YUFEI ZHAO: The goal for the next few lecturers is to prove Freiman's theorem, which we discussed last time. And so we started with this tool that we proved called the Plunnecke-Ruzsa inequality, which tells us that if you have a set A now in an arbitrary abelian group, and if A has controlled doubling, it has bounded doubling, then all the further iterated sumsets have bounded growth, as well. So this is what we proved last time. And so we're going to be using the Plunnecke-Ruzsa inequality many times.

But there are some other tools I need to tell you about. So the next tool is known as Ruzsa covering lemma. All right. So let me first give you the statement, and then I explain some intuition. I think the technique is more important than the statement. But in any case, here's what it says.

If you have X and B-- they're subsets of some arbitrary abelian group-- if you have the inequality that X plus B is at most K times the size of B-- so the size of X plus B is at most K times the size of B-- then there exists some subset T of X, with the size of T being, at most, K, such that X is contained in T plus B minus B. So that's the statement of Ruzsa covering lemma. So let me play to explain what it is about.

So the idea here is that, if you're in a situation where if it looks like-- so you should think of B as a ball. So if it looks like X plus B might be coverable by K translates of B-- so here, you're supposed to think of B like a ball in the metric space-- then X is actually coverable. So if it looks like, meaning just by size alone, by size info alone-- so just based on the size, if it looks like X plus B might be coverable by K different translates of B, then actually X is coverable, but you have to use slightly larger balls by K copies of B minus B. So you should think of B minus B as slightly larger balls than B itself. So if B were an actual ball in Euclidean space, then B minus B is the same ball with twice the radius.

And so Ruzsa covering lemma is a really important tool. The proof is not very long. And it's important to understand the idea of this proof. So this is a proof not just in additive combinatorics but something that happens-- it's a standard idea in analysis. So it's very important to understand this idea.

And the key idea-- here, I think the proof is more important than the statement up there-- the key idea here is that if you want to produce a covering, one way to produce a covering is to

take a maximal packing. A maximal packing using with balls, for instance, implies a covering with balls twice as large. So let me illustrate that with a picture.

Suppose you have some space that I want to cover. But you get to use, let's say, unit balls. So how can I make sure I can cover the space using unit balls? And I don't want to use too many unit balls.

So what you can do is, let me start by a maximal set of unit balls, so with centers-- so it's now half unit balls, so the radius is 1/2. So I put in as many as I can so that I cannot put in any more. So that's what maximal means-- "mal" here doesn't mean the maximum number. Although, if you put in the maximum number, that's also OK. But maximal means that I just cannot put in any more balls that are not overlapping.

If I have this configuration, now what I do is I double the radii of all the balls. So whoever takes today's notes will have a fun time drawing this picture. So this has to be a covering of the original space. Because if you had missed some point, I could have put a blue ball in. Yes?

AUDIENCE: What if a space has some narrow portion?

YUFEI ZHAO: Question-- what if a space has some narrow portion? It doesn't matter. If you formulate this correctly-- if you miss some point-- so here, it depends on how you formulate this covering.

The point is that, if you take a maximal set of points, when you expand, you have to cover the whole space. Because if you missed some point-- so imagine if you had-- for example, suppose you missed-- suppose you had missed some point over here. Then I could have put an extra ball in. So if you had missed some point over here, that means that I should have been able to put in that ball there initially.

So this is a very simple idea, but a very powerful idea. And it comes up all the time in analysis and geometry. And it also comes up here. So let's do the actual proof.

So let me let T be a subset of X be a maximal subset of X, such that the sets little t plus B are disjoint for all elements little t of a set big T. So it's like this picture. I pick a subset of X so that if I center balls around these t's, then these translates of B's are destroyed. Then you put in a maximal such set.

So due to the disjointness, we find that the product of the sizes of t and B equals to the size of the sum, because there's no overlaps. But t plus B has size at most X plus B, because T is a

subset of X. But also, from the assumption we knew that X is-- size of x plus B is upper bounded by the size B times K.

So in particular, we get that the size of T is at most K. So in other words, over here, the number of blue balls you can control by simply the volume. Now, since T is maximal, we have that for every little x, there exists some little t such that the translate of B given by x intersects one of my children translates. Because if this were not true, then I could have put in an extra translate of B.

