[SQUEAKING]

[RUSTLING]

[CLICKING]

BRYNMOR CHAPMAN: Hello? Can people hear me? Are we good in the back? Are we good? Can people hear me? Good, good, good.

OK, so today, we are going to move on from counting and start the last unit on probability. Probability is very important. It's one of the most important disciplines, not just in computer science, but in all of the sciences. Please don't tell your other science instructors that I said this.

You use it in-- like we've already seen it in this class with the Miller-Rabin algorithm. You also see it in other randomized algorithms that you'll see later in 6 1210, 1220, 5220 if you take that. We see it in game theory, information theory, signal processing, cryptography, machine learning, medicine, stats, forensics. The list goes on and on.

But unfortunately, it's also one of the least understood disciplines. So there's a pretty famous quote attributed to Mark Twain, Benjamin Disraeli, several other people. There are three kinds of lies. There are lies, there are damned lies, and statistics. So you can kind of lie with probability and statistics by saying things that are true that are also very misleading.

Common sense is often very demonstrably completely unreliable, which we'll see in a moment. There are many graduate students, like even in our department-- again, don't tell people I said this-- they can do all sorts of exciting maths. But you give them a basic probability question, and they don't know where to start.

There are many garbage papers that are based on faulty understanding of probability. Even some of our colleagues have some, let's say, issues. Today, we are going to look at an example. Who has heard of the Monty Hall problem?

Oh, nice. Decent number of people. So, yeah, we'll basically try and use the Monty Hall problem to illustrate that it's really best to just throw away your intuition and just fall back on rigorous step-by-step analyses.

So the Monty Hall problem, it kind of became famous in 1990 when this magazine columnist named Marilyn vos Savant received a letter from a Craig Whitaker, which read as follows. Suppose you're on a game show, and you're given the choice of three doors. Behind one door is a car and behind the others, goats. You pick a door-- say, number 1-- and the host, who knows what's behind the doors, opens a different door-- say, number 3, which has a goat.

He says to you, Do you want to pick door number 2? Now, the question is, Is it to your advantage to switch? And who thinks yes, it is advantageous to switch?

Who thinks no? Who thinks you're better off sticking with your original choice? Who thinks it doesn't matter?

OK, so it looks like most people think that it's advantageous to switch. We've had a few abstentions. I'm going to count those as it doesn't matter.

So Marilyn replied that it's indeed advantageous to switch. If the car is behind the originally picked door, then you lose by switching. But if it's not, which is twice as likely, then you win by switching. And, of course, she soon received a torrent of hate mail, including from professional mathematicians, adamantly asserting that she was wrong.

So the hypo that was described in this letter became known as the Monty Hall problem after the host, Monty Hall, of the game show which had games like this where contestants faced very similar situations. So today, we're going to analyze it. But first, can I get a handful of volunteers to come play Let's Make a Deal?

OK, so empirically, at least, it looks like it's better to switch, right? Most people switched, and most people won the prize when they did that. So even though there's only one door with the prize, it seems to be pretty promising, right?

So let's see if we can analyze this mathematically. So we're going to solve the Monty Hall problem using what we call the tree method, which is a very simple and elementary approach that doesn't rely on intuition. It's just going to rely on mathematical rigor.

So the tree method uses four-ish steps. We used to call it the four-step method. But I couldn't for the life of me tell you what the four steps are if they weren't written in front of me. So I prefer to call it the tree method.

Oh, I can't even count. It kind of has five steps. We've got a step zero. Before you even start, you really want to make sure that you know what question you're trying to answer.

So before we can even think about solving a mathematical problem, we need to know what the problem is. We need to understand the setup, make sure we know what we're trying to ask before we can start analyzing it. So the letter from this Mr. Whitaker was not entirely precise. So we've got to figure out what axioms we're working with.

So we're going to assume the following. A prize is equally likely to be in any box. So, in principle, it's possible that Monty Hall tries to be tricky and tries to get into the contestants' minds, figure out what box they're going to pick, and put the prize somewhere else or, alternatively, put the prize somewhere where they will pick it, because it's an exciting game show.

But we're going to assume that Monty is just picking the boxes randomly. So the prize is going to go into a randomly selected box, completely uniform, and then the other two will have the goats. We're also going to assume that Marilyn or your colleagues, whoever is doing the picking, is equally likely to pick any box for the first pick.

