[SQUEAKING]

[RUSTLING]

[CLICKING]

BRYNMOR CHAPMAN: Hello, everybody. Can people hear me in the back? Is it working? Cool.

OK, so I believe that last time, we started counting, and you saw the sum rule in particular. So let's take a look at a quick example.

So I've got a very standard deck of playing cards here, as you can see. How many kings and queens are there?

AUDIENCE: Eight.

BRYNMOR CHAPMAN: Eight. Why is that?

AUDIENCE: There's four kings and four queens.

BRYNMOR CHAPMAN: Yeah, so the answer was that there are four suits for each rank. We've got, therefore, four kings, four queens. Nothing is both a king and a queen. So you can just add them together. You've got eight cards that are either kings or queens.

What if the sets are not disjoint? So what if I instead wanted to count the number of cards that are either hearts or face cards? So then what do I do?

So I've got 12 face cards here. My hands are a little bit too small for playing cards, but hopefully, you can roughly see. How many hearts are there? Yeah, 13 hearts, so 13 hearts in one hand, 12 face cards in the other, so 25?

No? What's gone wrong? Can anybody tell me what's gone wrong? Yeah.

AUDIENCE: You've got some cards that are both hearts and face cards.

BRYNMOR CHAPMAN: Yeah, I've got two kings of hearts here, so I'm cheating. So what do?

AUDIENCE: Well, since you're double-counting the cards that are both hearts and faces, you can subtract off one of them.

BRYNMOR CHAPMAN: Yeah, exactly. So the answer was, I've got three cards that are both hearts and face cards. So these are the three cards that I've counted twice. I've got two each of these three cards.

So if I just discard one copy, then I'll be fine. So I subtract off everything that's in the intersection. So I had 25 before, subtract off those three, and I end up with 22 that are either hearts or-- well, these cards are too big. Maybe I should have actually brought standard cards.

I've got 22 cards that are either hearts or face cards. Does that make sense to people? So this is what we call the inclusion-exclusion rule or the principle of inclusion and exclusion, so often abbreviated as the PIE for Principle of Inclusion-Exclusion.

So if I have 2 sets A and B and I want to compute their union, the size of this union-- ew, well, this chalk is terrible. Sorry-- the size of their union is going to be the sizes of the individual sets minus the size of the intersection. Does that make sense to everybody? It may be a little bit easier to see if we draw a picture.

So hopefully, people remember Venn diagrams from grade school. Yeah? Maybe? No? Well, this is a Venn diagram, a couple of overlapping circles.

So maybe we've got hearts on one side, and we've got faces on the other. But then we've got some things that are in the middle. So this can be our set A. This can be our set B. And here, we've got king of hearts, queen of hearts, and jack of hearts.

So if we want to compute the size of the union, we're going to take the things that are only on the heart side, the things that are only on the face side, and then the things that are in the intersection, we only want to count them once. But that's counted in both of A and B. So we're going to add together the sizes of A and B and then subtract out the middle Does that make sense to everybody?

So what happens if we have three sets? Any ideas? Well, can we do something similar? So let's draw a three-way Venn diagram this time.

So we've got A here, B here, C here. So is it true that A union B union C is equal to the sum of the individual sets minus the pairwise intersections? Oops, that should be a B. Is this true? Yeah.

AUDIENCE: [INAUDIBLE] A and B, A and C, B and C [INAUDIBLE] overlap.

BRYNMOR CHAPMAN: Yeah. So your classmate's observation was that these intersections themselves overlap. So if we think about how many times we've counted each item-- everything here, we've only counted once.

Everything here, we've only counted once. Everything here, we've only counted once. That's fairly easy.

What about the things in these kind of weirdly shaped segments? So those appear in two of the individual sets and in one of the pairwise intersections. So we've also counted these ones.

But now we've got a problem in the middle. How many times have we counted this? Well, we've counted it once for A, once for B, once for C, and then we've subtracted it once each for A intersect B, A intersect C, and B intersect C. So we've counted this zero times in total.

So what we actually want is to add back in the three-way intersection, and that gives us the correct formula. Does that make sense to everybody? Yeah?

So let's take a look at an example then. How many numbers k in the range from 1 to n are coprime with n? How could we figure this out? And those of you who already know the totient function, don't invoke that.

How could we use the principle of inclusion-exclusion to compute this? Any ideas? Tentative hands? Maybe? Yeah.

AUDIENCE: I guess you can do the prime factorization of n and then check what numbers within-- well, I don't know the specifics but basically check which numbers are multiples of a specific factor and then count which are the ones that share factors between them.

