PROFESSOR: So today the basic idea is to try to understand the toggle switch. How such a thing can be made, why it represents a memory module. But then we'll relatively quickly get into two themes that are going to be useful throughout the rest of the semester. First is these dimensionless equations that often pop up in the analysis of these gene circuits. And it's absolutely essential that you understand how to get to the dimensionless equations, and also what these parameters end up meaning.

And then we'll, at the end, talk about stability analysis. How is it that you can determine whether a particular set of interacting pieces in a cell or in an ecosystem whatnot-- once you get it in the form of an equation, how's is it you can determine whether a particular fixed point or a particular location is going to be stable to perturbations? This is going to be useful for us to determine where the gene network is going to kind of move towards. Also it'll be useful for us to determine whether a gene network is going to oscillate. And later, it'll be relevant in the context of predator prey oscillations, and a bunch of things.

So can somebody maybe explain briefly the idea of the toggle switch? Yes, please.

AUDIENCE: You have these two genes that repress expression to each other. So when one of them is large, the other one is not expressed. [INAUDIBLE].

PROFESSOR: Perfect. So you'll have two genes that are going to mutually repress each other. So in this case, each of them will be a transcription factor of some sort that will bind to the other promoter and repress it. So there are various levels of abstraction that we might use to describe such things.

So we might, for example, just say A repressing B, B repressing A. Of course when

you write it that way, it doesn't have to be in the context of a gene network. This A's and B's could just be chemicals, they could be species eating each other. It could be almost anything.

Now in this framework, to get a basic sense of why this thing might have two alternative stable states, is that often we like to take the Boolean approximation. This thing will not always work, but it's a useful thing to do to just kind of first get a sense of what might possibly be happening. So we might say, 0 corresponds to some sort of low. 1 might correspond to high.

And of course, these things have to be put it in quotes, because we haven't specified what we mean by this. But it's useful to just make sure that we're all thinking about the same things. And then in the context of A and B, we can just say, well, there's a number of different states it could possibly be in. And you can ask whether this assignment of logic values keeps everybody happy.

And so you might ask, well, is 0, 0 a mutually happy state? And of course then you have to say well, if you're in the 0 state, you're not repressing the other guy, but maybe 0 is sort of your equilibrium anyway. So what we have to do is we have to make the assumption that when you're not being repressed, in that case the promoter will be actively making that protein. So then you'll go to some sort of high or 1 state. So in that case, you say, if you start out with both repressed, maybe both of them should start trying to increase their levels. So this, in some ways, is not a stable state.

Similarly here, this is also not a stable state. Because in this case, they're both going to be trying to repress one another. So then they'll both start coming down and then the situation may resolve into one of these two. So this is just where either A or B is on, and repressing the other one.

And so for example, in context of the repressilator, on Thursday, this is just a useful way to start imagining how this A repressing B, repressing C repressing A-- how such a loop can lead to oscillations. This kind of analysis does not at all prove that there are oscillations in any given manifestation of this thing. But it's useful to just

make sure that you're roughly getting the idea of what the system might be doing.

Now in many cases, we'll want to be a little bit more explicit, and draw the gene network in more detail. And there are multiple manifestations of the Thomas Switch. There are many of them that have been made. So the important thing is not too necessarily keep track of exactly what the components are, but in one case for example, we might have A corresponding to something here. It's coming back and repressing again.

Expression of this B. And this might all be on one piece of DNA. Whereas this B here will come back and repress A.

Now one thing that is-- and I just want to mention here, they also have a GFP. And this actually is a case where those two can be expressed off of a single promoter. So this is often done in bacteria where there's a single promoter-- so RNA polymerase actually will transcribe both of these genes. This repressor B, as well as this fluorescent protein. So there are going to be alternative loading sites for the ribosome in that case. And eukaryotes typically do not do this. Yeah?

AUDIENCE: [INAUDIBLE].

PROFESSOR: That's right. So--

AUDIENCE: And it will transcribe everything until--

PROFESSOR: Exactly. So here comes-- so an RNA polymerase down here made this whole thing. And now you might have two separate locations where the ribosome loads and makes this protein B. And then a different ribosome would make this GFP.

AUDIENCE: And when does it stop? It just keeps going down?

PROFESSOR: Yes. So there's a termination sequence.

AUDIENCE: No, no, but it would B and then GFP, and then it would just keep going?

PROFESSOR: No. The ribosome is told to basically start here and end here. So then it just makes

the B protein. And then another ribosome binds here.

AUDIENCE: If you had more proteins on the same strand after GFP--

PROFESSOR: And when you say protein, you're referring to the ribosome.

[INTERPOSING VOICES]

AUDIENCE: I mean G's right?

PROFESSOR: Oh! OK, you're saying if there were another gene?

AUDIENCE: No, if you have more genes coded, if you're coding for more proteins after GFP-- if you have more genes on a

[INTERPOSING VOICES]

--it will just keep going for an arbitrarily long--

[INAUDIBLE].

PROFESSOR: Arbitrary is always a dangerous word. But they can be more than two. And actually at the biophysics retreat that some of you guys were at just last two days, there was a great talk by Gene-Wei Li, who's going to be a new incoming biology faculty member. And he was talking about the FO F1 ATP synthase. So it's the thing responsible for making ATP.