She doesn't know anything about what Monty is doing. She's just going to toss some random coins in her head, and uniformly, at random, she'll just pick one of the boxes. It doesn't really matter where the prize is.

Now, after Marilyn picks her box, Monty must reveal a non-prize. And moreover, it cannot be in Marilyn's box.

So all that Mr. Whitaker said was Monty did reveal a non-prize that wasn't behind the chosen door.

It's not entirely clear that this must happen. And, in fact, I believe on the game show, there were instances where he didn't. So we're just going to assume that Monty is going to reveal a prize, regardless of whether Marilyn picked the correct box to begin with.

And, lastly, if Monty has a choice, he picks randomly. So if Marilyn has picked a non-prize to begin with, there's only one other non-prize. So Monty has to reveal that one.

But if Marilyn has picked the prize box to begin with, then Monty could, in principle, open either of the other two boxes. And that would still follow this rule. So when Monty has a choice like this, we're going to say that he tosses a coin, and he's going to open one of them at random.

So note that these are all axioms. We are just taking them on faith. They may or may not have merit. But if you accept them, then you also have to accept whatever conclusions we draw from them.

But I should point out that there are other perfectly reasonable sets of axioms that would lead us to different conclusions. Like, maybe Monty is very rude. And if Marilyn has picked the prize, then he offers her the option to switch. And if Marilyn has not picked the prize, then he forces her to go with whatever she picked.

This is a totally reasonable interpretation. And that would lead us to a different conclusion. Or, conversely, maybe it's benevolent. Maybe if Marilyn didn't pick the prize initially, he'll always try and encourage her to switch. But if she picked the prize, he'll just be like, nope, you got to stick with that.

So now that we've got these axioms, we can now pose the mathematical question, What is the probability that Marilyn wins by switching?

So yes, this is why I singled you out. I don't know where you have gone. Your colleague who did not switch, we're trying to analyze this probability, which is why I singled you out. So I apologize for that, but it's for the maths. Don't take it personally.

OK, so now we have a well-defined question. So now we can analyze it using the four-step method. The first of our four steps-- well, the tree method, which has four steps-- we need to figure out what we call the sample space.

So whenever we're thinking about a probability problem, underlying it is always what we call a probability space. And a probability space consists of two things.

These two elements are a set, which we call the sample space.

And for the purposes of this class, we're only going to be looking at discrete probability spaces.

So S is going to be a non-empty and countable, or finite, set. We're not going to look at sets like R or something like that. So this could be a set of three boxes, or it could be the natural numbers, something like that, something that's discrete and countable.

And then we also have a function, which we're going to denote by the letters Pr, which maps S into the unit interval 0 to 1. So this is going to be a total function with-- I probably shouldn't use S here. Let's call it x.

And, slightly oddly, when we're talking about probabilities, we use square brackets instead of parentheses to enclose the argument. But this is just a function. You can think of it as, like, Pr of x.

so We're going to have the requirement that if you sum up Pr of x over all of the x in your sample space, you should get 1. That's what it means for this to be a probability function. Oh, actually, why don't we leave both of these up and move onto this board?

OK, so this is called our probability function.

And elements of s are what-- we're going to call those outcomes.

So, we've got a sample space S, which is a set of outcomes. And we've got a probability function, which maps each outcome to a number between 0 and 1. And it should be the case that all of those probabilities should sum to 1.

Does that make sense to everybody? People reasonably Happy with this? Questions, comments, concerns? OK.

So our first objective is going to be to figure out what our sample space should be. So we're going to do that using what we call a tree diagram. So, for the moment, we're going to disregard the probability function and just focus on S.

And then once we've figured out what S is, we can figure out how to assign probabilities to everything. So the tree diagram, as you might expect, it's going to be a diagram in the shape of a tree, where each level of the tree represents some step in the random process we're trying to model.

So in the case of the Monty Hall problem, what might be the first step? What's the first thing that happens in the Monty Hall problem, the first source of randomness? Yeah?

AUDIENCE: You pick a first door.

BRYNMOR CHAPMAN: OK, there's something random that happens even before that. The answer was, you pick a door. Yeah?

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: Yeah. So the first random thing that happens is I, or Monty, picks which door the prize is behind. So that's going to be the first level of our tree diagram.

So let's start drawing out our tree diagram. We're going to have our root here. And let's label this with Monty hides prize.

