secret code. Start with 00000.

00000
10000
11000
11100
11110
11111

You can determine whether each bit position is 1 or 0 by whether the number correct has increased or decreased relative to the previous one.

- 4. How many bit questions do you need, if you hear the low info answers immediately after posing them?**

If the answers are given immediately after they are asked, Tom Rokicki suggests the following approach requiring only five questions altogether. The basic idea is to use the fact that three answers are possible (same (S), increase (+), or decrease (-)) whereas our method for the previous question used only + and -. Start with 00000 again. Next guess 3 (that is 00011 in binary). This changes two bits, so if the answer is S (same number correct), then one of these 1s is correct and one isn't. On the other hand, if the answer is + (increase), then both must be correct. Similarly, if the answer is -. Continue with this reasoning and you'll find the answer. If you get stuck, see the following:

```
[3,
  S: [29,
    SS: [4,
      SS+: [14,
        SS+-: {5},
        SS+S: {22},
        SS++: {14}],
      SS-: [14,
        SS--: {17},
        SS-+: {26},
        SS-S: {9}]],
    S-: [5,
      S-S: {10,18},
      S-+: [10,
        S-S: {6},
        S-++: {2},
        S-+-: {1}]],
    S+: [5,
      S+S: {13,21},
      S+-: [10,
        S+-+: {30},
        S+--: {29},
        S+-S: {25}]]],
  -: [5,
    -S: [9,
      -SS: {0,16},
      -S+: {8,24}],
    -+: [8,
      +- -: {4,20},
      -++: {12,28}]],
  +: [5,
    +-: [8,
      +-+: {11,27},
      +--: {3,19}],
    +S: [9,
      +SS: {15,31},
      +S-: {7,23}]]]
```

The answers to the left of the colon are the previous answers. The indentation gives precedence. For example, `+: [5` means “if the answer to the first question is +, then ask about 5 (00101).”

5. How does this generalize if the code is N digits long, but the digits could be base b for some $b > 2$? For example, b could be 10 and all digits 0 through 9 would be allowed. I have a solution when all answers come at the end, but that requires $1 + (b-1) \times N$ questions. I don't think it's optimal.

Part I: Mind Games

We can use the reasoning of question 3. Here are the first questions to find a secret that is base 10.

00000
10000
20000
30000
40000
50000
60000
70000
80000
90000
01000
02000
03000
04000
05000
06000
07000
08000
09000

This yields a set of questions of size $1 + N(b-1)$, but I await a clever reader to suggest a better approach.

Solution to Refuse and Reveal

1. Rosalba: "I have five integers that may or may not be distinct."

Quentin: "What is the minimum?"

Rosalba: "20."

Quentin: "Which of these would not allow me to infer all their values: number distinct, mean, maximum, or median?"

Rosalba: "Only the median."

Quentin: "Great. I know the numbers."

What are they?

They must all be 20. Any of the other statistics would reveal this fact for sure.

- 2. Rosalba: "I have seven integers that may or may not be distinct."**

Quentin: "What is the minimum?"

Rosalba: "20."

Quentin: "Which of these are you willing to tell me (i.e., would not allow me to infer all their values): mean, median, and maximum?"

Rosalba: "All of them."

Quentin: "Ok, what is the maximum?"

Rosalba: "21."

Quentin: "I know which of mean and median you're willing to tell me now."

Which? Why?

The mean would be enough to determine how many 20s and how many 21s there are. For example, if there are four 20s and three 21s, then the mean would be about 20.43. Knowing the median would not be enough because a median of 20, for example, could result from a situation in which there are three 20s and four 21s or five 20s and two 21s.

- 3. Rosalba: "Can you find some situation in which I would be happier to tell you the mean rather than the median?"**

Quentin: "Could you give me a hint?"

Rosalba: "In an example I can think of, there are three numbers, two of them distinct." Give it a try.

Suppose there are three numbers and two distinct ones. If you know that the minimum is 20 and the median is 23, then you know that the other number is also 23. On the other hand if you know that the mean is 22, for example, this could result from 20, 20, and 26.

- 4. Rosalba: "Can you find some situation in which all of minimum, maximum, mean, and median are necessary and sufficient to find the identities of five numbers that are all integers?"**

Minimum is 20, maximum is 22. Mean is 21.2. This could be explained by three 22s and two 20s or by two 22s, two 21s, and one 20. The median will tell you which.

