HERBERT GROSS: Hi. Well, I guess I really should have said "goodbye," because this is the last lecture in our course-- not the last assignment, but the last lecture. The reason I said "hi" was, why quit after all this time with saying that? And we've reached the stage now where we should clean up the vector spaces to the best of our ability and recognize that, from this point on, much of the treatment of vector spaces requires specialized concentration.

In fact, I envy the real people who make regular movies where they have stuntmen when things get tough. I would continue on with this course, except that I don't have a stuntman to do these hard lectures for me.

And also, the particular topic, as I told you last time, that I have in mind for today-- the subject called eigenvectors-- has two approaches to it. One is that it does have some very elaborate practical applications, many of which occur in more advanced subjects.

It also has a very nice framework within the game of mathematics idea. And my own feeling was that, since we started this course with the concept of the game of mathematics, mathematical structures, I thought that, rather than go into complicated applications, I would treat eigenvectors in terms of the structure of mathematics as a game.

In fact, as an interesting aside, I'd like to share with you a very famous story in mathematics, that, when the first book on matrix algebra was written by Hamilton, he inscribed the book with, "Here at last is a branch of mathematics for which there will never be found practical application." He did not invent the subject to solve difficult physics problems. He invented the subject because it was an elegant mathematical device. And so I thought that maybe, for our last lecture, we should end on that vein.

At any rate, the subject for today is called eigenvectors. And from a purely game point of view, the idea is this. Let's suppose that we have a vector space V and a linear mapping, a linear transformation, f, that maps V onto itself, say. And the question is, are there any vectors in V other than the zero vector, such that f of v is some scalar multiple of v?

You see, the reason I exclude the zero vector, first of all, is that we already know that, for a linear transformation, f of 0 is 0. And c times 0 is 0 for all c. So this would be trivially true if v were the zero vector. So what we're really interested in-- given a particular linear transformation, are there any vectors such that, relative to that linear transformation, the mapping of that vector is just a scalar multiple of that vector itself? In other words, does f preserve the direction of any vectors in V?

By the way, don't confuse this with conformal mapping that we talked about in complex variables. In conformal mapping, we didn't preserve any-- in general, we did not preserve directions of lines. We preserved angles. In other words, an angle might have been rotated so that the direction of the two sides may have changed. It was the angle that was preserved in the conformal mapping. What we're asking now is, given a linear transformation, does it preserve any directions?

And to illustrate this in terms of an example, let's suppose we think of mapping the xy-plane into the uv-plane under the linear mapping f bar, where f bar maps x, y into the 2-tuple x plus 4y comma x plus y. In other words, in terms of mapping the xy-plane into the uv-plane, this is the mapping-- u equals x plus 4y, v equals x plus y. It's understood here, when I'm referring to the typical xy-plane, that my basis vectors are i and j.

Notice, by the way, that if I look at the vector i, i is the 2-tuple 1 comma 0. Notice that when x is 1 and y is 0, u is 1 and v is one. So f bar maps 1 comma 0 into 1 comma 1. Again, in terms of i and j vectors, f bar maps i into i plus j. And i plus j is certainly not a scalar multiple of i.

Similarly, what does f bar do to j? j, relative to the basis i and j, is written as 0 comma 1. When x is 0 and y is 1, we obtain that u is 4 and v is 1. So under f bar, 0 comma 1 is mapped into 4 comma 1. Or in the language of i and j components, f bar maps j into 4i plus j. And that certainly is not a scalar multiple of j. In other words, 4i plus j is not parallel to j.

On the other hand, let me pull this one out of the hat. Let's take the vector 2i plus j-- in other words, the 2-tuple 2 comma 1. When x is 2 and y is 1, we see that x plus 4y is 6 and x plus y is 3. So f bar maps 2 comma 1 into 6 comma 3. And that certainly is 3 times 2 comma 1-- see, by our rule of scalar multiplication. In other words, what this says is that the vector 2i plus j is mapped into the vector which has the same sense and direction as 2i plus j, but is 3 times as long.

So you see, sometimes the linear transformation will map a vector into a scalar multiple of itself. Sometimes it won't. Sometimes there'll be no vectors, other than the zero vector, that they're mapped into-- scalar multiples of themselves, and things of this type. But that's not important right now. In terms of a game, what we're trying to do is what? Solve the equation [INAUDIBLE]-- well, let's give it in terms of a definition.