And we're going to have three edges, which go to the three possibilities for what happens here. So it could be in door A, door B-- actually, why don't I call these purple, gold, and silver? Where's the eraser? Here we go. So we've got our purple box, our gold box, and our silver box.

What's the next random thing that happens? Yeah?

AUDIENCE: You choose the door.

BRYNMOR CHAPMAN: Yeah.

AUDIENCE: Same as last time.

BRYNMOR CHAPMAN: Yeah. So as your colleague said last time, the next random thing that happens is Marilyn, or the contestant, picks a door. So M-- oh, dear. Monty and Marilyn both begin with the same letter, OK. Let's just call this Student.

So now each of these is going to split into three more possibilities. So our contestant can pick purple, gold, or silver for each of these three things. Regardless of where Monty has hidden the prize, the contestant can pick any of the doors.

What's the next random thing that could happen? Yeah? Were you raising your hand?

AUDIENCE: Sorry about that.

BRYNMOR CHAPMAN: Oh, no.

AUDIENCE: If Monty has a choice of which door to reveal [INAUDIBLE].

BRYNMOR CHAPMAN: Yeah, so if Monty has the choice of door to reveal, then he picks one of the two remaining doors randomly.

So, so far, we've got three branches where that could happen. We've got both Monty and Marilyn pick door P. So in that case, Monty can reveal either G or S. So let's make a branch for that. Similarly, if they both pick P, we could have either P or S. And if they both pick S, then here, we could have either P or G.

Now, for the remaining ones, we could say that there's a random process there, where Monty is randomly picking from a single element set. So it's always going to be the same thing. We could consider it a random outcome, though. So we can also draw these arrows here.

Kind of running out of space here, so let's put this up here, and let's label it on this board. So Monty reveals a zonk. So this is the technical term for a non-prize. I think the non-prizes on Monty's show were called zonks. They weren't always goats. Sometimes it was some other ridiculous prize, like a garbage can or a potato or something. But, yeah, they were just called zonks.

So now, are we done? Is there any other randomness that we haven't modeled yet? Who thinks yes? Oh, OK. So what are we missing, then?

Everybody seems to think we're missing somebody. Or are we just asleep? It is very early in the morning, I understand.

Are we missing anything? Who thinks we're not missing anything? Well, a few hands, OK. Well, yeah, it turns out that this is all of the randomness.

You could think about a different problem, where Marilyn has a random strategy. But for now, we're just-- our question is, Marilyn is always going to switch. What's the probability that she wins? So this is all the randomness that we've got.

So now that we have our tree structure, we can label our outcomes. And that'll tell us what the set S is. So the set S is just going to be the set of leaves of our tree.

Each path corresponds to some way that the random process could turn out. Those are going to be our outcomes. So formally, we could label these as tuples. So let's say that this one is PPP. So that represents Monty hides the prize behind door P, Marilyn picks door P, and then Monty reveals a zonk behind door G.

This one's going to be PPS. This one's going to be PG. And now, I haven't actually labeled this edge. But if Monty hides the prize behind door P, and then Marilyn picks door G, where is Monty revealing the zonk? S, yeah. So let's label that, and we can continue all the way down.

So S-- oh, I've just realized that I've overloaded it. It's a script S in my notes, and so it looks different. I'm sorry. [LAUGHS] OK, so our squiggly S is going to be the set of all tuples, so x, y, z, in PGS-- oops-- cubed. So it's going to be triples of P, G, and S, such that z is not equal to x, and z is not equal to y.

So that's how we're mathematically going to represent our sample space. So x is the door that Monty hides the prize behind. y is the door that Marilyn initially chooses. And then z is going to be the door that Monty reveals, which has a goat. And by our rule, that door cannot be either of the first two, but anything else is a valid outcome.

OK, do we have any questions about the sample space and how we compute it?

So our next step, now that we've got our sample space, we want to figure out our probability function.

So we'd like to figure out what the probability of each of these outcomes is. So to do this, we're going to assign a probability to every edge that we've got in our tree diagram. And so that's just going to be the probability that if we get to the source of that edge, what's the probability that we proceed along that edge?

So at the very beginning, if we're at the root, what's the probability that, given that we're at the start, Monty hides the prize behind door P? Yeah?

AUDIENCE: A third.

BRYNMOR CHAPMAN: One-third, yeah. So we're going to label this with one-third. What about G? What's the probability that Monty hides the prize behind door G?

Anybody?

I mean, yeah, arithmetic is pretty difficult. But I'll just do it for you. It's also one-third. And the probability that Monty hides it behind door S, also a third. Yeah, question?

