

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high-quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: So today, what we're going to do is we want to talk about this idea of autoregulation, so what happens when a gene regulates its own expression. And I think that this is a really nice topic because it really encapsulates several of the big themes that we're going to be seeing throughout the semester.

So first there's this idea of a network motif. So this motif, it's the simplest possible motif where a gene regulates itself. It occurs more frequently than you would expect by chance. We'll explain what we mean by that. And then possibly an evolutionary explanation is that this thing may be advantageous in one way or another.

Now, there are two basic kinds of negative autoregulation, or sorry, two basic kinds of autoregulation. There's negative and positive. So negative means that that protein represses its own expression. Positive autoregulation would mean that it is activating or upregulating its own expression. So down here is positive autoregulation. And these things have very different possible purposes.

All right, we'll see that there are basically two things that negative autoregulation may do for us. First is that it speeds the on time relative to simple no regulation. But the other thing is kind of an interesting thing, which is that the protein concentration is somehow robust. And robust means that it's comparatively non-sensitive or insensitive to something. And particularly, this means protein concentration is robust to say fluctuations in the production rate.

In various contexts, we're going to see how robustness is an important goal, that biology might have, a particular cell might have. And that you would like this protein concentration to be robust to a wide range of things. But in this case, anytime you talk about robust, you have to say all right, it's robust. What is it that's robust against

what? In this case, its protein concentration is robust, the fluctuations in the production rate.

Now, positive autoregulation, on the other hand, in some ways makes these two things worse. We'll discuss what we mean by that. But it perhaps performs some other functions. In particular, it allows for the possibility of bistability, and hence memory. And we'll explain how memory appears in this process.

But before we get into autoregulation, I just want to do a quick review of this thing about getting intuition or binding of activators, for example, to the promoter. But also how we can use this thing about sequestration in order to generate what we call this ultra-sensitive response.

Now, when we talk about this, we're often thinking about a question that we'd like the rate of expression of some gene to be sensitive or be ultra-sensitive to the concentration of something. So if you have x activating y , then what you might like, in many cases, is for there to be an essentially digital response, where if you think about the production rate of y as a function of this input x , you might like it to just be very sharp, to be essentially 0. Quickly to come up, and then saturate at some level, β maybe.

Now, the question is how is it that you can do that? One thing that we discussed that could get you something kind of like this. Does anybody remember what the base solution might be?

AUDIENCE: Cooperative regulation.

PROFESSOR: Right, cooperative regulation. So if it's the case that you have, for example, a tetramer of x that is somehow activating y , then you can get something that looks, maybe not this sharp, but much sharper than what you get, which is the simple Michaelis-Menten-type looking curve if you just had a single monomer of x activating y . So one solution is indeed cooperativity.

But I think that there's another interesting one that is this-- second solution is this idea of molecular titration. I kind of explained the basic idea at the end of the lecture

on Tuesday. But then what we want to do is try to get a sense of what the requirements might be for the bindings in order to generate this effect. So what we have is a situation where we have, again, just x is regulating y . But it's a little bit more complicated, because this protein x , in addition to binding the promoter and activating expression of y , we also have the possibility that x can bind to something else. In particular, there might be some other protein, maybe w , which combined reversibly into some complex wx .

Now, for some situations-- in particular, if we describe this by some K_w and this by some K_d we can ask in what regime will this generate an ultra-sensitive response, where as a function of the total concentration of x -- So we might want to call that x total, just be clear. We'd like it that there's very little expression of y , and that all of a sudden to get maximal, if you like, expression of y . And the question is, well, what do we need in terms of these K_wK_d 's, and so forth.

So there are going to be three different conditions that we're going to play with. First might be the relationship between this K_d , the affinity of binding of that transcription factor to the promoter, as compared to this K_w . I will write some options, but as always, you should start thinking. You don't need to watch me write.

Don't know again. So I'll give you that. Just 30 seconds to think about it. Yes.

AUDIENCE: Is K_w the dissociation constant? Like if the concentration w times concentration x over the--

PROFESSOR: Right, so it's the guy that has units of concentration. So lower in both these cases corresponds to tighter binding.

I'll let you think for 20 seconds maybe.