So in other words, there exist two elements, b and b prime, such that t plus b equals to x plus b prime. And hence, x lies in little t plus B minus B, which implies that the set X lies in T plus B minus B. OK. So that's the Ruzsa covering lemma. So it's an execution of this idea I mentioned earlier that a maximal packing implies a good covering.

Any questions? Right. So now we have this tool, we can prove an easier version of Freiman's theorem where, instead of working in integers, we're going to work in the finite field model. So usually, it's good to start with finite field models. Things are a bit cleaner. And in this case, it actually only requires a subset of the tools that we need for the full theorem.

So instead of working in finite field model, we're going to be working in something just slightly more general. But it's the same proof. So we're going to be working in a group of bounded exponent. Freiman's theorem in groups of bounded exponent. And the word "exponent" in group theory means the following-- so the exponent of an abelian group is the smallest positive integer, if it exists. So the smallest positive integer r so that rx equals to 0 for all x in the group.

So for example, if you are working Fp to the n, then the exponent is p. So every element has-- if you add to itself r times, then it vanishes the element. The word "exponent" comes from-- I mean, you can define the same thing for non-abelian groups, in which case you should write this expression using the exponent. Instead of addition, you have multiplication. So that's why it's called exponent. So the name is stuck, even though we're still working on the additive setting.

We're going to write this, so the angle brackets, to mean the subgroup generated by the subset A, where A is a subset of the group elements. And so the exponent of a group, then, is equal to the maximum number of-- so if you pick a group element, look at how many group elements does it generate. So it's equal to the max.

All right. I mean, the example you can have in mind is F2 to the n, which has exponent 2. So in general, we're going to be looking at a special case in a group with bounded exponent. And Freiman's theorem, in this case, is due to Ruzsa, who showed that if you have a finite subset in an abelian group with exponent r-- so a finite exponent-- if A has doubling constant at most K, then what do we want to say? We want to say, just like in Freiman's theorem, that if A has bounded doubling, then it is a large portion of some structured set. And here "structured set" means subgroup.

Well, if you're going to be looking at subgroups that contain A, you might as well look at the subgroup generated by A. So the claim, then, is that the subgroup generated by A is only a constant factor bigger than A itself. So let me say it again. If you have a set A in a group of bounded exponent, and A has bounded doubling, then conclusion is that A is a large proportion of some subgroup.

Conversely-- we saw last time-- if you take a subgroup, it has doubling constant 1, so if you take a positive proportional constant proportion of the subgroup, it also has bounded doubling. And this statement is in some sense the converse of that observation. This bound here is not optimal. I'll make some comments about what we know about these bounds. But for now, just view this number as a constant.

Any questions about the statement? All right. So let's prove this Ruzsa's theorem, giving you Freiman's theorem in groups with bounded exponent. So we're going to be applying the tools we've seen so far, starting with Plunnecke-Ruzsa. So by Plunnecke-Ruzsa inequality, we find that-- so then I'm going to write down some expression. You may wonder, why this expression? Because we're going to apply covering lemma in a second.

So this set by Plunnecke-Ruzsa, its size is bounded, because the set can also be written as 3A minus A, so its size is bounded by K to the fourth times the size of A. And so Plunnecke-Ruzsa is a very nice tool to have. It basically tells you, if you have bounded doubling, then all the other iterated sums are essentially controlled in size. And then we're using that here.

OK. So now we're in the setting where we can apply the Ruzsa covering lemma. So by covering lemma, we're going to apply the covering lemma-- so using the notation earlier-- with X equal to 2A minus A and B equal to A. By covering lemma, there exists some T being a subset of 2A minus A, with T not too large-- because of our earlier estimate-- and such that 2A minus A is contained in T plus A minus A.

So it's easy to get lost in the details. But what's happening here is that I start with A, and I'm looking at how big these iterated growing sumsets can be. And if I keep on doing this operation, like if I just apply Plunnecke-Ruzsa, I can't really control the iterated sums by an absolute bound. This bound will keep growing if I take bigger iterations.