AUDIENCE: So they have to add up to 1--

BRYNMOR CHAPMAN: Yeah, that's a great observation. So all of these three probabilities should add up to 1, because if we're at this vertex here, we should proceed somewhere. So it wouldn't make a whole lot of sense to have probability 1/3, probability 2/3, and probability 1, because then you have to proceed down this path, but then you also have to proceed down one of the others. It's a bit sticky.

So given that you're at this vertex, you've got to reach a leaf somehow. You've got to get to some outcome. So you've got to proceed. So all of these edges should sum up to 1. Great observation.

It's not quite true, though, that that holds for the entire column. So what happens if we look at the next column? So given that Monty hides the prize behind door P, what's the probability that Marilyn then picks door P? Any ideas? Yeah?

AUDIENCE: One-third.

BRYNMOR CHAPMAN: One-third, yeah. So Marilyn picks this with probability one-third. Marilyn doesn't know or care what's going on with where the prize actually is. She's just picking each door with probability one-third, regardless.

So we've kind of got the same thing going on here. Each of these three edges has probability one-third. But now the same holds true for this branch, right? If Monty has hidden the door-- or hidden the prize behind door G, now the same logic applies. Marilyn is still picking each door with probability one-third.

So we've also got one-third here, one-third here, one-third here. So it's not quite that all of the edges in a column add up to 1, but all of the edges coming out of a single vertex add up to 1. And same thing for this bottom branch.

Now, what about the last column? If Monty has hidden the prize behind door P, and Marilyn has also picked door P, what's the probability that we get to outcome PPG? Yeah?

AUDIENCE: One-half?

BRYNMOR CHAPMAN: One-half, that's right. And it's also one-half-- whoops, that's not a 2. That's a 3. There we go, one-half. So probability one-half for each of them. We've got probability one-half that the zonk behind G is revealed, probability one-half that the zonk behind S is revealed.

And as we were saying before, if Monty hides the prize behind P, and Marilyn picks G, Monty now has no choice. So we're just going to say that this is probability 1. And same with this one-- this one also has one PSG. And similarly, for the other three branches, we've got 1, one-half, one-half, 1, 1, 1, one-half, and one-half. So do these edge probabilities make sense to everybody?

So now, remember what the probability function was. We have to assign a probability to every outcome. The outcomes are the leaves of our tree, not the edges. So what is the probability that we end up with outcome PPG, for instance?

Yeah?

AUDIENCE: It would be 1/18.

BRYNMOR CHAPMAN: 1/18. And how are you getting that?

AUDIENCE: It's one-third times one-third times one-third.

BRYNMOR CHAPMAN: Yeah, one-third times one-third times one-third. We're going to multiply together all of the probabilities along the path from the root to that leaf.

So Monty hides the prize behind door P with probability one-third. And then, conditioned on that happening, Marilyn then picks door P with probability one-third. So then there's probability one-ninth that we make it to this vertex here. And now that one-ninth is split again among these two outcomes.

Does that make sense to everybody? So if we do that for every outcome, we've got-- oops, 1/18, not 1/8, another 1/18 for PPS. Then what about PGS? Yeah?

AUDIENCE: It'd just be one-ninth.

BRYNMOR CHAPMAN: Yeah, it's just one-ninth. So same deal for the first two layers, but then in the last layer, we've just got a probability of 1. We don't have to split again, so it remains at one-ninth.

And once again, we can continue. We've got 1/9 here, 1/18, 1/18, 1/9, and we can continue. So we're going to justify this more formally later. But for now, just take it on faith that in order to get the probability of an outcome, you multiply together the probabilities along the path to that outcome. Hopefully, it at least makes intuitive sense. But we'll reason about it more formally later.

So now, we have our sample space, we have our probability function, so we've got the entire probability space that we're trying to reason about. The next thing that we want, we want to figure out what we call events we care about.

So an event is simply a subset of our sample space.

So it's a set of outcomes.

So up here, we could think about the event that Marilyn wins. So what is that formally as a set?

Yeah?

AUDIENCE: Subset of [INAUDIBLE].

BRYNMOR CHAPMAN: Yeah, exactly. So the answer was, it's the subset of events where x and y are different. So the notation that we usually use for events is the following.