First of all, if we have a vector space V, and f is a linear transformation mapping V into itself, if little v is any non-zero element of the vector space V, and if f of v equals c times v for some scalar-- for some number c-- then v is called an eigenvector and c is called an eigenvalue. In other words, if a vector has its direction preserved, geometrically speaking, all we're saying is that if the direction doesn't change, the vector is called an eigenvector.

And the scaling factor-- which means what? Even though the direction doesn't change, the image may have a different magnitude, because the scalar c here doesn't have to be 1. That scalar is called an eigenvalue. And I'll give you more on this in the exercises, and perhaps even in supplementary notes if the exercises seem to get too sticky. But we'll see how things work out.

For the time being, all I care about is that you understand what an eigenvector means and what an eigenvalue is. Quickly summarized, if f maps a vector space into itself, an eigenvector is any non-zero vector which has its direction preserved under the mapping f-- that f maps it into a scalar multiple of itself.

There is a matrix approach for finding eigenvectors, and the matrix approach also gives us a very nice review of many of the techniques that we've used previously in our course. For the sake of argument, let's suppose that V is an n-dimensional vector space, and that we've again chosen a particular basis, u1 up to un, to represent V. Suppose, also, that f is a linear mapping carrying V into V.

Remember how we used the matrix approach here? What we said was, look at. The vectors u1 up to un are carried into f of u1 up to f of un. And that determines the linear transformation f because of the fact of the linearity properties. In other words, when you take f of a1 u1 plus a2 u2, it's just a1 f of u1 plus a2 f of u2.

So once you know what happens to the basis vectors, you know what happens to everything. But since we're expressing V in terms of the basis u1 up to un, that means that f of u1, et cetera, f of un may all be expressed as linear combinations of u1 up to un. And that's precisely what I've written over here.

We also know that the vector v that we're trying to find-- see, remember we're trying to find eigenvectors. The vector v, relative to the basis u1 up to un, can be written as the n-tuple x1 up to xn. And now I won't bother writing this, because I hope, by this time, you understand this. This is an abbreviation for saying what? The vector v is x1 u1 plus et cetera xn un, because I am always referring to the specific basis when I write n-tuples without any other qualifications.

Now, the question was, how did the statement f of v equal cv translate in the language of matrices? Remember that we took this particular matrix of coefficients and wrote-- well, we transposed it. Remember what we said? We said that the matrix A would be the matrix whose first column would be the components of f of u1 and whose nth column would be the components of f of un.

In other words, what we did was, is we said that to take f of v, we would just take the matrix-- notice how I've written it, now-- not a1,1, a1,2, a1,1, you see? a2,1, a3,1, et cetera-- you see, these make up the components of f of u1. These make up the components of f of un. That's what was called the matrix A. v itself was the n-tuple x1 up to xn, which became the column matrix X when we wrote it as a column vector. Remember that. And then what we said was the transpose of that would be c times this n-tuple.

But if I write this n-tuple as a column vector, I don't need the transpose in here. In other words, in matrix language, with A being this matrix and X being this column matrix, this translates into the matrix equation A times X equals c times X, where we recall that the matrix A is what's given, and what we're trying to find is, first of all, are there any vectors X that are mapped into a scalar times X? In other words, are any column matrices that are mapped into a scalar times that column matrix with respect to A?

And secondly, if there are such column matrices, what values of c correspond to that? Well, notice, by ordinary algebraic techniques-- because matrices do obey many of the ordinary rules of algebra-- AX equals cX is the same as saying AX minus cX is 0. Notice that we already know that matrices obey the distributive rule. In other words, I could factor out the matrix X from here.

Of course, I have to be very, very careful. Notice that capital A is a matrix. Little c is a scalar. And to have A minus c wouldn't make much sense. In other words, since A is an n-by-n matrix, I want whatever I'm subtracting from it to also be an n-by-n matrix.

So what I do is the little cute device of remembering the property of the identity matrix I sub n. I simply replace X by I sub n times X-- in other words, the identity matrix times X. That now says what? AX minus c, identity matrix n-by-n-- identity matrix times X equals 0. Now I can factor out the X. And I have what? The matrix A minus c times the identity matrix times the column matrix X equals 0.

Now, remember, back when we were first talking about matrix algebra, we pointed out that matrices obey the same structure. Matrices obey the same structure with certain small reservations that numbers obey. For example, we saw that if A was not the zero matrix, AX equals zero did not imply that X equals zero like it did in ordinary arithmetic. But it did if A happened to be a non-singular matrix.