But Ruzsa covering lemma gives me another way to control the bounded iterations. It says there is some bounded set T such that this iterated sum is nicely controlled. So let me iterate this down even further. We're going to iterate the containment by adding A to both sides. And we obtain that 3A minus A is contained in T plus 2A minus A. But now, 2A minus A was containing T plus A minus A. So we get 2T plus A minus A.

So now we've gained quite a bit, because we can make this iteration go up, but not at the cost of iterating A but at the cost of iterating T. But T is a bounded size. T is bounded size, so T will have very nice control. We'll be able to very nicely control the iterations of T.

So if we keep going, we find that for all integer n, positive integer n, n plus 1 A minus A is contained in nT plus A minus A. But for every integer n, the iterated sums of T are contained in the subgroup generated by T. So therefore, if we take n to be as large as you want, see that the left-hand side eventually becomes the subgroup generated by A, and the right hand side does not depend on n. So you have this containment over here.

We would like to estimate how large the subgroup generated by A is. So we look at that formula, and we see that the size of the subgroup generated by T-- and here is where we're going to use the fact that the assumption that a group has bounded exponent. So think F2 to the n, if you will.

So in F2 to the n, if I give you a set T, what's the subspace spanned by T? What's the maximum possible size? It's, at most, 2 raised to the number. So in general, you also have that the subgroup generated by T has size at most r raised to the size of T in a group of exponent r, which we can control, because T has bounded size. And the second term, A minus A, so we also know how to control that by Plunnecke-Ruzsa.

Therefore, putting these two together, we see that the size up here is at most r to the K to the 4 times K squared size of A, which is the bound that we claimed. Any questions? So the trick here is to only use Plunnecke-Ruzsa somehow a bounded number of times. If you use it too many times, this bound blows up. But you use it only a small number of times, and then you

use the Ruzsa covering lemma so that you get this bounded set T that I can iterate.

Let me make some comments about the bounds that come out of this proof. So you get this very clean bound, following this proof. But what about examples? So what kinds of examples can you think of where you have a set of bounded doubling, but the set generated, the group generated by that set, is potentially large?

So even in F2 to the n, so if A is an independent set-- so a basis or a subset of a basis, for instance-- then K-- so A is independent set, so all the pairwise sums are distinct-- so K is about the size of A/2. That's the doubling constant. Whereas, the group generated by A has size 2 to the size of A, which is around 2 to the 2K times the size of A. OK.

So you see that you do need some exponential blow-up from K to this constant over here. And turns out, that's more or less correct. So the optimal constant for F2 to the n is now known very precisely. And so if you give me a real value of K, then I can tell you there are some recent results that tells you exactly what is the optimal constant you can put in front of the A. So very precise. But asymptotically, it looks like 2 to the 2K divided by. K So that's what it looks like. So this example is basically correct.

For general r, we expect a similar phenomenon. So Ruzsa conjectured that-- in the hypothesis of the theorem, the constant you can take is only exponential in K, the r to the some constant C times K. And it has been verified for some values of r, but not in general-- for some r, for example primes.

OK. Any questions? All right. So this is a good milestone. So we've developed some tools, and we were able to prove a easier version of Freiman's theorem in a group of bounded exponent. And you can ask yourself, does this proof work in the integers? And well, literally no. Because if you look at this proof, this set here is infinite, unlike in the finite field setting. In the integers, well, that's not very good.

So the strategy of Freiman's theorem, the proof of Freiman's theorem, is to start with the integers, and then try to, not work in the integers, but try to work in a smaller group. Even though you start in a maybe very spread out set of integers, I want to work in a much smaller group so that I can control things within that group. And this is an idea called modeling. So I have a big set and want to model it by something in a small group.

So we're going to see this idea. But to understand what does it mean to have a good model for

a set in the sense of additive combinatorics, I need to introduce the notion of Freiman homomorphisms. So one of the central philosophies across mathematics is that if you want to study objects, you should try to understand maps between objects and understand properties that are preserved under those maps.