He analyzes process work. And there are many sub units. So there are half a dozen or so. And so it's a very long transcript. And then what he showed is that actually you have different rates of synthesis of the different genes on this one transcript.

And in some cases you want actually more copies of one of the subunits than another one. And so then actually if the final protein, if it needs 12 of these, only one of these, then actually you can make 12 times as much of this, because you just have more translation here than you did here. And then it's great, because then you have all the right ratios, all the components to make the protein. So you can actually have additional regulation even at that stage.

It's possible not everybody followed that discussion, and my apologies. But feel free to just erase it from your brain if you're too confused. But what you need to keep track of here is the level of GFP is going to be perhaps proportional to the level of B. Because they're being expressed at the same time.

Now in order for this thing to be a memory module, you also want to be able to reset the state. So if you were in this state, you'd like to be able to somehow get it to move to this state instead. Does anybody remember what the inputs were in the context of the sample toggle switch, that it was in that review?

AUDIENCE: [INAUDIBLE].

PROFESSOR: Right. OK, so there are multiple versions of toggle switch. And indeed in one case this was just a small molecule, IPTG, and that's because this was in that case, the lac I. And this then represses the repression. So in many contexts this class, you'll have to remember that a minus, minus is equal to a plus. And in different toggle switches you indeed have different ways of inhibiting this repression.

So this could be another small molecule, ATCN. In the example that they had in this review, it was actually heat that did this. But that's just because this transcription factor was a temperature sensitive mutant. So above some temperature, it could no longer repress.

So for example, if you start out in high GFP-- so in this case, where they-- all right. If the cells are in high GFP, and you want to switch it into the alternative state, what stimulus do you want to apply here? So I'll give you a guess. It's going to be either heat or IPTG. I just want to make sure that we can all read these diagrams. This is to switch from the high GFP, and we want to go to the low GFP. Give you 15 seconds to just try to read off this diagram.

Do you need more time? No. Ready. Three, two, one. So we have a majority are B. So a majority of people are saying, well, in this case, you have a lot of GFP. Means you have a lot of this protein B. In this case, the lac I.

And we don't have very much of this other protein. So that means that if we want to switch the state, and get a lot of A, we have to stop this repression. So we have to add IPTG. Any questions about what we mean by the various symbols up on this board?

So after we had IPTG, then indeed the GFP should go down. How long is it going to take for it to go down? Does this switch go down immediately? After you add the IPTG, how long do you think it's going to take for this repressor, lac I to fall off of that promoter?

Do you think it's going to be seconds or hours? It's actually seconds, and that's because this IPTG rapidly can go across the membrane, it'll rapidly bind to the inhibitor lac I, and then lac I will fall off. So in this case, lac I is actually still present, so it's going to take hours, actually for lac I to go away. But it takes seconds for the lac I to become inactive, ineffective.

So this is the separation of time scales idea. But how long is it going to take for the concentration of-- well, how long is it going to take for GFP to go away maybe? Is that going to be seconds or hours?

AUDIENCE: [INAUDIBLE].

PROFESSOR: Yeah, In this environment the bacterium-- they might be dividing every half hour. That's the characteristic time-- if these are stable, that would be the characteristic time scale for example, for GFP to go down. And of course, that's even after it's inhibited. After you stop making GFP.

In this case, once you had the IPTG, the lac I is going to fall off of this promoter, and then we have to first make A. And only after we make A will we start repressing expression of the B and GFP. That's the process that you expect to take, hours, based on the generation times.

But this is a memory module because after we've added the IPTG, now we can, in principle, take the IPTG away, and it'll stay low. So in principle, this IPTG signal can be a transient signal. And the cells will remember that they encountered IPTG.

That's what makes this a memory. This input can be transient. Transient signal is remembered.

Now a big part of why this why this paper was important is because this toggle switch was constructed out of components that previously in principle, had not even ever seen each other before. So maybe that lac I and this promoter had been put together, but this whole system, this whole gene network was composed of individual components that they were synthetic. And were put together because they thought, oh this should kind of work.

And it's led to, I think a real flowering of, this intersection between modeling and experiment. This thing was built based on a model that told them that maybe if we do this, these things they multiplarize in order to repress and so forth. So there was a real sense in which the modeling-- and then we're going to talk more about the modeling-- on how it was essential to get an idea of how we should construct this thing, and to guide our work. Because if any of you do work in the lab, you'll know that things are hard. And often components don't behave the way you think they should and so forth.

Now modeling can't save you from all that pain, but at least it can guide you in the right direction so that it limits the number of things you have to try. Are there any questions about this element, before we switch over to some of the dimensionless equations that we're going to be using?

So the reading for today was composed of a review that these two pieces, and one of which I think might be hard for some people, and the other one might be hard for other people. For those of you with maybe more limited experimental experience in biology, maybe the review was a challenge to try to understand all of the nomenclature and the words, whereas those of you that have not played with differential equations as much recently, may have found the reading on modeling the toggle switch and stability analysis to be more challenging. We will, in some cases, do this where we have two different hopefully shorter kinds readings.