BRYNMOR CHAPMAN: So the idea was look at which numbers share a prime factor with n. So for simplicity, let's assume for now that n is the product of, say, three primes, p, q, and r. So then how could we figure out how many things share a factor of p? Yeah.

AUDIENCE: The quotient when you divide n by p.

BRYNMOR CHAPMAN: Yeah. So there are going to be n/p that share a factor of p. So let's say that-- let's call this set A sub p, the set of numbers-- the subset of these numbers that share a factor of p. So there are going to be n/p of those.

What about the things that share a factor of q? Same thing. So similarly, we can define A sub q to be the set of numbers that share a factor of q. And then we've got this, and similarly, A sub r is going to be the set of things that share a factor of r.

Now, what exactly are we trying to compute? What we really want is the complement of the union of these three sets. So we're trying to compute the size of n minus Ap union Aq union Ar.

So now we can use inclusion-exclusion. We're trying to compute the union of this set. That's going to be the sum of the sizes of the individual sets minus the-- oh, I did this the wrong way around, didn't I? Oh, no, I didn't. That was the right one.

So the sum of the individual sizes minus the sum of the pairwise intersections plus the sum of the three-way intersection-- so what are the sizes of these intersections? What's the size of the intersection of Ap and Aq? Yeah.

AUDIENCE: n over pq.

BRYNMOR CHAPMAN: Yeah. So the answer was n over pq. So p and q are distinct primes. If something shares both the factor of p and a factor of q, it has to be a multiple of pq. So there are n over pq of those.

Similarly, we have n over pr and n over qr. What about the three-way intersection? How many things share a factor of p, q, and r? Yeah.

AUDIENCE: That's just n.

BRYNMOR CHAPMAN: Yeah, it's just n itself, so minus 1. And if you do some messy algebra, this turns out to p minus 1, q minus 1, r minus 1, which is also equal to n times 1 minus 1/p, 1 minus 1/q, 1 minus 1/r.

So hopefully, it's not too hard to see how this could generalize to get the more general form of Euler's totient. Does this application make sense to people? Any questions? Yeah.

AUDIENCE: [INAUDIBLE]?

BRYNMOR CHAPMAN: Oh, that's just the assumption that we made for the sake of example. So that's inclusion-exclusion with two or three sets, but more generally, you can continue the pattern.

So if you have, say, n sets A1, A2, all the way up to A sub n and you want to figure out the union of all of these, so as we did before, we're going to add up the individual sizes of the individual sets. Then we're going to subtract off the pairwise intersections.

Then we're going to add in the three-way intersections, and we're just going to keep on going until eventually we get to plus or minus the size of the n-way intersection, so A1 cap A2 cap all the way to A sub n. And the sign on that last term will depend on the parity of n. Does that make sense to people?

So perhaps a more formal way of expressing this is the following. So suppose we have some set u, which we're going to say is a union of n sets. So this is going to be our universe of discourse. And so this means that by convention, if we intersect no sets, this is just going to be u.

So just like if we union no sets, we start with nothing and then start unioning in more things, and we should get a union of sets. If we start with everything and then start intersecting, we should get an intersection of sets. So that's where this convention comes from. Yeah.

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: Oh, this notation? This notation? Both?

AUDIENCE: Both [INAUDIBLE].

BRYNMOR CHAPMAN: Oh, OK. So this is kind of like the big sigma. So just how this represents like an n-way addition, this represents an n-way union, and this represents an n-way intersection.

So we've got-- this is saying we've got a union of n sets, like A1 union A2 union A3 union An. And this is saying we've got an intersection of n sets, but n is zero. So by convention, we just take that to be the entire universe.

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: Sorry?

AUDIENCE: What was the [INAUDIBLE] intersection [INAUDIBLE]?

BRYNMOR CHAPMAN: The question is, why are we setting both the union and the intersection to be u? This is just how we're defining u. This is a convention.

So if you intersect no sets, you should get everything, just like if you multiply no numbers, you should get 1. It's kind of similar to that. You start with everything, and then you intersect one set. You should get the set that you intersected in.

So this is just a convention. If you prefer, you can just write in u itself. And what I'm about to write that might make it easier to understand. Sorry about that.

So the theorem is that if you look at all subsets of the numbers up to n and you take the i-way intersection of these Ai's and then multiply it by minus 1 to the size of i, this should give you 0.