So if you want to study groups, I don't really care how you label your group elements-- by 1, 2, 3, or A, B, C. What I care about is the relationships. And those are the data that I care about. And then, of course, then you have concepts like group homomorphisms, group isomorphisms, that preserve all the relevant data. Similarly in any other area-- in geometry you have manifolds. You understand not specifically how they embed into space but what are the intrinsic properties.

So we would like to understand what are the intrinsic properties of a subset of an abelian group that we care about for the purpose of additive combinatorics, and specifically for Friedman's theorem. And what we care about is what kinds of additive relationships are preserved. And Freiman's homomorphisms capture that notion.

So roughly speaking, we would like to understand maps between sets in possibly different groups-- in possibly different abelian groups-- that only partially preserve additive structure. So here's a definition. Suppose we have A and B, and they're subsets in possibly different abelian groups. Could be the same, but possibly different. And everything's written under addition, as usual.

So we say that a map phi from A to B is a Freiman s-homomorphism. So that's the term-- Freiman s-homomorphism, sometimes also Freiman homomorphisms of order s, so equivalently I can call it that, as well-- if the following holds. If we have the equation phi of A plus dot, dot, dot plus phi of a sub s equal to phi of a prime 1 plus dot, dot, dot plus phi of a prime s, whenever a through as, a prime through as prime, so a1 prime through as prime, satisfy the equation a1 plus dot, dot, dot plus as equal to a1 prime plus dot, dot, dot plus as prime.

OK. So that's the definition of a Freiman s-homomorphism. It should remind you of the definition of a group homomorphism, which completely preserves additive structure, let's say, between abelian groups. And for Freiman homomorphisms, I'm only asking you to partially preserve additive structure. So the point here is that if I only care about, let's say, pairwise sums, if that's the only data I care about, then Freiman homomorphisms preserve that data.

To give you some-- OK, so one more thing. If phi from A to B is, furthermore, a bijection, and both phi and phi inverse are Freiman s-homomorphisms, then we say that phi is a Freiman s-isomorphism. So it's not enough just to be a bijection, but it's a bijection and both the forward and the inverse maps are Freiman homomorphisms. So these are the definitions we're going to use.

Let me give you some examples. So every group homomorphism is a Freiman homomorphism of every order. So group homomorphisms preserve all additive structure, and Freiman homomorphisms only partially preserve additive structure.

A composition-- so if phi 1 and phi 2 are Freiman s-homomorphisms, then phi 1 composed with phi 2 is a Freiman s-homomorphism. So compositions preserve this property. And likewise, instead of homomorphisms, if you have isomorphisms, then that's also true, as well. So these are straightforward things to check.

So a concrete example that shows you a difference between group homomorphisms and Freiman homomorphisms is, suppose you take an arbitrary map phi from a set that has no additive structure. So it's a four-element set, has no additive structure. And I map it to the integers, claim that this is a Freiman 2-homomorphism. So you can check. So whenever this is satisfied, but that's never non-trivially satisfied. So an arbitrary map here is a Freiman 2-homomorphism.

And if furthermore-- so if you have, let's say, bijection between two sets, both having no additive structure, if it's a bijection, it's a Freiman isomorphism of here, order 2. Let me give you a few more examples. When you look at homomorphisms between finite groups, so you know that if you have a homomorphism and it's also a bijection, then it's an isomorphism.

But that's not true for this notion of homomorphisms. So the natural embedding that sends the Boolean cube to the Boolean cube viewed as Z mod 2 to the n. So what's happening here? This is a part of a group homomorphism. And so if you look at Z to the n, and I do mod 2, and I restrict to this Boolean cube, that's the group homomorphism If I view this as a subset of a bigger group.

So it is a Freiman homomorphisms of every order. And it's bijective. But it is not a Freiman 2-isomorphism. Because you have additive relations here that are not present over here. So if you read the definition, the inverse map, there are some additive relations here that are not preserved if you pull back.

Here's another example that will be more relevant to our subsequent discussion. So the mod N map, which sends Z to Z mod N, so this is a group homomorphisms. So hence, it's a Freiman homomorphism of every order. But it's not-- OK, so if you look at this map, and even if I restrict to here, so it's not a Freiman isomorphism just like earlier.