- 5. Rosalba: "There are 17 numbers that are not all distinct. Their minimum is 30, mean is 34, and median is 35."**

Quentin: "What is their total distance to 35?"

Rosalba: "I won't tell you, but the total distance to 35 is five less than the total distance to 38. Whoops! I shouldn't have told you that."

Quentin laughing: "You're right. Now I know the numbers."

What are they?

The point to which the total distance is guaranteed to be at a minimum is the median. That is, if there are an odd number n of values altogether, then the sum of the distances to the

median is smaller than to any other point. Here is why. Suppose the median is m . Let the total distance to m be x . Now consider the total distance to $m+1$. Because m is a median, there are $(n-1)/2$ values that are m or less. So, at least $1+(n-1)/2$ values increase their distance by 1 when calculating the total distance to $m+1$, including the median again. At most $(n-1)/2$ values decrease their distance by at most 1 when calculating the total distance to $m+1$.

So if there are 17 numbers and the median is 35, the mean is 34 and the distance to 38 is $x+5$, then there is one instance of 35, no instances of 36, and one instance of 37. Here is why: If there were one instance of 35, zero instances of 36 and 37, then the distance to 38 would increase by $3 \times 8 = 24$ for the numbers below 35 and decrease by 3×8 for the numbers greater than or equal to 38. The distance for 35 itself would increase by 3. This would give a net increase of 3. The fact that there is a net increase of 4 must mean that there is an instance of 37. In this case the net decrease for the upper eight values would be $3 \times 7 = 21$ for those greater than or equal to 38 and 1 for 37. This implies that there are seven values that are 38 or greater and one 35 and one 37. Because the mean is 34, the total distance to the mean of the numbers greater than the mean is at least $1 + 3 + 28$; we get those numbers from the 35, the 37, and the seven values that are 38 or greater. This total is 32. The values below the mean must match this, so they must all be 30s. Since they cannot be less than 30, the uppermost value cannot be greater than 38. So there are eight 30s, one 35, one 37, and seven 38s.

- 6. How would your answer to this question change if there were 1701 numbers but otherwise the same information as in the previous question?**

The reasoning is exactly the same as for 17 numbers with one 35, one 37, 849 38s, and 850 30s.

Tom Rokicki helped greatly to clarify the formulation of this puzzle.

Solution to A Biting Maze

- 1. The challenge is to find a route to the final page, and to decrypt the words and phrases on the way. There are hints in the encrypted words (they form a sentence of natural poetry) and in other parts of the Web page. Give it a try.**

Here is a list of the page ids along with the decryption of the words and phrases in those pages.

f42, if

f37, you

f33, would

f32, see

f39, all

f44, of

f28, nature

f52, gathered

f23, up

f51, at

f17, one

f13, point

f46, in

f53, all

f34, her

f29, loveliness

f14, and

f19, her

f36, skill

f0, and

f2, her

f30, deadliness

f16, and

f4, her

f10, sex

f49, where

f15, would

f31, you

f3, find

f12, a

f27, more

f7, exquisite

f8, symbol

f22, than

f9, the

f6, mosquito

f25, if you have found the whole route to get here, you have solved the maze.
congratulations.

The great natural biologist Havelock Ellis wrote this sentence in 1920. Take a look before you squash the next bug.

Solution to Mad Mix

- 1. We want a mixture having 2.7 liters of supertox in a three-liter mixture. We don't want to pour more than 10 liters of any liquid laced with supertox down the drain. How can we achieve this?**

First obtain 9 liters of supertox in the 10-vessel. (Fill the 10-vessel with supertox, pour the result into the 7-vessel until it fills, empty the 7-vessel into the reservoir, and then pour the three remaining in the 10-vessel into the 7-vessel. Now fill the 10-vessel again with supertox and fill the 7-vessel, which takes only 4 liters (since it already has 3). So the 10-vessel has 6 liters. Empty the 7-vessel into the supertox reservoir and pour the 6 liters from the 10-vessel into the 7-vessel. Now, fill the 10-vessel with supertox again. Pour into the 7-vessel until it is full, leaving 9 liters of supertox in the 10-vessel.) Then add one liter of water to the 10-vessel. Pour this mixture into the 7-vessel, leaving 3 liters in the 10-vessel having 2.7 liters of supertox.