And hopefully they're not both hard for you, because then you'll spend a lot of time

reading. But then it will be good for you. So don't run away quite yet. But from my standpoint, it's essential that you develop intuition behind these ideas.

So once you have this equation that somebody gives you, you have to be able to figure out what assumptions have they made, and what should the behavior be roughly, before you go off and you do simulations and full mathematical analysis. So a lot of what we're going to do in this class, and in particular during the lectures, is to try to work on that intuition. And the first thing you need to do is make sure that you know if you don't understand something. Sometimes things look really simple, especially these dimensionless equations.

Part of what's attractive about them is that they are simpler. But the connection to experiments can be quite challenging. And we'll kind of see some of this.

So these dimensionless equations-- dimensionless equations-- I think that they're both good and bad. The good is that you can figure out what are the essential features of the model. So you can focus your attention on the essential mathematical features.

The disadvantage is that you might not even know which of the parameters change when you do something experimentally. So the problem is that the connection to the biology or experiments is obscured in some cases. Connection to experiments.

And as an example of this, what we want to do is look at these equations for the toggle switch. I know that you guys just did reading on how to get to these final pair of equations. It's important to remember that in the context of the paper, they had a model. Here are the dimensionless equations that describe our system. And we use them to design the toggle switch.

But just from that, you don't necessarily realize what's been done. So we want to make sure we understand it. So the equations that we want to be comfortable with are the following. So it's u dot du dt. Now here they use alpha as the rate of expression.

So beware, in the past, we sometimes used alpha as a rate of degradation. So fair

warning. Alpha 1 1 plus that's v beta. v to the beta minus u v dot is equal to-- here is an alpha 2 divided by 1 plus u to the gamma minus v. Now in many, many cases, we're going to get-- equations that look like this are going to pop up time and time again over the rest of this semester, and you just have to be intimately familiar with them.

So first of all, can somebody say why this might be capturing the dynamics of a toggle switch? I can say something. Yes, please.

AUDIENCE: There's some more of each u or [INAUDIBLE].

PROFESSOR: That's right. So the more v you have, the less production you're going to have of u, and vice versa. Now in many cases, beta and gamma are going to be something that's larger than 1. This is capturing some element of cooperativity in the repression on each side. So this thing is indeed just a u repressing a v, and vice versa.

Now these things are wonderfully simple equations. See that there are four parameters that are completely specifying the dynamics. Now you'll notice that the world is always to be more complicated than the four parameters. Now there are two ways in which these complications went away. One is that we are modeling a simple-- it's a simple model of a complex system.

But the other is that even the simple model has been simplified by going to this dimensionless version. So I want to make sure that you understood the reading to the point where at least you understand what units of concentrations, times, and so forth are in this model.

So first I want to ask about the effective lifetimes of proteins u and v. Of u verses v. So we'll maybe call this tau u and tau v. And I want to know which one-- and how these things are related to each other. Tau u greater than [INAUDIBLE]. DK is again, don't know.

Now of course in principle, the lifetimes of u and v can be anything. The question is,

once we've written down that equation, have we actually said anything about that? Have we already made in assumption or not?

Do you need more time? I see a fair number of quizzical faces, which may mean that even with extra time, the quizzical faces would not go away. Question. Yes?

AUDIENCE: Are you asking about the lifetime of individual [INAUDIBLE]?

PROFESSOR: OK. All right. Yeah. So when I say effective lifetime, that's because if it's a stable protein, then the lifetime of that individual protein is maybe infinite, but you get an effective lifetime because of this dilution effect. So when I say effective lifetime in general, in this class, what I'm referring to is the sum of two effects of dilution to the cell growth, as well as actual degradation.

AUDIENCE: And you're saying [INAUDIBLE].

PROFESSOR: No. No, I'm saying given that the Collins Lab-- in their paper, wrote down those set of equations. I'm asking have they already specified anything about the effective lifetimes of these two proteins?

AUDIENCE: So in those equations, [INAUDIBLE]. Your question asks in units of that dimensionless time?

PROFESSOR: Yeah, I mean we can compare two times in some dimensionless-- yeah, we're going to also talk about that. Maybe I should have done that first. But this is a nice question because it's highlighting that you get a pair of equations, they look obvious, but then some of those basic things about the system are some how not quite clear. So I'm going talk about in this units of whatever time is being-- however time-- we're going to discuss how time is being measured in a moment as well. But however time is being measured, is there some relationship here or not?

Let's go ahead and vote. And it's fine, if you really do not know what I'm talking about, you can say E and that's fine. Let's see where we are. Ready, three, two, one.

So I would say that it's split between B's and D's. And that's great, because that

means we have something to talk about. There's broad agreement that it's one of these two.

So turn your neighbor. You should be able to find somebody that disagrees with you. If you can't find anybody that disagrees with you, you could think about how parameters are going to change as you vary other experimental things.

[CLASSROOM CHATTER]

PROFESSOR: Why don't we go ahead and reconvene? I just want to see where we are. So let's get our cards ready. And we're still working on this one, so you can ignore this.

Is it B or D? Ready? Three, two, one. So I'd say there's some migration towards B, but not 100% still. So I'm going to side with this here. Can somebody volunteer why they're saying that?