In other words, what we did show was that if this particular matrix had an inverse, then, multiplying both sides of this equation by the inverse of this, the inverse would cancel this factor, and we'd be left with X equals 0. In other words, if A minus cI inverse exists, then X must be the zero column matrix.

Now, remember what X is. X is the column matrix whose entries are the components of the v that we're looking for over here. Keep in mind that we were looking for a v which was unequal to zero. If v is unequal to 0, in particular, at least one of its components must be different from 0. So what we're saying is, if A-- if this matrix here, with its inverse, exists, then X must be the zero column matrix, which is the solution that we don't want. In other words, we won't get non-zero solutions.

Or from a different point of view, what this says is, if we want to be able to find a column vector X which is not the zero column vector, in particular, A minus cI had better be a singular matrix. In other words, it's inverse doesn't exist. And as we saw back in block 4, when a matrix is singular, it means that its determinant is 0. Consequently, in order for there to be any chance that we can find non-zero solutions of this equation, it must be that the determinant of A minus cI must be 0.

And by the way, what does A minus cI look like? Notice that I is the n-by-n identity matrix. When you multiply a matrix by a scalar, you multiply each entry of that matrix by that scalar. Since all of the entries off the diagonal are 0, c times 0 will still be 0. So notice that c times the n-by-n identity matrix is just the n-by-n diagonal matrix, each of whose diagonal elements is c.

And now, remembering what A is, and remembering how we subtract two matrices, we subtract them component-by-component. Notice that the only non-zero components of cI are the c's down the diagonal. What this says is, if we take our matrix A, A minus c In is simply, from a manipulative point of view, obtained by subtracting c from each of the diagonal elements. You see, it's a1,1 minus c, a2,2 minus c. But every place else, you're subtracting 0, because the entry here is 0.

So this is what this matrix looks like. I want its determinant to be 0. Last time, we showed how we computed a determinant. Notice that the a's are given numbers. c is the only unknown. If we expand this determinant and equate it to 0, we get an nth degree polynomial in c.

I'm not going to go into much detail about that now. In fact, I'm not going to go into any detail about this now. I will save that for the exercises. But what I will do is take the very simple case, where we have a two-dimensional vector space, and apply this theory to the two-by-two case. And I think that the easiest example to pick, in this case, is the same example as we started with-- example 1-- and revisit it. In other words, this is called "example 1 revisited." That doesn't sound-- Let's just call it example 2.

In example 2, you were thinking of a two-dimensional space relative to a particular basis. The 2-tuple x comma y got mapped into x plus 4y comma x plus y. In particular, 1 comma 0 got mapped into 1 comma 1. 0 comma 1 got mapped into 4 comma 1. The only reason I've left the bar off here is so that you don't get the feeling that this has to be interpreted geometrically. This could be any two-dimensional space relative to any basis.

But at any rate, using this example, remembering how the matrix A is obtained, the first column of A are these components. That's 1, 1. The second column of A are these components-- 4, 1. So the matrix A associated with f relative to our given basis that represents the 2-tuples here is 1, 4, 1, 1. If I want to now look at A minus c I2-- see, n is 2, in this case-- the two-by-two identity matrix-- what do I do? I just subtract c from each diagonal element this way so the determinant of that is this determinant.

We already know how to expand a two-by-two determinant. It's this times this minus this times this. This is what? c squared minus 2c. Plus 1 minus 4 is minus 3. This is c squared minus 2c minus 3. And therefore, the only way this determinant can be 0 is if this is 0. This factors into c minus 3 times c plus 1. Therefore, it must be that c is 3 or c is minus 1. In other words, the only possible characteristic-- or, eigenvalues, in this problem, are 3 and minus 1. And we'll see what that means in a moment in terms of looking at these cases separately.

You may have noticed I made a slip of the tongue and said characteristic values instead of eigenvalues. You will find, in many textbooks, the word eigenvalue. In other books, you'll find characteristic value. These terms are used interchangeably. Eigenvalue was the German translation of characteristic value. So use these terms interchangeably, all right?

But let's take a look at what happens in this particular example, in the case that c equals 3. If c equals 3, this tells me that my vector v is determined by f of v is 3 times v, where v is the 2-tuple x comma y. Writing out what this means in terms of my matrix now, my matrix A is 1, 4, 1, 1. v written as a column vector is this. And 3 times this column vector is this.