So now, if you want to think about this in the case of the two-way inclusion-exclusion here-- so this is saying, 0 equals-- we move the union over to the right-hand side, or I guess we're moving everything over to the left-hand side based on what I've written. Yeah.

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: So this formula is basically what we were doing here. We're taking i-way intersections, and we're flipping the sign on successive sizes of i. And we're just kind of moving this union over to the other side, so we get zero.

So if that formula is too much, you can think about the previous formula. But yeah.

AUDIENCE: [INAUDIBLE] i so big I would [INAUDIBLE].

BRYNMOR CHAPMAN: So the question was, what is the difference between big and little i? So big I is a set of indices. Little i is an individual index in that set. So we're summing over all subsets of the numbers up to n, and then for each of those subsets, we're taking the intersection of the relevant Ai's.

So over here, we've got A is just-- it corresponds to the subset just 1. B corresponds to the subset 2. A intersect B corresponds to the subset 1 comma 2, and A union B corresponds to the empty subset. So these are all of the subsets of the integers 1 and 2. Does that make sense? Question?

AUDIENCE: Can you use I [INAUDIBLE] for negative 1? Why is it absolute value?

BRYNMOR CHAPMAN: Size of I, because I is a set. Yeah, so the question is, why do we have absolute value here? That's the cardinality.

So how could we prove this? Well, let's look at some x. So say that we've got some element x of our universe. How many times do we count x on this left-hand side? How many of these terms does it appear in?

Well, it's going to be in all of the terms where x is in all of these Ai's. So if we say I sub x is the set of indices with x in Ai, it's going to be in the I sub x set. It's also going to be in every subset of I sub X.

Does that make sense? Do people see why? If x is in sets 1 and 2, it's going to be in the intersection where we've got I is 1, 2. It's also going to be in the intersection where I is 1 or where I is 2 or where I is empty.

But now, how much does it contribute to the left-hand side for each of these? Well, how many subsets of I sub x have even size, and how many have odd size? Yeah.

AUDIENCE: Half.

BRYNMOR CHAPMAN: Yeah, it's half each. So for the ones of odd size, we're contributing minus 1 to the left-hand side. For the ones of even size, we're contributing positive 1 to the right-hand side-- or to the left-hand side. So because they're half and half, those all cancel.

So x contributes plus 1 to the left-hand side for each even-size subset. I sub x, and it contributes minus 1 to the left-hand side for each odd-size subset. So when we add all of those together, we get 0. So every element of our universe contributes 0 to the left-hand side, so we get 0. Does that make sense to people? Yeah.

AUDIENCE: Both the left-hand side?

BRYNMOR CHAPMAN: Yes. We're looking at the-- sorry, question was, are those both supposed to be left-hand side? We're looking at what x contributes to this sum. How many of these terms does it contribute positively to? How many does it contribute negatively to?

So whenever the size of I is even, it's going to contribute positively if it's contained in that intersection. Whenever the size of I is odd, it will contribute negatively when it's in the intersection. Yeah.

AUDIENCE: How did you get the empty set? How do we construct the big I so that [INAUDIBLE]?

BRYNMOR CHAPMAN: Sorry. The question was, how do we construct a big I so that the intersection is the empty set?

AUDIENCE: Yeah, I may have understood it wrong, but [INAUDIBLE] intersection big U and then subtract it off to get [INAUDIBLE].

BRYNMOR CHAPMAN: Let me see if I'm understanding. You are trying to find an I such that this is empty or such that this is empty.

AUDIENCE: Second one.

BRYNMOR CHAPMAN: This one. So we don't care. We are looking at the sets that contain x, so they should never be empty. We're looking at what we had here. How many times have we counted x?

So if this is still confusing, please review the lecture notes. Or it's not important that the proof is not important. The most important thing is knowing how to use it. So you want to add together the single sets, subtract off the pairwise intersections, add together the three-way intersections, cetera.

So let's move on to the second topic for today, pidgeyhole principle. Has anybody heard of this before? Yeah.

So the pidgeyhole principle-- it gets its name from medieval times. You have a bunch of pidgeyholes that are supposed to contain pidgeys. But what if you have more pidgeys than pidgeyholes? What happens? Well, two of them have to share. Two of them get squished.

So that's basically what the pidgeyhole principle is saying. If we have two sets A and B where the size of A is greater than the size of B and f from A to B is total, then f cannot be injective.

So if is A set of pigeons, B is a set of pigeonholes, and f is saying, pigeon A lives in pigeonhole B, this should be a total function. Every pidgey should have a home. We don't want to have homeless pidgeys. So it's a total function, but that means that it can't be injective, meaning that there have to be two pigeons that share a pigeonhole.