Do you need more time? And it's OK if you are not 100% convinced of something. It's useful to just make a guess and then we can discuss. Shall we go for it? Everybody have your tools in front of you? All right, ready. Three, two, one. So we have, I'd say, a fair amount of disagreement, actually. I'd say we got some A's, B's,

C's. I don't know. There might be a one or two D's. Why don't we go ahead and turn to a neighbor there. You should definitely be able to find a neighbor that disagrees with you. And try to convince them that you--

AUDIENCE: What's the question?

PROFESSOR: Oh, yeah, sorry. No, no, all right, so the question is, what conditions do these things need in order to have an ultra-sensitive response, where the rate of production is what?

[SIDE CONVERSATIONS]

PROFESSOR: All right, why don't we go ahead and reconvene. And maybe I'll just get a sense of what the state of your thinking is right now. All right, so once again, the question is, we want to know what the relationship between all these binding affinities concentrations has to be in order to get an ultra-sensitive response, where the function of the total amount of x-- you add x. At first, you don't get really much of any expression of y, but all of a sudden, you get a lot. And first, we want to know the relationship between K_d and K_w . Let's go ahead and vote. Ready? Three, two, one.

Oh, I'm sorry. We're still on this one. So yeah, you could ignore these. It's just that it takes me time to write. So I wanted to take advantage of your discussion. We're still on this one. Do you guys understand what I'm trying to ask? Okay, all right, ready. Three, two, one. All right, I'm going to cover this so nobody gets confused.

All right, so we have a fair pocket of C's back there, but then some other distributions. So we've got some discussion over here. So maybe somebody can offer their insight?

What did your neighbor think?

AUDIENCE: My neighbor thought that at first, there should be the binding between x and to the gene. So x should bind much more equally with w. That's why K_w should be--

PROFESSOR: All right, so you're arguing that you want K_w to be much smaller than K_d here. Because as you add x, you want initially for them all to be sequestered by molecule

w. Is that OK? So if somebody's neighbor thought it should instead be B, do you want to offer their argument? Everybody's neighbors now convinced that C is--

So I'm going to side with that. So the idea here is that what you'd really like is-- The principle of this is that K_w is perfect, and then that corresponds to going to 0. Then that means that if you plot the free x as a function of x total, then what does it start out being if K_w is 0? 0, right? And when thinking about the concentration of free x , do we have to think about or worry about how much x is down to the promoter? No, we'll say that so long as you have any reasonable concentration of x , then the one x that binds to that promoter, binds that DNA, is not going to affect the concentration of free x .

So we really just have to worry about how much is bound to w . So it's going to be 0. And this is in the limit of K_w going to 0. And so when is it that something starts happening?

AUDIENCE: x is the concentration of w .

PROFESSOR: Right, so it's when you get n . And is it the concentration of w , or is it-- maybe we can be a little bit more precise? What shall I actually write here? Should I just write w , or should I write-- w total, right. Because the free w is changing all the time. And indeed, at you approach here, the free w goes to--

AUDIENCE: 0.

PROFESSOR: 0, right. So over here, the free w is equal to w_t . But then it decreases linearly until you get here. And now, there's no free w . And it's at that stage that this free x is going to go up with a slope of what? 1, right? Right, exactly. All the x that you add just turns into free x , so this goes up with slope 1.

So this is the idea behind this mechanism. What I'm drawing is the case of perfect sequestration. And in this case, what happens is, if you look at the production rate of y , it's a function again of x total. The production rate in here is to be equal to what? 0. Until you get to this w_t , and then it's going to come up. And how is it going

to come up? Does it immediately come all the way up to maximal? No, so what determines how rapidly that's going to happen?

AUDIENCE: The concentration of free x .

PROFESSOR: The concentration of free x . And we actually know what the concentration of free x is here now.

AUDIENCE: The K_d .

PROFESSOR: Right, so it's the K_d of the binding. So this is going to be some-- so this is again a Michaelis-Menten looking kind of curve, where the half maximal here is just going to be this w_t plus that K_d .

All right, so there's a sense that this is ultra-sensitive, because initially you don't get much of any expression, then all of a sudden you start getting significant amounts.

Are there any questions about the basic-- the intuition here? We still have come back and think a little bit more carefully, kind of quantitatively, about what this mechanism means for the other comparisons? But at least, based on this idea of the perfect situation, where the sequestration of w with x is complete, then this is the idea of what's going on. Is it clear right now? Yeah.

AUDIENCE: What's the direction of K_w in this situation? Is that--

PROFESSOR: In what I'm drawing here?