However-- let me go back to Z. So if A is a subset of integers with diameter less than N/s, then this map mod N maps A Freiman s-isomorphically onto its image. So even though mod N restricted to 1 through N is not a Freiman isomorphism of order 2, if I restrict to a subset that's, let's say, contained in some small interval, then all the additive structures are preserved.

So let me show you why. So this is not too hard once you get your head around the definition. So indeed, if you have group elements a1 through as, a prime 1 through a prime s, and if they satisfy the equation-- so if they satisfy this equations, so we're trying to verify that it is a Freiman s-isomorphism, namely the inverse of this map is a Freiman s-homomorphism, so if they satisfy this equation, so this is satisfying this additive relation in the image, in Z mod N, then note that the left-hand side-- so all of these A's are contained in a small interval, because the diameter of the set is less than N/s.

So if you look at how big a1 minus a1 prime can be it's, at most-- or, it's less than N/s in size. So the left-hand side, in absolute value, viewed as an integer-- so the left-hand side is less than N in absolute value, since the diameter of a is less than N/s. So you have some number here, which is strictly less than N in absolute value, and it's 0 mod N, so it must be actually equal to 0 as a number as an integer. So this verifies that the additive relations up to s-wise sums are preserved under the mod N map, if you restrict to a small interval.

Any questions so far? So in additive combinatorics, we are trying to understand properties, specific additive properties. And the notion of Freiman homomorphisms Freiman isomorphism capture what specific properties we need to study and what are the maps that preserve those properties.

And the next thing we will do is to understand this model lemma, the modeling lemma, that tells us that if you start with a set A with small doubling, initially A may be very much spread out in the integers-- it may have very large elements, very small elements, very spread out. But if A small doubling properties, then I can model A inside a small group, such that all the relevant data, namely relative to these Freiman homomorphisms, are preserved under this

model.

So let's move on to the modeling lemma. The main message of the modeling lemma is that if A has small doubling, then A can be modeled-- and here, that means being Freiman isomorphic-- to a subset of a small group. So first as a warm-up, let's work in the finite field model, just to see what such a result looks like. And it contains most of the ideas, but it's much more clean. It's much cleaner than in the integers.

So in the finite field model, specifically F2 to the n, what do we want to say? Suppose you have A, a subset of F2 to the n, and suppose that m is some number such that 2 to the m is at least as large as sA minus sA. So remember, from Plunnecke-Ruzsa, you know that if A plus A is small, then this iterated sum is small.

So suppose we have these parameters and sets. The conclusion is that A is Freiman s-isomorphic to some subset of F2 raised to m. So initially, A is in a potentially very large vector space, or it could be all over the place. And what we are trying to say here is that if A has small doubling, then by Plunnecke-Ruzsa, sA minus sA has size not too much bigger than A itself-- only bounded times the size of A itself. So I can take an m so that the size of this group is only a constant factor in larger than the size of A itself.

So we're in a pretty small group. So we are able to model A, even though initially it sits inside a pretty large abelian group, by some subset in a pretty small group. So let's see how to prove this modeling lemma.

So the finite field setting is much easier to-- it's not too hard to deal with, because we can look at linear maps. So the following are equivalent for linear maps phi, so for group homomorphisms. So phi is a Freiman s-isomorphism when restricted to A. So here, phi I'm going to let it be a linear map from F2 to the n to F2 to the m. The following are equivalent.

So we would like phi to be a Freiman s-isomorphism when restricted to A. Because this means that when I restrict to A and I restrict to its image, it Freiman isomorphically maps onto the image. So what does that mean?

So phi is already a homomorphism. So it's automatically a Freiman s-homomorphism. For it to be an s-isomorphism, it just means that there are no additional linear relations in the image that were not present earlier, which means that phi is injective on sA. So let's just think about the definition. And in the definition, if you know additionally that phi is a homomorphism,

everything's much cleaner.

It is also equivalent to that phi of x is non-zero for every non-zero element x of sA minus sA. So this is a very clean characterization of what it means to be in Freiman s-isomorphism when you are in an abelian group and you have linear maps of homomorphisms.