So let's look at a couple of examples. Do we still have 26 students in the room? Nice. Not too many students have dropped. So there are more than 26 students in the room, so at least two of you share the first letter of your first name.

So we've got, well, at least 27 students and 26 letters of the alphabet. So if we map every student to the first letter of their first name, that's a total function, so it can't be injective. So two students must share. Does that example make sense?

What if I've got n differently colored pairs of socks? Don't worry about it. How many socks do I need to pull out of my drawer to ensure that I get two of the same color? Sorry?

AUDIENCE: Plus 1.

BRYNMOR CHAPMAN: Yeah, the answer was n plus 1. So in my case, two. So if I've got n socks-- or, sorry, n colors of socks-- I need n plus 1 to guarantee a pair.

So I've been a little bit hand-wavy with this one. When you're trying to prove something with the pigeonhole principle, you generally need to be very careful about what your pigeons are, what your pigeonholes are, and what the map between them is. So what are those three elements in this case?

Maybe we could phrase it instead as, if I have n plus 1, then it guarantees a pair. Yeah.

AUDIENCE: And the [INAUDIBLE]

BRYNMOR CHAPMAN: So A is the set of socks, and B is a set of colors. And the map is going to map the sock to its color. Does that make sense to everybody?

Does anybody know what the population of Boston is?

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: Are you looking at the lecture notes?

AUDIENCE: Yeah.

BRYNMOR CHAPMAN: Oh, OK. Yeah, about 650,000 as of however many decades ago these lecture notes were written. I don't actually know that was very long ago. It could be current.

Is anybody a human biology nerd? Does anybody know how many hairs humans have? So it turns out humans have less than about 200,000 hairs.

So what does the pigeonhole principle tell us here? Any ideas? Getting some facepalms. Yeah.

AUDIENCE: That at least three humans have the same number of hairs.

BRYNMOR CHAPMAN: Oh, good point. So the answer was at least three humans in Boston have the same number of hairs. So as stated, all we know is that at least two Bostonians have the same number of hairs. So does that mean that there's somebody with the same number of hairs as me? Not necessarily.

So we know that there exists two Bostonians. We don't know who they are. This is non-constructive. It only tells us of existence.

But your colleague also brings up another good point. Because this 650,000 is so much larger than the 200,000, we can actually say something stronger. All we need is 200,000 plus 1. That'll give us a collision of two Bostonians.

We've actually got more than three times that number. So there are more than three Bostonians who have the same number of hairs. So in fact, we can conclude that there are at least four Bostonians who all have the same number of hairs. Does that make sense to everybody? Yeah.

AUDIENCE: How did you get the four [INAUDIBLE]?

BRYNMOR CHAPMAN: Ah, so the question was, how did we get the four? So this is what we call the generalized pigeonhole principle.

So up here, we just said, if A is larger than B, but if we have the stronger condition that if A is larger than k times B, then we've got a k-plus-1-way collision, not just a two-way collision.

So because we have 650,000, which is more than 3 times 200,000, we've got a 3-plus-1 way collision, not just a three-way collision. Does that make sense? Yeah.

AUDIENCE: So [INAUDIBLE] in this example, the number of possible hairs you might have [INAUDIBLE] every integer or natural number from 1 to [INAUDIBLE]?

BRYNMOR CHAPMAN: Yeah. I suppose we could include 0, but I guess everybody knows that there are bald Bostonians. So maybe we exclude that. But yeah, every natural number from 1 up to 200,000 is going to be our B, and our A is going to be the set of all Bostonians. And we are mapping a person to the number of hairs that they have.

So you can basically always use the generalized pigeonhole principle if you would-- just like strong induction will always work whenever weak induction does. You could always set k equal 1.

But the proof is basically the same, and if you use these to prove things about functions, the proofs can often be very short, very simple. But they can also sometimes require some cleverness. So it may not always be obvious what A, B, and f should be.

So here's a slightly more involved example. So we have an 8 by 8 chessboard and 33 rooks. Does everybody know how chess works, vaguely? So rooks can move along ranks or files, ranks or files. Sorry.

So we would like to place these rooks so that any set of five-- they don't all attack each other. Or no, I phrased that wrong. So we want to find five of them where none of them attack each other. So basically, we want them all to be in distinct rows and files.

So if we've got something like this, if we have a rook here, a rook here, and a rook here, those three do not attack each other. So we're trying to find a set of five like that. No two of them share a column or a row.