AUDIENCE: What you drew there.

PROFESSOR: This one? Oh, OK, so the definition of k_w is going to be that k_w -- and this is a definition. It has units of concentration, so that means we have to put up-- this a conservation of w , a concentration of x , the concentration of wx .

AUDIENCE: Oh, [INAUDIBLE].

PROFESSOR: Any other questions about this? Yeah.

AUDIENCE: Why does the flow go the free x 1, because isn't the free x now binding with the gene?

PROFESSOR: OK, right, so this is what we were saying is that there's just one copy of this gene. So given that there might be-- where in this case, this might be 1,000 of these proteins, and then it's just not a significant-- yeah, in the back?

AUDIENCE: So whenever you have a k , do have always concentration?

PROFESSOR: OK, so in this class whenever possible, at least in the lectures, we'll always try stick with the dissociation constant, which is the guy that has units of concentration. You can also talk about the association constant, and then they also use k for that. Horribly confusing, right? So whenever possible, I'll be referring to the dissociation constant.

AUDIENCE: Is that written a lower case k or a big K ?

PROFESSOR: Oh, that's right. No, it's a good thing that my K 's are so clearly big K 's that there's no possible source for confusion there. Yes. But if you're ever confused about which K or whatnot I'm referring to, please just ask. All right, any other questions about the basic mechanism? Yeah.

AUDIENCE: Just a correct one. So if we take the other limit for K_d going to 0, would it be the same-- so free x versus x total, would it be the same slope, but starting from 0 that should be left?

PROFESSOR: Oh, OK. The other limit of K_d going to 0. Yeah, this is an interesting point. Right, so I think this plot is independent of K_d , because binding to that promoter doesn't affect the free concentration of x anyways, right? It does affect this, though, because if K_d gets smaller and smaller, then this curve actually gets more and more steep. Does that answer your question?

AUDIENCE: Isn't that more sensitive if the curve is deeper?

PROFESSOR: OK, right, OK. Yes, OK, that's a good point. But now you're taking two limits of things going to 0, so we have to be a little bit careful, because it still maybe depends

on the relationship between those two as they go to 0. No, that's fair. But the idea is that if you're in this other limit, then you end up not having significant sequestration in the sense that as you start adding x , you start getting expression of that protein early on. And so then the whole mechanism doesn't work. So it's true that in principle that thing is steep, but it was never inhibited to begin with. Because the moment you start adding x , you start getting expression from the gene. Other questions about that?

Let's try out this next one. Now, this is a question of K -- so this is the binding affinity of x to that sequesterer as compared to the total amount of that w . All right, so what's going to be the requirement there? So we'll go ahead and give you 30 seconds to--

AUDIENCE: Excuse me, but don't we already have what we want? What's the more specific question?

PROFESSOR: Right, OK, so here, this is the case where w went to zero. Sorry, Kw went to 0. Right? And indeed, this works. And the question is, in general-- Kw is not actually going to be equal to 0. And it might be relevant for this question still. Yeah, because in some ways, this is the idealized version, and any real thing is going to have just some numbers for these things. The question is, will those numbers work.

Do you need more time? Let's go ahead and make our best guess. In order to get ultra-sensitivity response of y to the concentration of x , what is that you need here? All right, ready. Three, two, one. So we got some A's, some B's, some D's. C is not very popular. OK, so it seems like-- Well, why don't we just go ahead and spend 30 seconds? Turn your neighbor. You should be able to find someone who disagrees with you.

[SIDE CONVERSATIONS]

PROFESSOR: OK, why don't we go ahead and reconvene, and you can give your argument to the group? All right, does somebody want to volunteer an explanation? Yeah.

AUDIENCE: I haven't had time to test this explanation.

PROFESSOR: That's just fine. Well, your neighbor will appreciate it.

AUDIENCE: So I was thinking that the time scale of that axis really on the second graph is set by how quickly that slope, the curving slope, rises?

PROFESSOR: OK. so times scale I'm a little bit worried about.

AUDIENCE: Sorry, concentration scale.

PROFESSOR: OK, all right.

AUDIENCE: It's set by how quickly the second curve rises, so once you go above 0. That's the only feature on that graph. And that should be in a Michaelis-Menten curve that's determined by K.

PROFESSOR: K, which K?