Remembering that to compare two matrices to be equal, they must be equal entry by entry, the first entry of this product is what? It's x plus 4y. The second entry is x plus y. Therefore, it must be that x plus 4y equals 3x, and x plus y equals 3y. And both of these two conditions together say, quite simply, that x equals 2y. And what does that tell us? It says that if you take any 2-tuple of the form x comma y, where x is twice y-- in other words, if you take the set of all 2-tuples 2y comma y, these are eigenvectors. And they correspond to the eigenvalue 3.

And I'm going to show you that pictorially in a few moments. I just want you to get used to the computation here for the time being. Secondly, if you want to think of this geometrically, x equals 2y doesn't have to be viewed as a set of vectors. It can be viewed as a line in the plane. And what this says is that f preserves the direction of the line x equals 2y.

Oh, just as a quick check over here-- whichever interpretation you want-- notice that if you replace x by 2y-- remember what the definition of f was? f was what? f of x, y was x plus 4y comma x plus y. So if x is 2y, this becomes 6y comma 3y. In other words, f of 2y comma y is 6y comma 3y. That's the same as 3 times 2y comma y. And by the way, notice that the special case y equals 1 corresponded to part of our example number 1, when we showed that f bar of 2 comma 1 was three times 2 comma 1.

In a similar way, c equals minus 1 is the other characteristic value. Namely, if c equals minus 1, the equation f of v equals cv becomes f of v equals minus v. And in matrix language, that's AX equals minus X. Recalling what A and X are from before, remember, A times X will be the column matrix x plus 4y, x plus y. Minus X is the column matrix minus x minus y. Equating corresponding entries, we get this pair of equations.

And notice that both of these equations say that x must equal minus 2y. In other words, if the x component is minus twice the y component, relative to the given basis that we're talking about here, notice that the set of all 2-tuples of the form minus 2y comma y are eigenvectors, in this case, and the corresponding eigenvalue is minus 1, and that f preserves the direction of the line x equals minus 2y.

And I think, now, the time has come to show you what this thing means in terms of a simple geometric interpretation. Keep in mind, many of the applications of eigenvalues and eigenvectors come up in boundary value problems of partial differential equations. I will show you, in one of our exercises, that even our linear homogeneous differential equations may be viewed as eigenvector problems. They come up in many applications.

But I'm saying, in terms of the spirit of a game, let's take the simplest physical interpretation. And that's simply the mapping of the xy-plane into the uv-plane. And all we're saying is that the mapping f bar that we're talking about-- what mapping are we talking about? The mapping that carries x, y into-- what was it that's written down here? x plus 4y comma x plus y.

What that mapping does is it changes the direction of most lines in the plane. But there are two lines that it leaves alone in direction. Namely, the line x equal 2y gets mapped into the line u equals 2v, and the line x equals minus 2y gets mapped into the line u equals minus 2v.

By the way, notice we are not saying that the points remain fixed here. Remember that the characteristic value corresponding to this eigenvector was 3. In other words, notice that 2 comma 1 doesn't get mapped into 2 comma 1 here. It got mapped into 6 comma 3-- that the characteristic value tells you, once you know what directions are preserved, how much the vector was stretched out. In other words, 2i plus j get stretched out into 6i plus 3j.

Well, I'm not going to go into that in any more detail right now. All I do want to observe is that, if I was studying the particular mapping f bar, notice that the lines x equal 2y and x equal minus 2y are, in a way, a better coordinate system than the axes x and y, because notice that the x-axis and the y-axis have their directions changed under this mapping. But x equals 2y and x equals minus 2y don't have their directions changed.

In fact, to look at this a different way, let's pick a representative vector from this line and a representative vector from this line. Let's take y to be 1. In this case, that would say x is 2. In this case, it says x is minus 2. Let's pick, as two new vectors, alpha 1 to be 2i plus j and alpha 2 to be minus 2i plus j.

And my claim is-- I'll write them with arrows here, as long as we are going to think of this is a mapping. My claim is that alpha 1 and alpha 2 is a very nice basis for E2 with respect to the linear transformation f. Well, y? Well, what do we already know about alpha 1? Alpha 1 is an eigenvector with characteristic value 3. In other words, f of alpha 1 is 3 alpha 1. Alpha 2 is also an eigenvector with characteristic value minus 1. So f of alpha 2 is minus alpha 2.