Does the setup make sense? How could we do that? Rather, how could we prove the existence of such a set? Maybe it would help to draw out the chessboard. Let's see. 1, 2, 3, 4, 5, 6, 7, 1, 2, 3, 4, 5, 6, 7. There's our wonky chessboard.

So how could we do this? Any ideas? Yeah.

AUDIENCE: Maybe [INAUDIBLE] the five pieces. So that's 32 [INAUDIBLE] we have 33 [INAUDIBLE].

BRYNMOR CHAPMAN: So the answer is, if we place four rooks, they cover at most 2 to the 5 squares? Is that true? I don't think that's true. If we place rooks in these four squares, they're going to cover 3/4 of the chessboard. Yes.

AUDIENCE: I'm just curious. Do the [INAUDIBLE] have to not attack each other, but their position relevant to the other 28 doesn't matter?

BRYNMOR CHAPMAN: That's right.

AUDIENCE: OK. This is a simple problem.

BRYNMOR CHAPMAN: Pardon?

AUDIENCE: If there's only five that need to be distinct and we have eight rows and eight columns, then it's pretty easy to determine.

BRYNMOR CHAPMAN: So what is your solution then?

AUDIENCE: You can literally just place five of them diagonally, couldn't you?

BRYNMOR CHAPMAN: Oh, oh, sorry, I think you may have misunderstood the problem. We have 33 rocks. We're trying to find a set of five that don't attack each other. We're given the placement. We're not allowed to just place five rocks. If we could place them, then yeah, we could place them along a diagonal, and then we've got eight.

So some of these have rooks. We don't get to control where. And we're trying to find some subset of these, some subset of five of these that don't attack each other. Does the problem setup make sense? Maybe? Yes? No? Yeah?

So how could we use the pigeonhole principle here? Well, perhaps let's start with the easiest element. What should our pigeons be? What are we trying to make collide? Yeah.

AUDIENCE: Rooks.

BRYNMOR CHAPMAN: Rooks. Yeah. So our pigeons are going to be the rooks. And I've lost my eraser.

Now, what should our pigeonholes be? We want to set it up so that if we get a five-way collision, those five rooks do not attack each other. Yeah.

AUDIENCE: [INAUDIBLE] the rooks are being placed on the board, and the pigeonholes should be the squares that are on the chessboard.

BRYNMOR CHAPMAN: Yeah. So the answer was, we're placing these rooks on the chessboard, so let's make the pigeonholes the squares of the chessboard. So you are off to a good start. That's a good idea.

The problem right now is that we've got fewer rooks than squares on the chessboard. So what we really want is a five-way collision, so we want to ensure that we have more than four times as many rooks as pigeonholes. So is there a way that we can modify your solution to make fewer things where collision is still good for us? So can we subdivide our squares so that we have-- yeah.

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: Sorry?

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: The answer was, squares that don't have rooks. That's not going to be very helpful. We're assigning rooks to pigeonholes. So the map should map rooks to occupied squares of some sort.

But how could we maybe split up the chessboard into eight pigeonholes? If we have eight pigeonholes, 33 rooks, that's more than four times as many. So we get a five-way collision. What could our eight sets of squares be? Yeah.

AUDIENCE: Columns.

BRYNMOR CHAPMAN: Columns. So if we have five rooks in a column, is that helpful? All of those rooks attack each other then. You're getting close, though. Yeah.

AUDIENCE: It'd be like diagonals [INAUDIBLE] offset by one.

BRYNMOR CHAPMAN: Yeah. So the answer was diagonals but then offset them. So I'm going to number of these squares. So everything numbered 1 is going to be our first pigeonhole. So if we get five rooks along the main diagonal, all of them are going to be in different rows and different columns, so they won't attack each other. So that's good.

Number 2 is going to be offset by one, but then we're also going to wrap around down here. So once again, we've got eight squares that are all in different rows and different columns. So if we get five on two squares, they also don't attack each other.

And we can continue in this way. Three goes along here and then these two squares, et cetera.

So now we have eight pigeonholes, 33 rooks, so that's more than four times as many. We get a five-way collision. Whichever number of square has that five-way collision, those are five rooks that don't attack each other.

Does that make sense to everybody? So it took a little bit of creativity to come up with the pigeonholes here, but once we have that, the proof is fairly straightforward.

So let's change gears a little bit and move on to the last counting technique for today, which we call combinatorial proofs or double counting. Oops.