AUDIENCE: Well, if I remember correctly then they [INAUDIBLE] time by multiplying time by degradation rate. And they do that with the same degradation rate for both [INAUDIBLE].

PROFESSOR: Right. So the answer here is yeah-- so the way that we got to this non dimensional time, that we're going to discuss in a moment, is multiplying by-- or divided by some degradation rate. And you remember the derivation. You remember that it was [INAUDIBLE].

But in many cases you don't get to read the derivation before you have to answer it. In many cases you just get these equations, and you have to figure out what the authors have assumed. So I think your answer is very much correct, but it might be-- I'm glad that you're using the pre-class reading, but at the same time we have to be able to answer this question just from these equations. Because it is contained there. Yeah?

AUDIENCE: So if you remove production [INAUDIBLE].

PROFESSOR: Perfect.

AUDIENCE: If you solve that equation, it's exponential to E, and the [INAUDIBLE].

PROFESSOR: Yes. That's right. So the statement here is that it's nice to just imagine that we shut down production. So these first terms go away. All right now we just have u dot is equal to minus u.

So the concentration of u and B will both fall exponentially. And they're going to fall the same rate in whatever units of time-- and we'll discuss this in a moment-- but however time is being measured, they're going to fall at the same rate because it's whatever appears in front of these two that determines that rate. So we've already assumed that the effective lifetime of u and v are the same once we've written down these equations.

Now that doesn't have to be true. That means that if experimentally, you want to make a toggle switch with two proteins, with different st abilities, then you can't use these equations. You have to add a delta on one of the two equations to capture that dynamic that they have different lifetimes.

So these equations are simple because we've already combine things, but they've also already assumed some things to make it look more simple and more symmetric. So for example, they're allowing for different effective cooperativities, beta gamma, for the two repressors. They didn't have to do that if they didn't want to. If they wanted to, they could have just had beta in both equations.

And that case they would be assuming that the effective cooperativity of the repression is the same for the two. So depending on what you write down, you're making different assumptions about the system. And so you have to be able to look at these equations and figure out what assumptions have been made. Yes, question.

AUDIENCE: I'm having trouble seeing what the first term is doing. So the second one is [INAUDIBLE].

PROFESSOR: So broadly, the first term in these equations is the production rate of the protein.

And the second term is some sort of effective degradation, but it could be due to dilution. The nice thing about just ignoring the production term is that it's a way of focusing only on the effective lifetime portion. This effective lifetime is it's captured here, irrespective of whether there's production or not.

Of course the production is going to affect what the equilibrium is that we go to and so forth, but remember for example, that this effective lifetime that's set for time scale both to come up to some equilibrium, as well as come down to some equilibrium. That's telling us about this effective lifetime is essential to tell us about the rate that concentration is going to change within the cell, whether you're going up or down. Did that answer your question a little bit? Yes?

AUDIENCE: That last thing you said is not quite true because--

PROFESSOR: Because of the toggle. I agree. What I was really referring to there was if we get rid of one, and we're just talking about where we just manually set the production rate to one thing or another. The actual dynamics of this are more complicated, certainly. So is everybody happy with this idea that just by writing down those equations, we've made some assumptions constraining relationships between u and v in some ways, but not in other ways. Yeah?

So we've kind of already danced around this other question, but I want to make sure that we address it head on. We want to know what is this unit of time. If I say oh, at time T [INAUDIBLE] 1 versus time [INAUDIBLE] to 2-- so if delta t-- what is 1 in this case? Are we referring to one second? CGS.

Or is it one hour corresponding to another unit? Cell generation time, effective lifetime, or don't know? Ten seconds, since we've already kind of said this, but it's important enough to make sure.

Ready? No? Yes? I'll give you another ten seconds, just to make sure.

Do you need more time? Let's vote. Ready? Three, two, one. OK.

A majority of the group is agreeing in this case it's the effective lifetime. And it's not

the cell generation time necessarily because if we actually have a degradation tag on this protein, if it's not a stable protein then the effective lifetime is going to be shorter than the cell generation time.

So we've normalized time by dividing out that-- whatever that delta thing would have been on the right. Now this is great because it makes the equations simple. But remember what that means is that if we change the degradation rate, then all of a sudden it's not obvious what's going to happen. Experimentally we're always allowed to affect the degradation rate. But the question is, which parameter or parameters change if you add a degradation tag? So if you increase the degradation rate.

That's what I want to do. We'll do that in just moment there. I have another question that I wanted to do. Do you guys remember the equations up there? Maybe I'll come over to the other side just so that I-- it's useful to be able to stare at those equations as we discuss.

So the question is, what is the unit of concentration of protein u? And of course, these are dimensionless. What I really mean is what does u equal to 1 mean in real units? I'll just write that.

What does u equal to 1 mean? What does u equal to 1 correspond to? In something that is recognizable to an experimentalist in the lab?

You guys understand what I mean? So if in this model I say, u is equal to 1, or u is equal to 10. What do those numbers-- what do they mean? So are there any questions about my question first of all? Yes?

AUDIENCE: Are we [INAUDIBLE] solely on seeing those equations? Or based on the reading of [INAUDIBLE]?