AUDIENCE: Oh, the K of the Kd.

PROFESSOR: Yes.

AUDIENCE: Oh, no.

PROFESSOR: OK, right, so I think I like everything you're saying, although I think you're about to answer the next one, although the next one, I think, is in some ways the hardest one, so you're ready.

AUDIENCE: Right, which we determined should be a lot more than Kw.

PROFESSOR: Right, so we know that Kd should be a lot a lot more than Kw.

AUDIENCE: And I think that this should be a decent amount less than-- Sorry, so I think that Kd should be a decent amount smaller than w2.

PROFESSOR: All right, so you think Kd should be a decent amount smaller than w2. OK, so you really are answering the next question.

AUDIENCE: But no, I think this is important.

PROFESSOR: Oh, no, I agree it's important. But--

AUDIENCE: Basically speaking, if you let w_t become too small, then it compares into everything else, all the relevant quantities, then this feature, which is really determining the over-sensitivity strengths have 0 side. That's what I wanted to say.

PROFESSOR: OK, yeah. I very much like that explanation. And since we have it on tape, we can now play it again in a few minutes and then-- I think you're making a very nice argument, and actually, we'll go ahead and even--

Does anybody want to argue against. Yeah. So what he's arguing is that what you-- So he's actually arguing that it's this over here, which is that what sets the scale that this thing increases by is the K_d . And you want this thing to come up quickly relative to this other scale, which is the total w . So if total w is just too small here, then it's not like you get any ultra-sensitive response, because you want it to be low, low low, and then come up kind of quickly. So there's some sense that you want this scale to be of the order of this scale or maybe even a bit shorter. Is that-- and I agree that this actually one that generates a lot of argument an discussion in general, because I think that reasonable people can disagree. I think this is actually-- I will side with you on this. But what about this, because you haven't said anything about KW .

AUDIENCE I think that combining one and three.

PROFESSOR: OK, OK, fine.

AUDIENCE: That's why I was trying to answer--

PROFESSOR: OK, no. OK, fair. But it's just too many logical leaps. I think it's true. And there's a reason I ordered the questions in this way, so that is so this one you could-- Otherwise, I agree that if you combine these things, you do get-- and which one do you get?

AUDIENCE: I think A.

PROFESSOR: Is that what other people got? OK, I agree that actually you can get to A here from the combination of C and B here. But it's a little bit crazy. Is there a more direct explanation we can get? Yeah, go ahead.

AUDIENCE: Well, if you don't have a lot of w, then your binding is too sensitive in a way. You'll add a little bit of x, and then the binding will be all used up, and you'll be making y immediately. You want a lot of w to be able to soak up a sizeable amount--

PROFESSOR: Yeah, and then maybe in the back.

AUDIENCE: So originally, I agreed, but that just makes you switch, like Evy said, you can sequester less, but that still gives you a sense of [INAUDIBLE]. You're more sensitive to fluctuations if you have a lower sequestration capacity, but it doesn't really change-- You still have that switch--

PROFESSOR: Sure, OK. So I think that this is actually [INAUDIBLE], because there are two ways of having more of this wt helps. One is that it pushes this boundary further to the right. But the other is that it makes it a better sequestering agent because there are just more of these w's to bind. And if you go and you do you do the calculations, and I encourage you to do it, what you actually find is that the concentration of free x is given approximately by $x_{\text{total}} \text{ divided by } wt \text{ over } K_{\text{sub } w}$. So what you see is that the free concentration of x-- and this is in the region. This is for x_{total} less than and maybe significantly less than w total.

But the idea is that when you're in this sequestering regime, what we see is that the concentration of x is going to increase linearly with the total x. But it's going to be sort of suppressed by this amount here, where this ratio is much larger than 1. You want to be much more than 1, $wt \text{ over } Kw$. And this is telling you how good of a sequesterer this guy, this w is, because you want to bind tightly and you need to have kind of enough to keep the free concentration of x from growing.

Yes.

AUDIENCE: So that's what I was thinking as well, but then I was wondering if there's too much of w, then wouldn't x never bind completely, but always get sequestered?

PROFESSOR: So this is a very important point, which is that it needs to be possible for the cell to make enough x to overcome the w . If there's so much of the sequestering protein that you cannot even get beyond that, then you're never-- It's true that you might have a nice ultra-sensitive switch out there, but you just can't get there. So yeah, this is certainly relevant as well.