- 2. Can we get a mixture that is $\frac{2}{9}$ supertox and the rest water without pouring any supertox down the drain?**

First obtain 2 liters of supertox in the 10-vessel. To obtain 2 liters of supertox in the 10-vessel, fill the 7-vessel with supertox, and empty it into the 10. Repeat; but now there are 4 liters left in the 7-vessel. Return the 10 liters to the supertox reservoir. Now pour the 4 liters into the 10-vessel and refill the 7-vessel with supertox. Now pour from the 7-liter vessel to the 10. When you are done, there will be 1 liter left in the 7-vessel. Empty the 10-vessel back into the reservoir and then transfer the single liter from the 7-vessel to the 10-vessel. Fill the 7-vessel from the reservoir and pour into the 10-vessel, leaving space for two. Fill the 7-vessel again from the supertox reservoir and pour into the 10-vessel. This fills the 10-vessel and leaves 5 in the 7-vessel. Empty the 10-vessel back into the supertox reservoir and pour the 5 from the 7-vessel to the 10-vessel. Refill the 7-vessel. Pour from the 7-vessel into the 10-vessel, leaving two liters in the 7-vessel. Empty the 10-vessel back into the supertox reservoir and then pour the two liters from the 7-vessel to the 10-vessel. Finally, fill the 7-vessel with water and pour that into the 10-vessel.

- 3. Can we get a 26% concentration of supertox in the 7-liter container?**

With a big extra bucket, one can obtain any mixture satisfying a ratio $p/(p+q)$ of supertox where p and q are whole numbers whose sum is 100 or less. Here's what you do: Create p liters of supertox and dump them into the big extra bucket and then q liters of water and dump. Now you have that ratio. Suppose you want 53 liters, for example. You fill the 10-liter vessel five times and put the results in the big bucket. Now fill the 10-liter vessel one more time, pour into the 7-liter vessel. The three liters left over go into the big bucket. So, in this case, we want a mixture of 10 liters or more having a concentration of 26% of supertox. So, if we had 13 liters of supertox and 37 liters of water, that would do it. Fill the 10-liter vessel with supertox; pour into the big bucket. Fill the 10-liter vessel with supertox

and pour it into the 7-liter vessel. The 3 remaining liters go into the big bucket. Empty the 7-liter vessel into the supertox reservoir. Obtaining 37 liters of water is easy, of course (three 10s and a 7). Once you have the mixture, fill the 10-liter vessel. There will be 2.6 liters of supertox in there.

This solution is due to Vlad Simionescu.

- 4. Suppose that we want to get mixtures that are all supertox, $1/2$ supertox, $1/3$ supertox, $1/4$ supertox, ... $1/50$ supertox. You are allowed to specify the size of the two vessels. You may even include fractional sizes as long as the smallest vessel contains at least one liter and the largest contains at most 30 liters. In addition, you have the big bucket again with capacity over 100 liters. Which two vessel sizes would you choose in order to obtain each of these mixtures? Given those vessel sizes and given a target mixture, how would you obtain it?**

Suppose we want all fractions up to $1/k$. In our case, $k = 50$. Let the basic unit be 0.6 liters and then the smaller bucket can be 17 basic units or 10.2 liters, and the larger bucket can be 50 basic units or 30 liters. Because 17 is a prime number, we can generate any number k of units between 1 and 49 from these even without the extra bucket. Start by generating 1 unit of supertox and putting it in the big bucket (which can be very small in fact). Then generate k units of water and mix.

This extremely elegant solution is due to Ivan Rezanka. Ivan has shown how to solve this problem without the extra bucket. The basic hint is to put an amount of supertox in the large vessel such that when the vessel is subsequently filled to completion with water, the desired concentration is achieved. This can be done if the larger vessel is viewed as consisting of L units where L is the least common multiple of numbers 2, 3, ..., k .

Solution to Dig That!

In these solutions, call the north-south roads columns and the east-west roads rows. The center line is column 4 between Start and End.

- 1. If the tunnel is at most eight blocks long and begins at Start and ends at End, then what is the minimum number of probe devices you would need to guarantee to determine the precise route of the tunnel in one hour?**

Because the straight line road from Start to End is of length six, the most the tunnel can deviate from the center line is by one. For this reason, place probes one away from the center line at the intersections indicated by the circles in Figure 1-25.




Figure 1-25: Because you know the direction of the roads as they leave an intersection, each pair on each row determines the fate of the tunnel on that row and the next.