Notice, then, therefore, from an algebraic point of view, if I pick alpha 1 and alpha 2 as my bases for-- well, I should be consistent here. I called this V. I suppose it should have been E2, simply to match the notation here. But that's not important. Suppose I pick alpha 1 and alpha 2 to be my new bases. Notice, you see, that f of alpha 1 is 3 alpha 1 plus 0 alpha 2. f of alpha 2 is 0 alpha 1 minus 1 alpha 2.

So my matrix of f, relative to alpha 1 and alpha 2 as a basis, would have its first column being 3 and 0. It would have its second column being 0 and minus 1. In other words, the matrix now is a diagonal matrix, 3, 0, 0, minus 1. It's not only a diagonal matrix, but the diagonal elements themselves yield the eigenvalues. Notice how easy this matrix is to use for computing-- A times X-- if X happens to be written relative to the alphas, because the easiest type of matrix to multiply by is a diagonal matrix.

And I'm not going to go through this here, but when you write-- when you pick the basis consisting of eigenvalues, eigenvectors, and write this diagonal matrix, the resulting diagonal matrix gives you a tremendous amount of insight as to what the space looks like. And I'll bring that out in the exercises. All I want you to get out of this overview is what eigenvectors are and how we compute them.

And I thought that, to finish up with, I would like to give you a very, very profound result, which I won't prove for you, but which I will state-- has, also, a profound name. But I'll get to that in a moment. I call this an important aside. It really isn't an aside. It's the backbone of much of advanced matrix algebra.

But the interesting thing is this. Remember, given an n-by-n matrix A, we were fooling around with looking at the determinant of A minus cI equaling 0 and trying to find a scalar c that would do this for us. That's how we get the characteristic values. A was the given matrix. I was the given identity matrix. c was a scalar whose value we were trying to get the determinant of this matrix to be 0.

The amazing point is that if you substitute the matrix A for c, in this equation, it will satisfy this equation. And what do I mean by that? Just replace c by A over here, and this equation is satisfied. By the way, that may look trivial. You may say to me, gee, whiz. What a big deal. If I take c and replace it by A, this is A times the identity matrix, which is still A. A minus A is the zero matrix. And the determinant of the zero matrix is clearly 0.

The metaphysical thing here is, notice that c is a number. It's a scalar. And A is a matrix. Structurally, you cannot let c equal A. All we're saying is a remarkable result, that if you mechanically replace c by A, this equation is satisfied.

And before I illustrate that for you, I've made a big decision. I'm going to tell you what this theorem is called. I wasn't originally going to tell you that. It's called the Cayley-Hamilton theorem. And by my telling you this name, you now know as much about the subject as I do. That's why I didn't want to tell you what the name was, so I'd still know something more than you did about it. But that's not too important.

Let me illustrate how this thing works. Let's go back to our matrix 1, 4, 1, 1, all right? The determinant of A minus cI, we already saw, was c squared minus 2c minus 3. My claim is, if I replace c by A in here, this will still be obeyed, only with one slight modification. See, this becomes what? A squared minus 2A. And I can't write minus 3, because 3 is a number, not a matrix. It's always understood, when you're converting to matrix form, that the I is over here.

And if you want to see why, you can think of this as being A to the 0, and think of the number 1 as being c to the 0. In other words, structurally, this is A squared minus 2A minus 3A to the 0 power. This is c squared minus 2c minus 3c to the 0. And my claim is that this equation will be obeyed by the matrix A.

Let's just check it out and see if it's true. Remember that A was the matrix 1, 4, 1, 1. To square it means multiply it by itself. If I go through the usual recipe for multiplying two two-by-two matrices, I very quickly see that the product is 5, 8, 2, 5. Notice that, since A is 1, 4, 1, 1, multiplying by minus 2 multiplies each entry by minus 2. So minus 2A is this matrix. Notice that minus 3 times the identity matrix is a diagonal matrix that has minus 3 as each main diagonal element-- in other words, this matrix here.

And notice, now, just for the sake of argument, if I add these up, what do I get? 5 minus 2 minus 3, which is 0, 8 minus 8 plus 0, which is 0, 2 minus 2 plus 0, which is 0, 5 minus 2 minus 3, which is 0. In other words, this sum is the zero matrix, not the zero number. You see, technically speaking here, in this equation here, the 0 refers to a number, because the determinant is a number. But here, we're talking about, it satisfied in matrix language.