PROFESSOR: I would say that the reading actually has a somewhat more complicated model, and involves multiple steps. So this is really just looking at those equations, given that we can see that there is some statement about how u and v are repressing each other. So we're saying that there is some function that describes-- it's a

phenomenological function-- describes the input output relationship in terms of some concentration of u leads to some repression of v and vice versa. From that actually, we should be able to say something about what we've already assumed in these equations.

Do you need more time? Yeah, question.

AUDIENCE: So the k of u promoter is-- that's defined as the--

PROFESSOR: That's the binding of v to the u promoter-- the piece of DNA in front of the u gene. Do we need more time, or shall we give it a go? Ready? Three, two, one. All right.

So we got some A, B, C, D's. No E's. At least it means that your neighbor has an opinion. He or she cannot say that-- Turn to your neighbor and discuss.

[CLASSROOM CHATTER]

Did you guys all decide you're comfortable? Why don't we go ahead and reconvene. It seems like it may be we're coming to consensus here. Ready? Three, two, one. So now we got a clear majority agreeing that it should be D.

So what happened was that over here in the original equations we had a real concentration of u divided by some k. And that was the k for repressing v. Because v is this guy that is over here. And what we did is we then just turned u divided by this k for binding this v promoter into just u. So in particular, when that concentration of u is equal to its associated k, that corresponds to half repression.

Same thing here. One way to think about this is just that when u is equal to 1, we're getting half of maximum possible repression. Similarly, that's what v equal to 1 means. U and v, do they have to have the same-- I mean, does u equal to 1 and v equal to 1 mean the same thing in terms of the number of the proteins in the cell? No. Not necessarily.

So we've allowed for the possibility that those things are measured in different real units. But in both cases, it's telling us about how the strength of repression. And u

and v equal to 1 tells us about that crossing point where the other promoter is half repressed. Are there any questions about what happened there? This is especially confusing because it doesn't enter into the equations at all.

Things just went away. But you know that this is the dimensionless versions of the equations because we have u here, and we're adding 1 to it. Are we allowed to add things with different units? No. Never.

What that means is that-- since they're being added, that means that we already know that we made-- this is the dimensionless version of the equation. And it actually then immediately tells us what u equal to 1 means.

So now what we want to do is we want to make sure that our intuition on this is tip top shape. In particular, we want to know if I want to change the dynamics of the system-- so let's say that we go and we spend lots of time calculating the fixed point stability, and we do everything on these equations. And then we know what we need to do is we need to change some parameter in order to get say, a toggle switch. We need to know how we do that. We need to know how the parameters in real life, or even in the context of the model, how is it that the parameters you can actually change experimentally, how is it that the affect or not the parameters in this model?

So the question is, let's say we increase the degradation rate of these two transcription factors u and v, which of the parameters are going to change? So I'll give you 30 seconds to think about what should be happening here. Do you need more time? Let's see where we are.

Ready? Three, two, one. And this is the possibility where you can put up two things. Remember our fabulous card system? So I think most people are saying it's going to be A and B are going to change.

So beta and gamma are capturing how cooperative that transition is. And the cooperativity is not affected by the questions of the exact concentration, or time scale and so forth. Because that has to do with the molecular nature of the interactions with the promoter. So in some ways beta gamma are the simplest

things in this system.

Now the question is, do alpha 1-- do they go up or do they go down? Now that I've told you that they change. If degradation rate goes up, alpha 1 and alpha 2, do they go up or do they go down? I'll give you 10 seconds to think about it.

Do you need more time? Let's see it. Ready? Three, two, one. It's a majority are saying it's going go down. Can somebody offer up an intuitive explanation for why this might be? [INAUDIBLE].

AUDIENCE: The degradation rate goes up, it means the times scale goes down. As you produce a fixed number of per unit time, the unit time gets smaller and produces less.

PROFESSOR: Right, so if the degradation rate goes up, that's kind of reducing this unit of time. So you just are not going to make as much protein in that unit of time. The way that I like to think about this maybe is that if the degradation rate goes up, then the real concentration of the protein should go down.

And that means that the repression should be less effective. And then wait, is this helping me? No, now that I'm saying this-- I don't like to think about it that way.

[INTERPOSING VOICES]

That's right. Yeah, so-- that's right. You should decrease the sort of-- So if the degradation rate goes up, you decrease the real concentration, which it means indeed that you are at steady state. Say a less effective repressor. That means the concentration goes down in these units of how repressive are you.

I think the way of thinking about it was correct, just the words were not. You can also then go-- and it's useful in the context of when you actually go and you do the math of removing all these parameters, just to make sure that-- because it's easy to do this thing where you just divide everything out, and you're happy, and whistling and so forth. But at the end of the day, you really just have no sense of what happened. Of how the parameters that come out of the model are affected by the real things that you can change.

And in some ways what's funny about these equations is that essentially everything is in the alphas. Because beta and gamma, that's this cooperativity parameter. That's what it is. That's all it is. So that means that everything else ends up in the alphas.

So the strength of the promoter, the concentration you need to repress, the lifetime, everything rolls up in these alphas. And this is the beauty of the dimensionless equations. Is it's telling you that you don't have all these different, separate knobs. You can't change one and change another, and go into some funny regime. Because it's all rolled into one fundamental parameter, these alphas.