All right, incidentally, if you plot free x as a function of x_T on a log-log scale, the question is what is it going to look like? Now we know that on a linear scale, it looks something like that. So if it's on a log-log scale, question is-- And log-log is nice because you can really get the full dynamic range of what's going on. If you plot x and this is free x , as a function of x total in this regime for large X and x total, what is it going to end up looking like?

AUDIENCE: It's a straight line.

PROFESSOR: It's a straight line. And eventually, it's going to be straight line with slope 1, actually, because they're just linearly related to each other. Now, down in the low regime, what is it going to look like well below this region of w ?

AUDIENCE: You should also get a straight line with 1 slope.

PROFESSOR: With what slope?

AUDIENCE: With 1 over [INAUDIBLE].

PROFESSOR: OK, so why does everybody agree? No, we have a disagreeable-- Right, so what's the--

AUDIENCE: It's going to be also slope 1, but just lower.

PROFESSOR: OK, right. So it's actually also slope 1, but it has a lower-- This is dangerous. This is why I brought this up, because it's really easy to look at these things and think that-- Yeah, because they're still linearly related to each other. So when you take the logs, it just affects the level. So you still get the same slope, but it's going to be down here somewhere. And then in this regime, you get this ultra-- this thing where it goes up

suddenly.

One thing that I strongly encourage everyone to do is, in these sorts of problems, it's wonderful to spend time-- Oh, sorry. This is \log . Free $x \log$ of xT . It's really very valuable to plot things in multiple different ways, by hand or by the computer or both or whatnot, just to make sure that you're keeping track of what's going on. Because often what you see and what you think is very different depending on what you plot. And you'd like to be able to see your problem from as many different angles as possible.

All right, I think that we've probably spent about as much time on this as we ought to. But are there any other questions about how this is going to come about? Yes

AUDIENCE: Are there any negative autoregulation that can use sequestration to get a switch-like [INAUDIBLE]?

PROFESSOR: OK, that's an interesting question, although we have to be careful about-- If you really want it to be more switch-like, you'd probably use positive autoregulation. And I'm not aware of a case where this has been combined, although-- It's likely there are some. I just don't know them.

I'm going to switch gears to this autoregulation, which is something, of course, that you guys just read about. And it looked like your understanding of it was solid. But we want to move through these different ideas. First, this idea of a network motif. And this is just the simplest example of a network motif. And it's so simple, we often don't even call it a network motif. But the idea here is that we have some network, and it has maybe N nodes and E edges. And the example that they give, that Uri gives in his book, N was 420, and E was 520. ,

And there's a basic question that if you have a network with N nodes and then you have directed edges-- these are edges that have an arrow pointing on one end. Now, in that case, and if you allow self-directed , edges how many possible edges are there? Does anybody remember what this looks like?

AUDIENCE: More than 20 possible, right?

PROFESSOR: What's that?

AUDIENCE: There'd be more than 20 possible.

PROFESSOR: Right, not even in terms of the actual number. Just in terms of N , for example.

AUDIENCE: Total or just self-directed?

PROFESSOR: Total self-directed. Or Sorry, total directed edges, total number of possible directed edges, if we included self edges, I guess. N squared. And you can think about this in multiple ways. One is just that, well, you can start at any of the N nodes. And you can end at any of the N nodes, and that gives you N squared. Right?

But there's another way that you could-- For instance, this is Emacs. It's NE N squared. You could also think about this, if you like, as just-- say, well, as they point out, there's sort of-- $1/2 N$ times N minus 1 is the total number of pairs in the network. And then for each of the pairs, you can have an edge pointing either direction, so that gives you a 2. Plus, you can have N self edges, right? Now, of course, these are just different ways of counting N squared. But it's useful to just think about it in different ways to make sure that you're comfortable with the combinatorics of it. In particular, because next week or the week after that, we're going to be talking about network motifs in more detail, in particular the feed-forward loop. And then we really have to keep track of these sorts of combinatorics better.

The way that Uri thinks about this is he says, all right, well, we're going to invoke this null model of a network, this Erdos-Renyi network, where we just simply say we're going to assign the E edges randomly across the Emacs possible. So of all possible edges, we're going to place the edges randomly and then I generate some random network. So this is basically what we typically call a random network. And that's going to allow us to define some null model that if everything were random, we can ask, for example, how many self edges do expect to get?