We put the description of the event in square brackets. So this is just going to be the set P, G, S, P, S, G, G, P, S, G, S, P, S, G, P, S, P, G. So the 6 outcomes where Marilyn and Monty originally picked different doors, and then Monty is forced to reveal the last one.

Similarly, we could also define other events like Monty opens door P. So this is going to be the set G, G, P, S, S, P, G, S, P, and S, G, P. So does this make sense to everybody? Any subset of events-- or, sorry, any subset of outcomes is called an event. We can take our English description and then look at which actual outcomes this corresponds to.

So what we want to do now is figure out what event we care about, or possibly what events. So what are the events that we care about for this problem? Anybody?

What was the probability that we're trying to compute? Yeah?

AUDIENCE: [INAUDIBLE] when Monty loses-- or Marilyn loses.

BRYNMOR CHAPMAN: Pardon?

AUDIENCE: When Marilyn loses.

BRYNMOR CHAPMAN: Yeah. So the events Marilyn wins, and Marilyn loses. These are the events we care about. So we could say Marilyn loses is the complement of Marilyn wins, which we denote with an overbar. So these are the two events we care about.

So the final step is to compute the answer. We've identified the event that we care about. So now we want to figure out the probability that that event occurs.

So we can extend the probability function two events. So if E is an event, we can think about the probability of an event just as we think about the probability of outcomes. And we're going to define this to be the sum over all outcomes omega in E of the probability of omega.

So hopefully, intuitively, that makes sense. All of these outcomes, exactly one outcome is going to happen. So if we're trying to figure out the probability that one of a set of outcomes happens, it's just going to be the sum of the probabilities that each individual one occurs.

So we know the probabilities of each of the outcomes. That's what we just computed here. So now we can use that to compute the probability of the events we just talked about. So what is the probability that Marilyn wins?

So we have our set of outcomes here. What's the probability of each one?

Now, what's the probability of outcome PGS? Yeah?

AUDIENCE: One-ninth.

BRYNMOR CHAPMAN: One-ninth, yeah. That one's written up there explicitly. What about the probability of PSG? Same, one-ninth.

What about the probability of GPS? Also a ninth. All of these are symmetric. What's that, 2/3? So we've got six events-- or six outcomes, each of which occurs with probability one-ninth. So if you add those all together, you get 2/3.

So this seems to support what we found empirically. Marilyn wins with better than probability one-half. The probability that she loses is going to be the complement of this.

So it's going to be the sum of the probabilities of all of the other outcomes. And we said that all probabilities fit out to 1, so this is going to be the complement with respect to 1.

So as it turns out, Marilyn was correct. We didn't have to rely on intuition or anything like that, no ingenious analogies, et cetera. It's just arithmetic over Q.

And I think the hardest part of this for many people is resisting the temptation to jump to some obvious conclusion. It might be tempting to get to the end of the tree diagram and be like, oh, we've got six outcomes where Marilyn wins, six where she loses, So probability 6/12 is 1/2. So if we actually set all of that intuition aside, go through everything formally, it's all quite simple and straightforward. Does that make sense to everybody?

So let's move on to a different example. Suppose I have three strange dice. So excuse my drawing-- whoops-- there. OK, so I've got one die that has 2 written on two of the faces, so top and bottom, 6 written on these two faces, and 7 written on those two faces. Hopefully, that looks vaguely cubical.

I've got a second die, which is similar. But instead of 2, 6, and 7, I've got 1, 5, and 9. Then I've got a third one, which has 3, 4, and 8.

So consider the following game. I pick a die. Now you pick one of the other dies-- or dice-- after seeing what I pick. And we each roll our dice. And whoever gets the higher number wins.

Who should win? Should I win, or should you win? Who thinks I win? OK, who thinks you win? Who thinks we should just do the tree diagram and figure it out? Yeah, good call. OK.

So let's call these dice R, G, and B. Yeah, let's draw out our tree diagram. So what happens if, say, I pick the red die, and you pick the green die? Who should win? Let's see.

So we've got-- when I roll my die, there are three possibilities. I could either get a 2, a 6, or a 7. And now you roll the green die, and you could have either a 1, a 5, or a 9-- not minus 5, 5-- 1, 5, 9, 1, 5, 9. And now what is the probability-- or, sorry, what are the outcomes?

We have 9 leaves, so 9 outcomes. How are we going to represent them? Well, we can do it basically the same as how we represented them up here. We can just look at a tuple of edge labels.