That's telling you that once you understand how these equations behave, then you in principle understand everything that could possibly happen in that simple model. But that's what's wonderful about the dimensionless equations. But the problem is that you sometimes lose track of how it's related to the real experimental things. So I would say that I very much like these dimensionless equations. They clarify things for you.

But you have to spend the time to make sure that you understand where all of reality went. Because it all ends up in this mathematical equation, and that simplifies things. You're not floating in a sea of symbols anymore, but it's really easy to lose track of the connection to real measurements. And that's why we want to do modeling so we can make that connection. So do the dimensionless equations, but make sure that you play with them a bit so you know what changes when you change what.

And we'll actually see a bit later, some other cases where it's actually quite tricky to figure out what's happening. On exams I always want to just ask a equation like this, if this parameter goes off-- and then the TAs always say, it's too hard of an question. I feel like it's like the most basic thing you would want from an equation like this. Is that if you increase the strength of this--

But it actually is-- it's surprisingly difficult. So maybe this year I'll convince the TAs that it's an OK question. Are there any questions about where we are right now?

So what I want to do for the last half hour is talk about stability analysis. The first context in which stability comes up in this class is indeed in this toggle switch. But it's not-- I would say-- the most satisfying application of it in some ways. Just because you end up with equations that all you can do is plot them. And it's useful to be able to recapitulate, to understand how the figures from that paper come about.

But at the same time it's not always-- after you find the solution, you don't feel so happy about it either. Because in terms of complexity as a function of time, things always start out simple. And then in the course of the calculations things get complicated. And the problems that are fun to solve are the cases where it comes simple again. This one-- it never quite converges.

I do want to talk about the stability analysis so that you can understand the calculation that was in the notes. But also so that you can just get some more intuition about some of the other problems that we're going to be solving in the next few weeks. Just to make sure that we're all talking about the same thing, it's useful to start by just making sure that in a one dimensional problem, we understand what we mean by stability. We're going to be fast.

X equals 0 is stable if and only if what? I'll give you 10 seconds. Ready? Three, two, one. All right.

So this is always-- So there's actually a fair number of answers here. And this is tricky because the temptation is always to jump into the two dimensional stability analysis, or the n dimensional stability analysis. And the thing is that we have to make sure that we are completely comfortable with a one-- talked about one dimension before you talk about multiple dimensions, because then everything is lost. I'm not going to have you guys discuss, the but it's going to be C. Many people are saying A or other things. There is a context in which this guy comes in. But let's just make sure.

So x equal to 0. Now first of all, is this thing-- is x equals 0, is it always a fixed point

of the system? Yes. So fixed point means that if you go right there, then in principle you don't move off it. So it doesn't say anything about whether it's stable or unstable. But if x equals zero, then indeed x dot is equal to 0. So that's a fixed point.

But if you have positive x-- in order for that to be stable, you have to have a negative change in x. So if you talk about the behavior as a function of time, this function of x here. If you start out above 0, the definition of stable is if you go a little bit away, you should come back. And that means that x dot has to be negative. So if positive x-- and similarly if x is negative, you want it to be stable, then you need the x dot to be positive.

So this is-- A less than 0. Now there is a context in which a stability condition looks something a little bit more like this. And can somebody say when it is that you get something to look-- condition around 1? Yes?

AUDIENCE: Discreet.

PROFESSOR: Yeah, in discrete maps then indeed the condition for stability looks something like this. If you have something that looks the x of t plus 1, is equal to axt, then if the condition for x equal to 0 being stable is for indeed a to be less than 1. Or it's in this case it's really the magnitude of a being less than 1. Because in that case you might get hopping. But then the condition is around the 1 thing, whereas here 0 [INAUDIBLE] indeed.

Solution is different-- it's going to go exponentially to-- so x is a function of time is just going to be some x naught e to the-- and is it at, or a is less than 0. It's just at, right? And in this case, we just have one dimension. So the eigenvector [INAUDIBLE] is really just x. There's just one eigenvalue, and a is indeed the eigenvalue.

So the stability for x equals 0 to be stable is that all-- and in general, for n dimensions, is that you need all the eigenvalues to be less than 0. And in this case, a is indeed the eigenvalue. And you need that to be less than 0.

So if you're not comfortable with this, or you got something other than C, then I think

it's essential that you take the time now to go over all of these ideas. First in one dimension. Make sure you're all comfortable with that. But then to look again at these two dimensional, high dimensional stability analyses. Because if these things you're finding tricky just because it's been a while since you looked at this stuff, that's fine. But if you don't spend the time to iron out now, then everything gets much more painful later.

I highly recommend Strogatz's book on dynamical systems. It's a beautiful book with just as clear as a textbook could be. So that book is available at various libraries and reading rooms and so forth. So it's a great reference.

Now we want to do is generalize this idea to think about in n dimensions. So now what we have a vector x. And there's going to be some matrix A that-- oh, and x dot. So the change in this vector x is going to be described by a linear set of equations, so that they can be specified by some matrix.

To determine the stability, we're going to then look at the eigenvalues of this matrix. Now in many cases in this class, what we're going to do is we're going to find that some set of non-linear equations has a fixed point somewhere, and then we're going to linearize around that fixed point to convert into a linear problem that looks like this. But for now what we're going to do-- the general statement is for n n dimensions. You want all the eigenvalues-- lambda I in particular-- to be the real part to be less than 0. And that has to be true for all the eigenvalues.