So I'd like to think of this as somewhat similar to-- I believe you've seen the bijection rule earlier. So if you have two sets and you can come up with a bijection between them, then those two sets are the same size. Well, similarly, if you have two sizes in a single set that shares those sizes, those sizes must be equal.

So basically, this is a proof technique where we take a set and count it in two different ways, and that can give us some nice combinatorial identities. So for instance, we have this identity. Have people seen this one before? No?

So if you sum up n choose k over all k from 0 to n, you just get 2 to the n. So if you haven't seen it before, it's perhaps a little bit counterintuitive. We've got a summation of a bunch of things involving factorials, but it turns out to be some nice exponential.

So how could we prove this? How could we come up with a set that we can count both as this expression and as this expression? Any ideas? Yeah.

AUDIENCE: [INAUDIBLE] any length n binary sequence.

BRYNMOR CHAPMAN: So the solution was count length n binary sequences. So how does that give you each side of the equality?

AUDIENCE: On the right side, [INAUDIBLE] a 1 or a 0 in each position. [INAUDIBLE] can imagine, well, first of all, how many are there that have 1 and 1 [INAUDIBLE] be like, choose one. How many others have [INAUDIBLE] choose two. And you have to consider all sequences that have any number of 1's in there [INAUDIBLE]

BRYNMOR CHAPMAN: Exactly. Yeah. So the right-hand side is fairly straightforward. If we're counting length n binary sequences, every position, you can either have 0 or 1. So product rule gives us 2 to the n.

And for the left-hand side, we're going to partition based on how many 1's there are in our sequence. So our sequence could have any number of 1's from up to n inclusive. And all of these are disjoint. If it has three 1's, it can't also have eight 1's.

So by our sum rule, we can just sum up all of the sizes of those sets. What is the size of the set of sequences that has, say, three 1's? Anybody? How many binary sequences of length n have three 1's?

AUDIENCE: n choose 3.

BRYNMOR CHAPMAN: n choose three. Yeah. Or more generally, there are n choose k with k 1's, and k from 0 to n form a partition of our set. So if we sum up all of these binomial coefficients from 0 to n, we're also going to get the number of length n binary sequences. Does that make sense to everybody?

Does the proof technique make sense to everybody? We're taking a single set, and we're counting it in two different ways. And each way of counting gives us a different expression. So those two expressions must be equal. So the key here is choosing the right set to count.

So as it turns out, that's actually a special case of a nice theorem. So we've been using binomial coefficients quite a lot, these n choose k terms. Does anybody know their eponym? Any ideas?

Has anybody encountered the binomial theorem before? No? Well, let's take a look.

So this here is what we call a binomial, a sum of two monomials. And the binomial theorem says that if we take a binomial, raise it to the n-th power, then we get basically the same thing as we had on the left-hand side over there, but we've got these powers of x and y. And by convention, 0 to the 0-- we just say that that's equal to 1.

So the proof-- does anybody know how we might prove this? It's fairly similar to what we just did. Well, what happens if we expand out the left-hand side? What do we get? Yeah.

AUDIENCE: You get coefficients of Pascal's triangle [INAUDIBLE].

BRYNMOR CHAPMAN: So the answer was, we got the coefficients of Pascal's triangle. You've skipped a couple of steps. More immediately, what do we get?

Say, if n equals 2, what do we get? We get xx plus xy plus yx plus yy. We've basically got a bunch of binary sequences of x's and y's. So this gives us a sum of 2 to the n terms of the form. And I'll call it b1, b2, up to bn where each b is either x or y. So now how do we get Pascal out of that?

AUDIENCE: You write [INAUDIBLE] the triangle as a different value [INAUDIBLE] start with the coefficient of the first [INAUDIBLE].

BRYNMOR CHAPMAN: So you're talking about coefficients. What are these coefficients? What are the coefficients of? Where are we getting the coefficients?

AUDIENCE: Of x or y.

BRYNMOR CHAPMAN: Not quite of x and y. So we're going to want to group like terms. So both multiplication and addition are commutative, so we can fiddle around with each of those terms and put all the x's at the beginning and the y's at the end.

And then we can also rearrange the terms so that we put like terms together. So we can group them into a bunch of terms of the form x to the k times y to the n minus k and then multiplied by some coefficient c sub k.

So now, as your colleague was saying, those c sub k are going to be the entries of Pascal's triangle. They're going to be the binomial coefficients. Why is that? Well, how many terms are there of this form?