So our sample space is going to be the set 2, 6, 7, a Cartesian product with 1, 5, 9. So this is going to be ordered pairs where the first element is one of these three things, and the second is one of these three things. So this one, for instance, is 2, comma 1. This one is 2, comma, 5, et cetera.

Now what are the probabilities on each edge? Yeah? One-ninth. So those are the probabilities on the outcomes.

Yeah, one-third on each edge. So when I roll my die, I've got a one-third chance each of getting 2, 6, or 7. When you roll yours, you've got a one-third chance each of getting 1, 5, or 9. And it's going to be the same for the other two branches. So when we multiply together the probabilities along paths to leaves, we get one-ninth for every outcome.

So when you have a probability space like this that's particularly nice, this is what we call a uniform probability space. So let's say S, comma, Pr-- so this is our probability space-- is uniform if Pr is constant. So in this case, it's a constant function that maps every outcome to a probability of one-ninth. So it's a uniform probability space.

And uniform spaces are particularly nice because we don't even have to worry about the addition that we were doing before. In order to figure out probabilities of events, you just count up how many outcomes are in the event and then divide by the total number of outcomes. So it just becomes a counting problem.

So using that, what's the probability that red wins? Yeah?

AUDIENCE: 5/9.

BRYNMOR CHAPMAN: 5/9, yeah.

So let's break this down a bit. So we're looking at the probability of the set, 2, comma, 1, red wins; 6, 1, 6, 5, red wins; and 7, 1, 7, 5, red wins. So 2, 1, 6, 1, 6, 5, 7, 1, 7, 5. We have five of these. So that's 5/9.

Now, what if instead we had red versus blue instead of red versus green? So if instead of red versus green, we have red versus blue, this column now gets replaced with 3, 4, 8 instead of 1, 5, 9. So our tree diagram is going to look like the following.

So we've still got 2, 6, 7 here. But this time, this is going to be 3, 4, 8-- 3, 4, 8, 3, 4, 8. And I'm going to omit the probabilities because, as we just established, it's uniform. So it's just a counting problem.

But let's write down the outcomes.

So in how many of these outcomes does red win? Well, not this one, not this one, not this one. So if red rolls 2, all of blue's numbers are going to be higher, so blue is always going to win. If red rolls 6, red wins if blue rolls 3 or 4; loses if blue rolls 8, and same for 7.

So we have four outcomes where red wins. So red wins with probability 4/9.

So we just saw that red beats green but loses to blue. What does that tell us about blue and green? If I pick blue, and you pick green, what happens? Who wins? Yeah?

AUDIENCE: Green.

BRYNMOR CHAPMAN: Green. Why?

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: OK, so green beats red. So I think we just established that red beats green, right? So here, we've got red beats green. Here, we've got red loses to blue.

So if we have green and blue, who thinks that green should win? Who thinks that blue should win? Who has absolutely no idea? Who wants to do the tree diagram?

Oh, still pretty enthusiastic about that, even though we've got some intuition. That's great. Intuition is bad. Discard the intuition.

OK, so if we've got green versus blue, so this can be the green roll here. We've got 1, 5, 9. Blue roll can be 3, 4, 8, 3, 4, 8, 3, 4, 8. So in which of these does green win?

Well, if green rolls 1, then green is screwed. So all of these are bad for green. If green rolls 5, well, 5 will beat 3 and 4, but not 8. So green wins these two. And if green rolls 9, then green beats whatever blue rolls.

So it looks like green is winning in five of these.

So this should be like that. So red beats green, green beats blue, and blue beats red. What's gone wrong?

Yeah?

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: Pardon?

AUDIENCE: There is only a hierarchy [INAUDIBLE].

BRYNMOR CHAPMAN: Yeah. So nothing's gone wrong. All that's gone wrong is I kind of expected that there would be a linear order on them. If red beats green and green beats blue, it kind of intuitively makes sense that red should beat blue, right?

No. No, no intuition. Just do the tree diagram. So yeah, this is perfectly fine, right?

And, in fact, it gets even weirder. What do you think happens if we each choose two dice-- or, sorry, we each roll twice? So what if I pick the red die; I roll it twice; you pick the green die; you roll it twice? Red probably beats green, so is it true that two red rolls probably beat two green rolls if we sum them up?

Who thinks yes? Who thinks no? Who wants to draw the tree diagram? Oh, wow, so eager. I don't. [LAUGHS]

What is the tree diagram going to look like if we do that? I'm rolling twice. You're rolling twice. Each roll can have three possibilities. So we're going to end up with a pretty massive tree.