We're going to be talking about the conditions in a two dimensional system where there are in principles and shortcuts, but it's all the same thing. It's all that you need. The real part of all the eigenvalues be less than 0 for the fixed point be stable.

So for now what we're going to do is we're just going to imagine that we have-- assume we have two dimensional problem. This vector x, we're going to write as x and y. So we can write the dynamics like this. This is x. This is x dot, y dot.

And this is really the same thing as saying that x dot can be described by some ax plus by, and y dot is some cx plus dy. And so we want to be as comfortable as we

can with all the possible things that could happen in these two linear and differential equations. And particularly we want to know what the condition for 0, 0 being stable? 0, 0 is stable if and only if-- now there's a rule that you found in your reading having to do with the trace and the determinant.

So trace is the sum of these two. Determinant is product minus the product here. Now this is not the kind of thing that you have to memorize. But it's useful to know that there is a simple rule, because it allows you to quickly determine whether something could possibly be stable. And you don't have to memorize it, but you should be able to figure it back out later.

So what we have is the trace of this thing being less than 0-- so I'm going to give you some options. Now of course, you can always look at your notes and that is not going to help you. You're not being graded on your answers right now. I'm going to encourage you to try and think about it and see if you can recapitulate what this thing should be. So it's less than, greater than--

So for example you can start to think about-- you should be able to write down some equation that you know is stable. And figure out, OK, what conditions would it satisfy. That's a common, useful trick to be able to do in life. To dredge these things out of memory. Do you understand the question?

I'm going to give you 30 seconds because it's well worth trying to figure out what this rule should be. So from the reading you know there's some rule about the trace and the determinant. And indeed they both have to be true in order-- this is an and sign. May be an and sign.

Do you need more time? It's OK if you haven't actually been able to recapitulate this. But it's useful to think about it. Should we go ahead and vote just to see? I'm just curious where we are.

Ready? Three, two, one. So we have a lots of B's. And can somebody give me an example of a matrix A that really ought to be stable? Yeah?

AUDIENCE: Negative identity.

PROFESSOR: Negative identity, OK. So if the matrix A is something that's minus 1 here. 0, 0 minus 1. Then x and y are uncoupled. x decays exponentially, y decays exponentially. Now we can just directly say, well, this, the trace is definitely negative, and the determinant's positive.

It's 1 minus 0. So that gets us here. Because I think there are many situations in life that are like this, where you know that there's some rule and you can't remember what it is. You don't need to actually remember, once you know that there's the rule, then you can figure out what it had to have been.

This is not a proof. The proof is only a few lines. You can do it, but the point is that this kind of situation comes up a lot. It's useful to just have simple things that you know what the answer has to be, and then that allows you to figure out where things were.

And indeed what you'll find is that for a two dimensional system, this condition that the trace of the matrix is less than 0 and the determinant is larger, that is equivalent to the statement that the real part of both eigenvalues is less than 0. And there's the derivation is simple, and it's in your notes. Are there any questions about where we are right now?

So what I want to just say a few more things about the trajectories here. Can somebody explain the notion of what's an intuitive statement about the eigenvectors that you get out here? Why do we like eigenvectors? What are they useful for?

AUDIENCE: Decoupling the differential equations.

PROFESSOR: Right, so they're decoupling, and there's a sense that-- I guess the way I like to think about this is that you have a fixed point say, like this. Now if these things are stable, that means the arrows are all say coming in. Are these eigenvalues-- are they purely real, or are they complex? Reminder, somebody?

Real and negative. There's no imaginary component. This is a stable fixed point with real negative eigenvalues. Now the idea of these eigenvectors is that these are

the two directions in which if you start out on one of them, you'll stay on them.

And in general then you can-- any other trajectory you can decompose as a combination of the pads on the two. The position as function of time can always described as the sum of the eigenvectors where you grow or you shrink along each eigenvector exponentially. Now for this thing to be stable, all these lambda I's are going to be negative. We have this property that-- that the dynamics of this matrix kind of keep you along the direction of the eigenvector.

What this is saying is that if you start somewhere random, that is off one the eigenvectors, then you can decompose the trajectory along each. But I just want to highlight that it's often useful to draw what these things end up looking like. Now I want to make sure I get this one correct. I'm a little bit worried I'm going to do something funny.

So here-- we're going to say this is v1 and this is v2. So this direction is one eigenvector, this direction is the other. So let's imagine that the trajectories look like this. The question is which eigenvalue is closer to 0?

Is it A eigenvector v1, or B eigenvector-- So this is really lambda 1 verses lambda 2. 1 is closer to 0 than is 2. Which of the eigenvalues is closer to 0? I'll give you 20 seconds to think about this because it's useful to be able to extract these things from the trajectories.

Do you need more time? And it's fine if you don't really understand how you can get to this question from this figure, then go ahead and flash C, D, or E, just so I know where we are. Let's vote. Ready? Three, two, one.

All right, so there's at least a majority are saying that it's going to be lambda 2. Can somebody offer why that is? Yes?