Well, it could come from putting the kx's in any k places from among the n possible places. So you've got n choose k possible places or possible ways to place the x's. The y's will be the remaining ones, and each of those gives you a different term.

So c sub k is going to be n choose k, similar to above. Does that make sense to everybody? So if we specifically put in x and y are both equal to 1, then we recover what we had originally.

You may also have encountered a form of this theorem where y equals 1. Have people seen that before, like 1 plus x to the n? Possibly? Maybe not?

Well, it is a useful theorem. The coefficients on the terms are the n choose k, the binomial coefficients, and this is where the name comes from.

And more generally, you can also get the multinomial theorem in the same way. If you have x plus y plus z for instance, raise that to the n-th power. You start getting multinomial coefficients instead of binomial coefficients.

Yeah. Sorry, I'm moving all the way over here. I didn't plan this out very well.

So for the multinomial theorem, if we have a multinomial here, like we've got mx's instead of just two, if we sum all those up and then raise it to the n-th power, we can use essentially the same proof to show that this is going to be the sum over all sequences k1, k2, all the way up to k sub m, where these all sum up to n of the multinomial coefficient n choose k1, k2, up to km of similar to the x to the ky to the n minus k term we had over there.

This time, we've got a product of m things. So essentially the same proof. If we think of the binomial coefficients as multinomial coefficients where we've only got two things, you can see how it expands.

We're going to look at all possible ways that you could place x1's, x2's, x3's, et cetera among the n different terms, and each of those is going to give us a different term that we can then combine when we're grouping like terms. Does that make sense to people? So the number of those, the number of ways that you can permute those n things, that's just going to be this multinomial coefficient here.

Are people happy with the binomial multinomial coefficients and the binomial multinomial theorems? Yeah.

AUDIENCE: So the coefficients, what is that? Like n over k1 comma k2 comma [INAUDIBLE]?

BRYNMOR CHAPMAN: So--

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: The question is, what does this mean? Is this the notation that you saw in the bookkeeper recitation? Perhaps not. If not, this is equal to n factorial divided by k1 factorial, k2 factorial, all the way up to km factorial.

So it's kind of an extension of the binomial coefficients. With the binomial coefficients, you just have two k's that sum to n. And we also simplify the notation by just dropping one of the-- dropping one of the k's because it's redundant in that case.

But for the multinomial coefficients, we have a bunch of k's that sum to n instead of just two. We're going to divide n factorial by the product of all of them. Does that make sense? Does that answer your question?

Any other questions about binomial multinomial theorems? Yeah. Sorry?

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: Oh, question-- what is the pi? So this is an analog of sigma but for products instead of sums. So we're saying xi to the ki multiplied over all possible i, so x1 to the k1 times x2 to the k2 times x3 to the k3, cetera. So pi for product. Any other questions?

So why don't we move on to another useful identity that we can prove using a combinatorial proof? n choose k equals n minus 1 choose k minus 1 plus n minus 1 choose k.

Who has seen this before? I know that at least one of you has. Few hands.

So we could prove it algebraically, just writing them out with factorials and doing a bunch of fiddly garbage, but that seems a bit difficult. So let's use a double counting proof instead.

How might we use a double counting proof to prove this identity? What set could we try and count? Any ideas? Yeah.

AUDIENCE: [INAUDIBLE] once again, consider the binary sequence [INAUDIBLE] that has k1's in it and [INAUDIBLE]?

BRYNMOR CHAPMAN: Yeah, good answer. So the answer was, once again, we can look at binary sequences of length n, but this time, we're going to restrict it to binary sequences of length n that specifically have k1's. So count binary sequences of length n with k1's.

Now, how many of those are there? Does that give us one of our two expressions easily? If so, which? Yeah?

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: Sorry?

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: Yeah. So essentially by definition, like this is equal to-- the size of this set is going to be n choose k. So let's call this set S.

How could we also make S equal to the right-hand side? Does this give us any ideas? Yeah, answer?

AUDIENCE: [INAUDIBLE] it can be either 0 or 1.

BRYNMOR CHAPMAN: Yeah.

AUDIENCE: [INAUDIBLE] you want to have k1 and then pick an option [INAUDIBLE] original [INAUDIBLE].

BRYNMOR CHAPMAN: Yeah. So the answer was, essentially, we're going to partition our set. So what are we partitioning on? Well, we're going to look at the last bit of the sequence.

So we've got two cases. The last bit can either be 0 or 1. It can't be both. So we've got a partition.

That means that we can use the sum rule. We can count each of these sets individually and then add their sizes together. That will give us the size of our original set. So how many length n sequences with k1's have a 1 at the end? Yeah?