Here is the reasoning. If no probe detects a tunnel, then the tunnel goes north up the center. It has no choice, because there are no loops. If a tunnel goes through a probe (represented by a circle), we can tell whether it is coming from the south and whether it is going north, west, or east. For example, if neither probe at row 2 detects the tunnel, then the tunnel must go north up the center at least till east-west cross street 3. By contrast, if the tunnel approaches the probe at row 2, column 5 from the west and leaves going north, then we know the following: The tunnel went north from Start one block, then turned east one block (where it encounters the probe), then turned north one block (to arrive at column 5 and row 3). In general, each pair of probes at row i can determine which column (north-south road) was used in coming from row $i-1$ and which column is taken to go to row $i+1$. That information is enough to determine the east-west directions on each row.

2. If you had only one probe, what is the minimum time it would take to guarantee to determine the outcome, assuming you could move the probe every hour?

You can guarantee to map out the tunnel in six hours. Test the points in Figure 1-25 in any order. Suppose every probe is negative in your first five attempts. No matter which probe point you leave out, there will be an ambiguity in the root (straight north up the center or a detour at the probe point still unexplored). You may be able to cut out early if a detour and a return to the central path has been detected, but this is not guaranteed.

3. If the tunnel is at most eight blocks long and begins at Start and ends at End, then what is the minimum number of point probe devices you would need to place to guarantee to determine the precise route of the tunnel in one hour?




Figure 1-26: The center column's point probes tell whether the tunnel goes up the center line. If it misses the center, it must go either east or west. Because of the length limitation, it cannot go in both directions. Without reading the explanation, you may be able to figure out which column is taken between every pair of rows.

The problem can be solved with nine probes as in Figure 1-26.

From row 1 to row 2, the tunnel takes column 3 if neither probe A nor B detects it. The tunnel goes up column 4 if A detects it, and column 5 if B detects it but not A. From row 2 to row 3, the tunnel goes up column 3 if neither C nor B detects it, column 4 if C detects it, and column 5 if B is the only one that detects it. From row 3 to row 4, the story is similar. The tunnel takes column 3 if neither D nor E detects it, column 4 if D detects it and column 5 if E detects it. Row 4 is just like row 2, so the pattern continues.

4. How few point probes could you use over two hours to guarantee to find the tunnel?

In the first hour, just probe the center streets as in Figure 1-26. In the second hour, probe the eastern side of the first miss. If the tunnel is there, then it stays east until it rejoins the center. It cannot go to the western side because the tunnel is only eight blocks long.

Peter Carpenter of Wales contributed to these solutions. Tom Rokicki found the solutions for 3 and 4.

This problem becomes much more challenging if the tunnel could be longer. Solving it in general is an open problem.

Solution to Preferential Romance

1. Are Bob and Alice compatible, passably compatible or neither? After dropping the smallest number of edges in zero, one, or both spouse's preferences, try to find a consistent ordering.

They are passably compatible. Drop Bob's MC edge. (He should care more about culture, it seems.) Here is a consistent ordering: BRKCFWFOEMHTJ.

2. Can you describe an algorithm for the counseling company to use to help marriages in distress? That is, try to find a method so spouses have to drop as few preferences as possible.

The technique uses graph theory. Each circle in the figure is a "node" and each arrow represents a one-way edge. In graph-theoretic notation, let graph $G1 = (N1, E1)$, where $N1$ is a set of nodes and $E1$ is a set of directed (one-way) edges on those nodes. Similarly, $G2 = (N2, E2)$.

In the pseudo-code below, $\cup$ means set union. For example $\{X, Y, Z\} \cup \{Z, W\} = \{X, Y, Z, W\}$. Also, $\cap$ means set intersect. For example $\{X, Y, Z\} \cap \{Z, W\} = \{Z\}$. Finally, the minus sign ($-$) means set difference. For example $\{X, Y, Z\} - \{Z, W\} = \{X, Y\}$.

A cycle is a tour that starts at some node X and ends at X such that it always follows the one-way edges. A root is a node such that no one-way edge points to that node.

Here is the pseudo-code:

```

Given  $G1 = (N1, E1)$  and  $G2 = (N2, E2)$ 
 $H := (N1 \cup N2, E1 \cup E2, \text{EdgeMarks})$ 
  where EdgeMarks is of the same size as  $E1 \cup E2$ 
  If  $e$  is in  $E1 \cap E2$ , then  $e$  has a null mark.
  If  $e$  is in  $E1 - E2$ , then  $e$  has a mark of 1
  If  $e$  is in  $E2 - E1$ , then  $e$  has a mark of 2.
To find a consistent set of preferences, find the fewest
edges having marks of 1 or 2 such that dropping them eliminates
all cycles.
To find a consistent ordering, start at a root of  $H$ 
and write down its label. Then eliminate that root and
edges leaving that root from  $H$ , then repeat.