I don't want to draw that, at least not on the board. Does somebody else want to come up and do it? Please? No, nobody?

OK, well, instead of drawing out the entire tree diagram, let's start drawing it and see if we can find a pattern. So this is R2 versus G2. So let's start by drawing as we did before. So my first roll, that was red. It could be 2, 6, or 7. Second roll could also be 2, 6, or 7.

Now, for green, we're basically just going to be drawing the same tree nine times, so let's not. So green could roll 1, 5, 9, and then can also roll 1, 5, 9 for the second roll. And rather than duplicating this tree nine times, I'm just going to leave this as it is. And we can see if we can find a pattern. Does that make sense to everybody?

So an outcome now is going to be a 4-tuple, instead of a pair. And in some sense, we can think of it as combining a leaf of this tree with a leaf of this tree. We're going to take a red leaf and a green leaf. Question? Or no? No, OK.

So we can then abbreviate it even further. Instead of looking at the pair of numbers that we've rolled, let's instead look at just the sum. So if red rolls two 2's, we've got 4 here. 2 and 6 is 8. 2 and 7 is 9. 6, 2 is 8. 6, 6 is not 9. That's 12. Never trust me to do arithmetic.

6, 7 is 13. 7 and 2 is 9. 7 and 6 is 13. 7 and 7 is 14. And if we do the same thing for green, we have 2, 6, 10, 6, 10, 14, 10, 14, 18.

So now we can start counting outcomes. If we take this leaf of the red tree, how many outcomes are there where red wins? If red rolls four in total, the only way that red wins is if green gets snake eyes. So we've just got one.

If red rolls 8, which could be either of these two, how many outcomes are there where red wins? Well, we've got the 2 here, a 6 here, and a 6 here. So we've got three. And we've got two red leaves that have 8.

What about 9? It looks like we've got two more that have 9. Same deal-- we don't have any 8's or 9's over on this side.

What else do we have? 12. OK, well, if red rolls 12, then red is good if blue rolls 2, 6, 10, but not 14 or 18. So we've got 1, 2, 3, 4, 5, 6 outcomes. And we have one red leaf where red has 12.

And then same thing for 13-- there are no 12's or 13's over here, and we have two 13's.

Now, what about 14? Well, if red rolls 14, then same deal. But we've got these two ties here. So red wins. Draw, we've just got two. And then red loses is going to be the complement.

So has anybody added all of these up yet? I haven't. 5, 30, 42, 43, is that right?

AUDIENCE: Yes.

BRYNMOR CHAPMAN: OK, so 43 outcomes for red wins.

Is something wrong? This doesn't match my notes. Let's see.

So we've got 1 for this subtree. For each of these, we have 3. For each of the 9's, we also have 3, so that's those, 12, 13, and 13.

Oh, do I just have too many terms? I just have too many terms. OK, not 43. It should be 37. Is that right? Yeah, or 12, 13, 13, and 14 each have 6.

AUDIENCE: Yes.

BRYNMOR CHAPMAN: OK. OK, yeah, never trust me to do arithmetic. It's always a bad plan.

OK, so red wins in 37 of the outcomes. There's a draw in 2 of them. And how many outcomes are there in total? Yeah?

AUDIENCE: Is that 81?

BRYNMOR CHAPMAN: Yeah, 81. So red loses in 81 minus 37 minus 2, which is that 42? I think that's 42. So because this is a uniform space, we have 37/81, 42/81, which is going to be-- well, actually, let's not reduce it. It's easier to see like that. And 2 out of 81 for a draw.

So if we only had one roll, red had the advantage. Red wins with probability of 5/9. If we have two rolls, the probability that red wins is smaller than the probability that red loses. So now suddenly, green has the advantage instead.

And as it turns out, the same thing happens with all three pairs of dice. So red beats green with one roll. Green beats blue with one roll. But blue beats green with two rolls. And blue beats red with one roll, but red beats blue with two rolls.

And yeah, so the dice continue to not have this transitivity property. But every time, it switched. And as it turns out, you can even extend this. Like if you have more dice with different numbers, you can get all sorts of weird things that could happen.

So there's a link to a rather interesting paper on arXiv if you want to see it in the lecture notes. But yeah, for now, I think that's all we have time for. Feel free to come up if you have questions. And otherwise, we will see you on Thursday.