So if we start by taking phi to be a uniformly random linear map-- so for example, you pick a basis, and you send each basis element to a uniformly random element-- then we find that if 2m is at least A at minus sA, then-- so let me call these properties 1, 2, 3-- so then the probability that 3 is satisfied is positive. Because each element of sA minus sA-- I can also ignore 0-- so each non-zero element of sA minus sA violates this property with probability exactly 2 to the minus m, everything is uniform.

So if there are very few elements, and the space is large enough, then the third bullet is satisfied with positive probability. So you get Freiman s-isomorphism. Any questions?

To get this model, in this case in the finite field setting, it's not so hard. So you kind of project the whole set, even though initially it might have or be very spread out into a lot of dimensions. You project it down to a small dimensional subspace randomly, and that works. So then with high probability, it preserves all the additive structures that you want, provided that you have small doubling.

Now, let's look at what happens in Z. So in Z, things are a bit more involved. But the ideas-- actually, a lot of the ideas-- come from this proof, as well.

So Ruzsa's modeling lemma tells us that if you have a set of integers-- always a finite set-- and integers s and N are such that N is at least sA minus sA, then so it turns out you might not be able to model the whole set A. But it will be good enough for us to model a large fraction. So then there exists an A prime subset of A, with A prime being at least an s fraction of the original set. And A prime is Freiman s-isomorphic to a subset of Z mod N.

So same message as before, with an extra ingredient that we did not see before. But the point is that if you have a set A with controlled doubling, then well, now you can take a large fraction of A that is Freiman isomorphism to a subset of a small group. This is small, because we only need n to exceed sA minus sA, which are only constant factor more than the size of A. Yeah?

AUDIENCE: Is s greater than 2 comma N?

YUFEI ZHAO: Sorry. S greater than 2 separate and some integer. Thank you. So in our application, s will be a constant. s will be 8. So think of s as some specific constant. Any questions?

So let's prove this modeling lemma. We want to try to do some kind of random map. But it's not clear how to start doing a random map if you just start in the integers. So what we want to do is first place ourself in some group where we can consider random automorphisms.

So we start by-- perhaps we're very wastefully choosing a prime q bigger than the maximum possible sA minus sA. And so just choose a large enough prime. I don't care how large you pick. q can be very, very large. Pick a prime.

And now I work inside Z mod q. I noticed that if you make q large enough, then A sits inside Z mod q Freiman isomorphically, or s-isomorphically. So just pick q large enough so that you don't have to worry about any issues. So Yeah. So the mod q map from A to Z mod q-- so this is Freiman s-isomorphic-- onto its image.

So let's now consider a sequence of maps. And we're going to denote the sequence like this. So we start with Z. That's where A originally sits. And maps to Z mod q-- so that was the first map that we saw. And now we want to do a random automorphism, kind of like the random map earlier. But in Z mod q, there are lots of nice random automorphisms, namely multiplication by some non-zero element.

And finally, we can consider the representative map, where every element of Z mod q, I can associate to it a positive integer from 1 to q which agrees with the Z mod q. So the final step is not a group homomorphism. So we need to be more careful.

So let me denote by phi this entire map. So from the beginning to the end, this composition I'll denote by phi. And lambda here is some element between 1 and q minus 1. Now, remember what we said earlier, that this map, this final map here, might not be a Freiman homomorphism, because there are some additive relations here that are not preserved over here.

But if I restrict myself to a small interval, then it is a Freiman homomorphism if we restrict to that interval. If you restrict yourself to an interval, you cannot have extra relations over here, because they cannot-- the interval is small enough, you can't wrap around. So let's consider restrictions to small intervals.

I start with A over here. So I want to restrict myself to some interval so that I still have lots of A

in that restriction. And you can do this by pigeonhole. So by pigeonhole, for every lambda there exists some interval we'll denote by I sub lambda inside q.

So the length of this interval will be at most q/s, such that if I look at the restriction of this interval, I pull it all the way back to the beginning then I still get a lot of elements of A. So A sub lambda, namely the elements of A whose map gets sent to this interval, has at least A/s elements. So for instance, you can chop up q into s different intervals. So one of them will have lots of elements that came from A.