```

This pseudo-code guarantees to find a consistent ordering, because any cycle-free graph will be consistent. To obtain a cycle-free graph, it is sufficient to eliminate marked edges, because unmarked ones are shared and we know that each spouse alone has produced a graph without cycles.

This is not quite an algorithm in that I did not tell you how to find the fewest edges to drop. In fact, that problem is NP-complete (a topic we'll discuss in Part II), so it is hard to find the fewest such edges, but heuristics such as simulated annealing (also discussed in Part II) work well.

Solution to No Change for the Holidays

1. If you did not know how much the purchase price is except that it is a whole number amount between \$1 and \$100 inclusive, which amounts would you write on your three checks in order to minimize Claude's change?

If the three checks are made out in the amounts of \$15, \$30, and \$60, then you reduce the maximum change left with Claude to \$14 no matter what the price. To see this, note that you can combine checks to sum up to \$15, \$30, \$45, \$60, \$75, and \$105. Because these sums differ by only \$15, any amount between \$1 and \$100 is no more than \$14 less than one of these sums.

2. Suppose Claude publishes his four whole number prices in an advertisement that you see. Can you show how he can guarantee to do so in such a way that at least one item will yield him non-zero change no matter which three check amounts you write?

Claude might publish \$50, \$51, \$53, and \$59. Suppose you try to find a set of checks that eliminates all change no matter which item your teenager selects. Some check must be \$50 or less or the \$50 item would require change. If \$50, then another must be \$1, so the third must be \$8, because every item must be purchasable. In that case, the \$53 would require change. If your teenager has no check for \$50, but one for a lesser amount x , then your teenager also requires a check for $50 - x$. The remaining check must reach \$59, so \$51 and \$53 will require change. Claude wins.

Solution to Quiet in the Depths

1. Assuming blood pressures are three digits long, can each department report one three digit number so that the captain can infer a blood pressure among the 10 middle blood pressures of the 25 sailors?

Each department sends back its median. Order them in ascending order and denote the medians as M1, M2, M3, M4, and M5 and their respective departmental groups as G1, G2, G3, G4, and G5. Figure 1-27 shows this ordering




Figure 1-27: The groups have been ordered in ascending order based on their medians. Thus, $M1 \leq M2 \leq M3 \leq M4 \leq M5$. Within each group imagine low blood pressures to be at the bottom of the line segment corresponding to that group and high blood pressures near the top.

The number M3 must be one of the middle nine of the 25 blood pressure values. Here is why: Two sailors in G3 have blood pressures less than or equal to M3, the three lowest blood pressures of G2 must be less than or equal to M3, and the three lowest pressures of G1 must be less than or equal to M3. This means that at least $2 + 3 + 3 = 8$ sailors must have lower or equal blood pressures to M3. Those are colored black in Figure 1-28.




Figure 1-28: The dark rectangles represent values that are less than or equal to M3.

Symmetrically, at least eight sailors have blood pressures greater than or equal to M3. This means that M3 must be greater than or equal to at least eight values and less than or equal to at least eight values, so M3 is in the middle nine.

2. The captain wants to find some blood pressure value that is sure to be less than or equal to the median and some blood pressure value that is sure to be greater than or equal to the median. Using the assumptions of the first question, can the captain obtain such values after one three-digit report from each department?

Again the captain asks for the medians of each department. Ordering them again as M1, M2, M3, M4, and M5, the captain knows that M1 is less than or equal to the median of the 25 blood pressure values and M5 is greater than or equal to the median of the 25 values. Here is why: M1 is less than or equal to two blood pressure values of G1 and three from each of the groups G2, G3, G4, and G5 for a total of 14 values. These are shown in Figure 1-29. So M1 is certainly less than or equal to the median. Symmetrically, M5 is certainly greater than or equal to the median.




Figure 1-29: The dark rectangles represent values greater than or equal to M1.

***Note:** The idea of using the median of the medians comes from “selection algorithms” for “order statistics.” The paper that introduced this idea was written by some of the greatest algorithmists of all time: Blum, Floyd, Pratt, Rivest, and Tarjan.*