Well, one way to construct this sort Erdos-Renyi network-- Yeah.

AUDIENCE: But how do you know if you have other kind of constraints in the system, how do you know that transcription might require in some cases autoregulation--

The answer that was give that was a good answer was that evolution is the only thing that can-- if you find a network motif that has to do with evolution--

PROFESSOR: So that is the argument that Uri makes, and you're maybe saying maybe that's not a good argument. And what Uri is saying as well, if you see these networks more frequently than you would expect by chance-- and of course, you can define what you mean by chance-- then you can say oh, maybe it was select for, it was evolved. And I think that, in the case of this autoregulation case, I think the results are not very sensitive. But I think that this question of what the right null model is, is a real issue, especially when you're talking about some of these other networks.

And what we'll see for the feed-forward loop is that you have to decide-- well, one thing we're going to see is an Erdos-Renyi random network is very much not an accurate description of real transcription networks. So then you could say, well, that's not a good null model to use. And so we'll definitely spend some more time thinking about this.

In the context of the Erdos-Renyi network, though, one way that you can generate it is that for each of the Emacs possible, each of these total number of edges, there is some probability that you're going to actually place a real edge there. And that probability is just E over N squared. E is the number of edges. N squared is the number of possible edges. So if you just create a random network in that way, then this is a manifestation of a random network that has at least the basic properties, the same number of edges as our network.

So from this, you say, well, how many self edges would you expect in this world? And you'd say well, in that case-- There are two ways of thinking about this. So you can either say, we're going to take, for each of the N edges, there's one possible self-directed arrow. And for each of those cases, we can just multiply this by P . And

this gives us E over N .

You could also think about it as-- There are multiple ways once again, of doing the counting. In an Erdos-Renyi network, you can say, all right, you would expect to get roughly E over N autoregulatory loops.

And this is of order 1. So this is 1.2, in the case of the network that Uri analyzes in his book and his paper. Whereas, how many were actually observing in the network that he studied?

AUDIENCE: 40.

PROFESSOR: There were 40, right? So in the observed transcriptional network-- and this is in *E. coli*-- he found that there were 40. And the basic statement here is that 40 is just much larger than 1.2.

And you can quantify this a little bit better, because, it's really you would expect 1.2 plus or minus the square root of this, in a random network like this. But that is, you'd expect 0, 1, 2, maybe 3. So 40 is definitely not what you would expect based on an Erdos-Renyi network. So this is the sense in which it's a network motif. It's that the observed network just doesn't look like a random network in this very particular sense.

And of these 40, does anybody remember kind the distribution between negative autoregulation and positive autoregulation?

AUDIENCE: It was 30-10 or something.

PROFESSOR: Yeah, I think it was like 34 and 6 was my recollection. I didn't write this down. So most of these guys have a form of x inhibiting x . But some had the form of x activating x . So this was something like 34 and 6. What you would say then is that negative autoregulation is a very strong network motif, whereas positive autoregulation is a weaker network motif, but still something that occurs perhaps more than you would expect by chance. Are there any questions about that where that argument came from, other concerns about it?

All right, then from this, then Uri says, OK, well, maybe these things evolved for a reason. And so what we'll do in the next half hour just argue or discuss what possible evolutionary advantages such a social network motif might have. Yeah.

AUDIENCE: Can you really propose adversarial explanations like that?

PROFESSOR: Oh, you can propose whatever you want.

AUDIENCE: I know, but it doesn't have much value. It's unquantifiable.

PROFESSOR: Yes. No, I'd say this is a major issue in a lot of evolutionary arguments. And I would say that the purpose of ideas and hypotheses is to get you to go and make new measurements. And so right now, we're the stage, OK, well, maybe these things are occurring more frequently than you would expect by chance. So now, we can just sit down and think, oh, what advantage might it give. And then we can go and try to experimentally ask whether those advantages are at least manifested in real systems. It doesn't prove that that's why they evolved, but it makes you more comfortable with the argument. feel

Ultimately, we assign some-- we have some agent probability somewhere in our brain. And the more evidence that we can accumulate that's consistent with these ideas, the more likely that we think it is. But in general, you don't prove things in this sort of evolutionary space the way you prove things in many other fields. Yeah.

AUDIENCE: I feel like it's hard to call this an argument. It feels more like just an observation.