AUDIENCE: I sort of see it as whichever one is bigger is going to be more effective at squeezing things toward the origin along that axis.

PROFESSOR: OK and bigger-- and you're saying more negative, in this case, you're saying? Yes,

right. So yeah that's right. So there's some notion that lambda 1 has to be more negative than lambda 2 because you first collapse along the direction of eigenvector 1, and then you slowly come in along the direction of eigenvector 2. That's saying that you have these two exponential decays, and the directional on the eigenvector 1 is more rapid than the directional on eigenvector 2.

I think in many of these cases in dynamical systems, differential equations, it's hugely valuable to be able to draw the trajectories. So in many cases, what we're going to do over the course of the semester is we're going to have a simple pair of differential equations that are going to be two proteins, or they're going to be rabbits and foxes or whatever. So we're going to locate where the fixed points are. And then we're going to figure out the stabilities and the eigenvectors.

And then we can really understand the entire dynamics of the system without solving the full thing, without using a computer. But just by figuring out where the fixed points are, and then the dynamics around there. Then you can basically understand all the dynamics of the system. But you have to make sure that you develop intuition of how systems behave near their fixed points.

Now this is a case where both of the eigenvalues were real. But of course, if you have complex eigenvalues, if you do the calculations, you'll see that the eigenvalues are going to be a complex conjugates of each other. And there you get spirals. So trajectories might look like so.

So there are two qualitatively different ways that the fixed point can be stable. You can have spiraling through a state of the fixed point, or you can come in via these straight lines. Of course, the specific trajectories in some cases, can be curved. And so if you look at a particular trajectory, sometimes even these can look a little bit spirally, so be careful. Those are the two basic ways it works.

There's a simple way of getting a sense of the dynamics of a system, as a function of the trace and the determinant. So what we have-- something here. If you go ahead and you do the calculation, what you find is that the two eigenvalues are going to be described by this. And then the basic dynamics are going to be the

following. So down here, you have a case where the eigenvalues are real, and they're of opposite sign. And in this case, is that stable or unstable?

Unstable. So in order for it to be stable, all the eigenvalues have to have negative real components. So if they have opposite sign, one of them is going to be positive. So this is all unstable down here.

Over here you have a case where the eigenvalues are real, greater than 0. So in this case, the trajectories are also coming out somehow. All right, here, this is the case where they are real, but now both less than 0. So this is again, this is like what we drew here, where everything is stable coming in. And up here is where we get the spirals.

But over here, it's the real part is less than 0. So this is where we have the spirals coming in, whereas over here, the real part is greater than 0 and we have spirals coming out. Oh, I'm sorry. Uh, that would have been useful. So this is in the trace of A and this is in the determinant.

So these are the variety of possible outcomes when you have a two dimensional system that's already a linear two dimensional system. And this is indeed consistent with it, to have a stable fixed point, you need to have the trace less than 0, and the determinant greater than 0. So it's in this quadrant.

Now it's-- finally it's useful to-- so let's just for now, stick with the linear system. And just imagine a case where we have in this A-- so remember it's A, B, C, and D. So we saw one way in which this thing could be stable. Was that if b and c were both equal to 0, so there's no cross interaction, but a and d were both less than 0, then it's all trivially stable. Now the question is, if something that looks like this, a is equal to minus 2 and d is equal to plus 1.

Questions is, could such a system be stable? So I'm seeing some nods. And on the face of this, you say, oh, that's a little bit surprising. Because d being plus 1, what that's saying is that y on its own is unstable. So if you start out with no x and no y, and you add a little bit of y, y starts growing exponentially.

So y on its own is unstable, but x is stable its own. And the trace being less than 0-- that's saying that there's some sense in which if y is unstable on its own, then x has to somehow be more stable than y is unstable. Because the sum of those things still has to be negative. Now that was necessary, but not sufficient, in order to have stability. Because we also need to have a condition on the determinant.

Now so the trace of a here-- that's minus 1. That's less than 0. OK, that's great. Now the determinant-- now it's going to be a times d. That's minus 2.

But then we also have to say minus b times c, right? And this thing has to be greater than 0. So what you see here is that b and c-- they have to have opposite signs order for this thing to work, in order for the origin to be stable. And the product has to somehow be strong. Now this makes sense, because of course, if b and c, or even for matter, just one of them were 0, then it would be impossible to get the stability.

But for example, if we have some situation where we have some x that's inhibiting itself-- that's a here-- but then y is activating itself. So y here is somehow on its own, unstable. Then what you need is you need maybe something that looks like this. Some cross activation and or repression. And what's interesting is that it doesn't matter which of the two, b or c is negative.

You can get actually, the origin to be stable in either way. But in this situation, you need the direction of the regulation to be in opposite directions. And if you'd like, you could then play with-- you could think for example about the directions of these trajectories around the origin, and so forth. But it's useful, I think, to play with the simplest toy systems that you can imagine, just so that you can get a sense of what are the basic ingredients that you need in order to get stability in something like this.

Because once you start doing the whole linearization around the fixed point, and then calculating traces and determinants, you're not going to have it. You're going to lose all your intuition about that about things at that stage. So it's useful to make sure that you nail things down in this context. We are out of time, so I'll let you go.