And this is why in the end we only get a large subset of the A. So we're going to forget about everything else and focus our attention on the set here. So thus, by our earlier discussion having to do with the final map being a Freiman s-homomorphism when you're working inside a short interval, we see that phi, if you restrict to this A sub lambda, is a Freiman s-homomorphism. Each step is a Freiman homomorphism, because it's a group homomorphism. And the final step in the restriction is also a Freiman s-homomorphism, because what we said about working inside short intervals.

So this part is very good. All right. So now let me consider one more composition. So at the end of the day, we would like to model this A lambda inside some small cyclic group. So far, we don't have small cyclic groups here. But I'm going to manufacture a small cyclic group.

So we're going to consider the map where, first we take our phi all the way until the end, and now you take mod m map. So if I don't write anything, if it goes to Z mod m, it means the mod m map. So let me consider psi, which is the composition of these two maps.

All right. So we would like to say that you can choose this lambda so that this A sub lambda gets mapped Freiman s-isomorphically until to the end. So far, everything looks pretty good. So you have Freiman s-homomorphism, and you have a group homomorphism. So the whole thing is a Freiman s-homomorphism. So psi is restricted to A sub lambda, is a Freiman s-homomorphism.

But now the thing that we really want to check is if there are some relationships that are present at the end in Z mod m that were not present earlier. And so we need to check that-- we claim that if psi does not map A sub lambda Freiman isomorphically, then something has to have gone wrong. So if it does not map A sub lambda Freiman isomorphically onto its image, then what could have gone wrong?

Claim that there must be some d which depends on lambda in sA minus sA, and d not 0, such that phi of d is 0 mod m. So we'll prove this. But like before, it's a very similar idea to what's happening earlier. The idea is that if you have-- we want to show that there are no additional additive relations in the image.

So we would like-- so if it's a Freiman isomorphism, then there has to be some accidental collisions. And that accidental collision has to be witnessed by some d. So this requires some checking.

So suppose-- indeed, suppose the hypothesis-- suppose that the psi does not map A sub lambda Freiman s-isomorphically onto its image. Then there exists a1 through as, a1 prime through as prime, in A lambda such that they do not have additive relation, but their images do have this additive relation. Their images all the way until the end having the additive relation means that phi has this additive relation mod m.

OK. So how can this be? Recall that since the image-- so all of these elements-- lie inside some short interval. The interval has length less than 2 minus s. So we saw this argument, very similar argument earlier, before the break. Because everything lies in the short interval, we see that this difference between the left- and the right-hand sides, this difference is strictly less than q.

Now, by switching the a's and a primes if necessary, we may assume that this difference is non-negative. Otherwise, it's just a labeling issue. Otherwise, I relabel them. So then this here-- so what's inside this expression-- we call this expression inside the absolute value, we call it star. So star is some number between 0 and strictly-- so at least 0 and strictly less than q.

Right. So if we set d to be this expression, so the difference between these two sums, on one hand this d here-- sorry, that's what I want to say. Suppose you don't have-- if you are not mapping Freiman isomorphically onto the image, then I can exhibit some witness for that non-isomorphism, meaning a bunch of elements that do not have additive relations in the domain. But I do have additive relation in image.

So if we set this d, then it's some element of sA minus sA. And it's non-zero, because we assume that d is non-zero. And so then, what can we say about phi of d?

So phi of d, I claim, must be this expression over here, the difference of the corresponding sums in the image. Because the two sides are congruent mod q-- two sides are congruent

mod q. And furthermore, they are in the interval from 0 to strictly less than q. So this is a slightly subtle argument. But the idea is all very simple. Just have to keep track of the relationships between what's happening in the domain, what's happening in the image.

Somehow, I think the finite field case is quite illustrative of-- there, what goes wrong is similar to what goes wrong here. Except here, you have to keep track a bit more things. OK. So consequently, thus phi of d is congruent to 0 mod m, which is what we're looking for. So that proves the claim.

Any questions? Yeah?

AUDIENCE: [INAUDIBLE]

YUFEI ZHAO: OK. So this part? All right. So we set d to be this expression. I claim this equality. So why is it true? First, the left-hand side and the right-hand side, they are congruent to each other mod q. So why is that?