And what this means is that matrices can now be reduced by long division. So I'll give you a very simple example. But what the main impact of this is, I can now invent power series of matrices. In other words, I can define e to the x, where x is a matrix, to be 1 plus x plus x squared over 2 factorial plus x cubed over 3 factorial, et cetera, the same as we did in scalar cases. And the main reason is that, once I'm given a matrix, and I find the basic polynomial equation that it satisfies, I can reduce every matrix to that.

Let me give you an example. The key thing I want you to keep in mind here is that we already know, for this particular matrix A, that A squared minus 2A minus 3I is the zero matrix. Suppose, now, I wanted to compute A cubed. Now, in this assignment here I'm going to show you how I can reduce matrices by long division. And in the exercises, I'll actually do the long division for you.

But what is long division in factoring form? What I'm saying is, I know that A squared minus 2A minus 3I is 0. So I would like to write A cubed in such a way that I can factor out an A squared minus 2A minus 3I. The way I do that is I notice that I must multiply A squared by A in order to get A cubed. The trouble is, when I multiply, I now have a minus 2A squared on the right-hand side that I don't want, because it's not here. And I have a minus 3A on the right-hand side that I don't want, because it's not here.

So to compensate for that, I simply add on a 2A squared, and I add on a 3A. In other words, even though this may look like a funny way of doing it, a very funny way of writing A cubed is this expression here. And the reason that I choose this expression is, notice that this being 0 means that A cubed is just 2A squared plus 3A. Moreover, notice that I can still get the structural form A squared minus 2A minus 3I out of this thing by writing it.

Seeing that my first term here is going to be 2A squared-- so I put a 2 over here-- this gives me my 2A squared, which I want. This gives me a minus 4A. But I want to have a plus 3A, so I add on 7A to compensate for that. This gives me a minus 6I, which I don't have up here. So to wipe that out, I add on a 6I. In other words, another way of writing I cubed, therefore, is this plus this. And since this is 0, this says that A cubed is nothing more than 7A plus 6I.

In fact, in this particular case, notice that any power of A can be reduced to a linear combination of A and I, because as long as I have even a quadratic power in here, I can continue my long division. It's just like finding a remainder in ordinary long division. You keep on going until the remainder is of a lower degree than the divisor. In this particular case, I've shown you that A cubed is 7A plus 6I. And I picked an easy case just so we could check it.

Notice that we already know that A squared is 5, 8, 2, 5. A is 1, 4, 1, 1. So A cubed is this times this. Multiplying these two-by-two's together, I get this particular matrix. On the other hand, knowing that A is 1, 4, 1, 1, 7A is this matrix, 6I is this matrix-- and if I now add 7A and 6I-- how do I add? Component-by-component. I get 7 plus 6, which is 13, 28 plus 0, which is 28, 7 plus 0, which is 7, 7 plus 6, which is 13. In other words, 7A plus 6I is, indeed, A cubed, as asserted.

And I think you can see now why I wanted to end here. From here on in, the course becomes a very, very technical subject and one that's used best in conjunction with advanced math courses that are using these techniques.

So we come to the end of part 2. I want to tell you what an enjoyable experience it was teaching you all. If nothing else, as I told you after part 1, I emerge smarter, because it takes a lot of preparation to get these boards pre-written. I couldn't do it alone, and I would like-- in addition to the camera people, the floor people, there are three people who work very closely with this project that I would like to single out.

I would like to thank, especially, John Fitch, who is the manager of our self-study project, who also doubles in as director and producer of the tape, the film series, and is also my advisor for the study guide and things of the like. I would like to thank Charles Patton, who is the one responsible the most for the clear pictures and the excellent photogenic features that you notice of me, the sharpness of the camera.

I would also like to thank Elise Pelletier, who, in addition to being a very able secretary, doubles in in the master control room as the master everything, from running the video tape recorder to making hasty phone calls and things of this type.

I would also like to thank two other colleagues, Arthur [INAUDIBLE] and Paul Brown, administrative offices at the center, who have provided me most excellent working conditions, and finally, Harold [INAUDIBLE], who was the first director of the center and whose idea it was to produce Calculus Revisited Part 1 and Part 2. It has been my pleasure. I hope that our paths cross again. But until such a time, God bless you all.

Funding for the publication of this video was provided by the Gabriella and Paul Rosenbaum Foundation. Help OCW continue to provide free and open access to MIT courses by making a donation at ocw.mit.edu/donate.