PROFESSOR: Which thing is an argument versus--

AUDIENCE: I guess the thing is it should be evolutionarily advantageous, that's an argument, but essentially, the whole thing is an observation, and then there's a little bit of an argument in the end.

PROFESSOR: Yeah, I will let each person decide what fraction and observation. Yeah, I don't feel especially strongly about it. My guess is that it did evolve it because it provides some useful function. And therefore, I think it's valuable to explore what those useful

functions might be. But for example, it's very hard to know which of these explanation-- this thing about increasing the response time, or sorry, increasing the response rate as compared to increasing robustness, how do you decide which one's more important? Then I think, once again, reasonable people can disagree about these things, yeah.

So first negative autoregulation, because this is the one that is the stronger network motif. I think that the book does a nice explanation of why it decreases the response time. OK, we can just ask. OK, response time-- and this is for a negative autoregulation. Response time goes down. And is this for turning on, off, both, or maybe neither, or E, don't know.

And I'll give you just 10 seconds to think about this. It's nice if you just remember it, but it's also maybe even better if you can figure it out. Because in a week, you're probably not going to just have it memorized, but you should be able to think through the logic of it and understand why this is going to be what it is.

All right, so the question is, negative autoregulation, maybe it does something. Maybe it decreases the response time. But does it decrease the response time for turning a gene on, for turning it off, for both, neither, or don't know

AUDIENCE: When you say turning it off, what exactly is the process you're imagining.

PROFESSOR: I'm imagining a process where the expression turns off immediately. So there's a signal that just stops--

AUDIENCE: what transcription can go ahead.

PROFESSOR: Right, so then it's just I chop up all the polymerases, and no more expression. But so a signal comes and tells the polymerases to stop making. Yeah.

All right, so do you need more time? No. Ready, three, two, one. All right, so we actually are all over the place on this. OK, turn to your neighbor. And you should be able to explain one way or the other why thi-- what is going on.

[SIDE CONVERSATIONS]

Let's go ahead and reconvene. I just want to remind everybody that when I say simple regulation, there's no autoregulation. It's just responding to a signal. That for a stable protein, the time to get to say, for example, half saturating concentration here is defined by the cell generation time. And that's true for turning on and for turning off. And what was the strategy that you could use if you wanted to decrease the response time in this situation?

AUDIENCE: Increase the degradation rate.

PROFESSOR: Right, so you could increase the degradation rate. And does that the on, off, or both?

AUDIENCE: Both.

PROFESSOR: Both. But there's a cost, which was what?

AUDIENCE: You have to make protein.

PROFESSOR: Right, you have to make a bunch of protein, and then you're just going to chop it up right after you make it. There is a reasonable-- there is a way to make things faster, but it has a significant cost.

The question is, if you have negative autoregulation-- so in this case, you have x that is repressing itself-- what is it that it's going to do? Is it going to affect the on time, the off time, or both. Let's just see where we are. Ready, three, two, one.

OK, so it's interesting. We're moving towards C, it seems. OK, so can somebody give me an explanation for C. Did we read the chapter?

AUDIENCE: Well, the chapter doesn't discuss the effect of negative autoregulation and turning off. I don't think it does.

PROFESSOR: Wow, it's a good thing we're doing that here then. All right. So first of all, can somebody give the explanation. Does T go up, down, or sideways.

AUDIENCE: Up.

PROFESSOR: So T on-- It's the time that goes down. I always get this confused. So time is the one that goes down, so the rate goes up. Negative autoregulation is faster turning on, we decided. Right? And does somebody want to give the explanation for why this is?

AUDIENCE: Well, your equilibrium level is lower.

PROFESSOR: Yeah, right. Yeah, exactly. Yes, this is actually surprisingly difficult to explain even though it's not a deep concept. But the idea is that you start out expressing a lot, so that if you had kept on expressing that high level, you would have done some exponential-- It would have take cell generation time from way up here. But instead, what happens is that you shoot on up. But then, once you get up here, you repress expression. So then you get an effective thing, where the time it takes you get half of your equilibrium, that goes down. So TI in here is shorter than here. Yes.

AUDIENCE: So in negative autoregulation, for decreasing what the book calls beta, to have the same steady state?

PROFESSOR: That's right. The initial beta, that rate, that maximal rate of expression, that goes up in a case of negative autoregulation. But then you start repressing expression once your concentration of x here gets to some reasonable level.