AUDIENCE: [INAUDIBLE]

YUFEI ZHAO: Sorry? Yeah.

AUDIENCE: [INAUDIBLE] every part of phi should preserve-- the first two parts are group homomorphisms. And then we just take [INAUDIBLE] mod q.

YUFEI ZHAO: Exactly. If you look at phi-- where was it?-- up there, you see that everything preserves. Even though the very last step is not a group homomorphism, mod q is preserved. So even though the last step is not group homomorphism, it is all mod q. So if you're looking at mod q, everything is group homomorphism. So here we're good.

Both are in this interval. The right-hand side is in the interval because of our assumption about short-- everything living inside a short interval. And the left-hand side is by definition. Because the image of phi is in that interval, especially if given that d is not equal to 0.

Is that OK? So it's not hard, but it's a bit confusing. So think about it. All right. So we're almost done. So we're almost done proving the Ruzsa modeling lemma in Z mod m. So let me finish off the proof.

So for each non-zero d in this iterated sumset, basically, we would like to pick a lambda so that that map up there does what we want to do. If it doesn't do what we want to do, then it is

witnessed by some d, like this. Those are the bad lambda. So if there exists a d, then this lambda is bad.

So for each d that potentially witnessed some bad lambda, the number of bad lambda, i.e., such that phi of d is congruent to 0 mod m, so here we're no longer even thinking about group homomorphisms anymore or the Freiman homomorphisms. It's just a question of, if I give you a non-zero integer, how many lambdas are there so that phi of d is divisible by m?

Remember that we picked q large enough so that, initially, you are sitting very much inside-- everything's really between 0 and q. So this dot lambda up there lacks uniformity. So the number of such bad lambdas is exactly the number of elements in this interval that are divisible by m.

So everything's more or less a bijection if you restrict to the right places. And the number of such elements is, at most, q minus 1 over m. So therefore, the total number of bad lambdas is, at most, for each element d of sA minus sA, a non-zero element.

We have, at most, q minus 1 over m bad lambdas. So the total number of bad lambdas is strictly less than q minus 1. So there exists some lambda such that psi, when restricted to A sub lambda, maps Freiman s-isomorphically onto the image.

Somehow, I think it's really the same kind of proof as the one in the finite field case, except you have this extra wrinkle about restricting too short diameter intervals, to short intervals. But the idea is very similar. OK. So that's the Freiman model lemma in the integers.

And let me summarize what we know so far. And so that will give you a sense of where we're going in the proof of Freiman's theorem. So what we know so far is that if you have a subset of integers, a finite subset, such that A plus A is size A times K at most, doubling constant at most K, then there exists some prime N at most 2K to the 16th times the size of A and some subset A prime of A such that A prime is Freiman 8-isomorphic to a subset of Z mod NZ.

So it follows from two things we've seen so far. Because by the Plunnecke-Ruzsa inequality 8A minus 8A is, at most, K to the 16 times A. And now we can choose a prime N between K to the 16 and 2 times K to the 16 and apply the modeling lemma.

So that's where we are at. So you start with a set of integers with small doubling. Then we can conclude that by keeping a large fraction of A, keeping-- I forgot to-- so very important. So there exist some A, which is a large fraction. Keeping a large fraction of A, I can model this

large subset of A by some subset of a cyclic group, where the size of the cyclic group is only a constant times more than the size of A.

So now, we are going to work inside a cyclic group and working with a set inside a cyclic group that's a constant size, a constant fraction of the cyclic group. Question is why 8? So that will come up later. So basically, you need to choose some numbers so that you want to preserve the structure of GAPs. So that will come up later.

And now we're inside a cyclic group. And you have a constant fraction of a cyclic group. Where have we seen this before? So when we proved Roth's theorem, that was the setting. So in cyclic loop, you have a constant fraction of cyclic group, and you can do Fourier analysis.

Initially, you could not do Fourier analysis starting with Freiman's theorem, because the set may be very, very spread out. But now, you are a large fraction of a cyclic group. So we're going to do Fourier analysis next time to show that such a set must contain lots of structure, just from the fact that it is large.

So that will be the next step. Good. Happy Thanksgiving.