AUDIENCE: [INAUDIBLE]

BRYNMOR CHAPMAN: Yeah. So the answer was n minus 1 choose k minus 1. So let's partition this S, S0 union S1. S0-- or let's call it S sub b-- is the set of binary sequences with n1's-- sorry-- n bits, k1's, and b at the end.

So S sub 1 is going to have size n minus 1 choose k minus 1. If we know that we've got a 1 at the end, we've got k minus 1 additional 1's that we need to place. Where can they go? Into any of the first n minus 1 positions. So there are n minus 1 choose k minus 1 ways to choose those.

And what about S0? Yeah.

AUDIENCE: n minus 1 choose k.

BRYNMOR CHAPMAN: Yeah. This is going to be n minus 1 choose k. And why is that?

AUDIENCE: Because there are still k1's.

BRYNMOR CHAPMAN: Yeah. If we've got a 0 at the end, we still need to place those k1's somewhere. So now we've only got n minus 1 places where we can put them, but we still have k1's that we need to place. So there are n minus 1 choose k ways to place them.

So these two cases are exhaustive. You can't have-- you can't have a sequence that ends with neither a 0 or a 1. It's a binary sequence.

They're also disjoint. You can't have both a 0 and a 1 at the end. So this forms a partition. So by the sum rule, the size of S is going to be the size of S0 plus the size of S1, which, we just established, is-- I guess I flipped them around-- n minus 1 choose k plus n minus 1 choose k minus 1.

And as we established at the beginning, this is also equal to n choose k. Does that make sense to everybody? Do people have any questions about either this identity or combinatorial proofs? Yeah.

AUDIENCE: So the reason why we didn't just write [INAUDIBLE]?

BRYNMOR CHAPMAN: Yeah. The question was, is there any reason we didn't write out the definition of n choose k using the factorial, simplify things? Because I'm a lazy fuck, and I don't want to.

AUDIENCE: [INAUDIBLE] a lot longer than that [INAUDIBLE].

BRYNMOR CHAPMAN: Yeah. So that kind of depends on-- so the concern was that maybe this takes a lot longer than just putting your head down and going through all of the messy algebra.

Yeah. At the beginning, it probably will. If you are not familiar with this proof technique, it may not be obvious what set you should be counting or how you should be counting it, but as you get practice, it could be much easier than just trying to brute-force it. And I think in possibly in recitation, in homework, you will see other identities where it may be less viable to actually go through the messy algebra.

So yeah, this is kind of a silly example, and it's one where you could definitely solve it in other ways. And in fact, those of you who have seen this identity before, I imagine you probably did prove it in other ways. But it's just kind of illustrative of this particular proof technique. Does that answer your question?

AUDIENCE: Yeah.

BRYNMOR CHAPMAN: OK. Any other questions, comments, concerns?

So as your colleague was hinting at earlier, this kind of gives us the following. So I imagine many of you have seen Pascal's triangle before, but if you just write out all of the binomial coefficients in rows of a triangle-- so this is n equals 0, n equals 1, n equals 2, n equals 3. And then successive entries in a row are going to be n choose 0, n choose 1, n choose 2, et cetera.

So if you write all of these out, what we just proved is that every entry of this triangle is the sum of the two entries above it. So if we take this 4 and the 6 here, they sum to 10. That is essentially what this identity is saying. So this is what we call Pascal's triangle.

We also showed that the sum of all of the numbers in a single row is going to be a power of 2. It's going to be 2 to the n. If we sum everything in row 5, we get 1 plus 5 plus 10 plus 10 plus 5 plus 1 equals 32. I think we proved it over there. Yeah, top of the rightmost board.

And you will see, I believe, in homework that if you sum the diagonals, you get the Fibonacci numbers. So if you kind of-- oh, let's see. So go like that. If you sum along these diagonals, you get the Fibonacci numbers.

Am I correct? Is that still a homework problem? Maybe?

Sorry, I've not actually looked at your homework. Full disclosure. It's a fun exercise if you want to try it.

But I will leave you with one more fun exercise. Prove that the sum from k equals 0 to n n choose k squared is going to be 2n choose n. So if you look at all of the entries of a row, instead of just summing them up directly, square them and sum them up. You get the entry that's twice as far down in the middle. So see if you can prove that using a double counting proof.

So that's all we have time for today. If you have any questions, comments, concerns, please feel free to come up. Otherwise, good luck on the exam. Yeah.