So now we're just talking about production rate of x. And that's as a function of x. And of this logic approximation is when it just is maximal until it gets to some K and then is completely repressed. So real versions will be much smoother, but this is just useful to start getting the intuition.

And the idea is that you shoot up to this K, and then you stop expressing. In this limit, actually, it's not even-- it's like a kink here. Because it just shoots up and then it turns around. But any real system will be smoother. Yes, question.

AUDIENCE: So if you get to a certain equilibrium level, then in autoregulation, you would need a stronger promoter.

PROFESSOR: Yes, you want a stronger promoter, because you really want to have high expression initially and then later repress that. So negative autoregulation allows you to speed up turning on, so T on goes down.

AUDIENCE: Without increasing the promoter, which is a good thing, because someone would die if you increase the promoter.

PROFESSOR: This is a very important point, which I was about to get to, which is that this is something that was done-- We could have done that without negative autoregulation by increasing the degradation rate. So the question the that we're bringing up here is, is there that same cost that we were referring to before of this futile expression of protein at equilibrium.

AUDIENCE: No.

PROFESSOR: So it's actually not. In a cell, you start out expressing a lot, but then later, you actually bring down your rate of expression. And in any case, there's no degradation. in this. The only effective degradation is due to the dilution, the growth of the cell. So if you have the same concentration, then actually, you don't make any more protein than you did here, because you have the same concentration at equilibrium.

So this is neat because this speeds up the response when you're turning on, without the associated cost of making that protein then degrading it. Any questions about that statement?

So now what about off? Is the off time the same as the on time here?

So what sets how fast--

AUDIENCE: Should the off time be slower because you have lots of degradation.

PROFESSOR: Right, and in principle, is there any active degradation that we've invoked on this?

AUDIENCE: I think not.

PROFESSOR: Of course, we could have both negative autoregulation and active degradation. But in principle right now, you can have the negative autoregulation without any active degradation. In that case, how long does it take for that concentration to go away when you stop expressing?

AUDIENCE: The cell degeneration time.

PROFESSOR: The cell degeneration time. So this thing actually looks the exact same as this.

So these guides are the same, whereas this one is faster than that one. Because the idea is that the best that-- unless you're inactively degraded, all you can do is you can shut off expression. But then if you turn off expression on the negative autoregulation, it's the exact same thing as turning off expression in the absence of the neg-- in either case, you just stop making protein. So the concentration just goes down because it's being diluted away during cell growth. So this is saying that response time was down only when turning off in the case of negative autoregulation.

Are there any questions about that idea? Yes.

AUDIENCE: With the negative autoregulation, in order to reach the same protein levels, you'd need much greater production rates, correct?

PROFESSOR: Yeah, so the idea is that this beta might be-- so this is the beta of negative autoregulation. It could be much larger than the beta of simple regulation in order to get to the same equilibrium.

OK, so what about this idea of robustness? Well, this is production rate and then degradation rate. So this is an αx . And so my question here is, I told you that robustness-- something is robust-- Yeah, question.

AUDIENCE: My question is in this case, you're saying that the off means signal disappears, right?

PROFESSOR: OK, T off is this idea. It's the $T_{1/2}$. So this is the time that it takes for the protein concentration to reach half-- to go from halfway the distance from where you were

to where you're going to end up.

AUDIENCE: But what if the signal not disappear, but to half of the original signal?

PROFESSOR: So the signal could do a range of different things. And it could be that the signal just changes so that instead of going down to 0, you go down to some other value. Is that what you're imagining?

AUDIENCE: Yes.

PROFESSOR: In that case, you still go exponentially to this new value, so actually the T-- the response time there is still actually the cell generation time. So it doesn't matter, in the absence of any these, for example, autoregulation. The time, the characteristic timescale, is always the cell generation time if it's a stable protein. It doesn't matter whether you're going up, down, or all the way to 0 or not.

So the question here is-- OK, x equilibrium is robust to what? And this is to small changes in what? It's going to be A , α .

So this is going to be our first example of an advanced use of our cards. So the way that it works is that you can choose more than one. OK, now, this requires some manual dexterity. So what you have to do is if you think that the answer is more than one of these things, then what you have to do is show me more than one card. These cards are amazing, right? You can do so many different combinations. I'll give you 20 seconds to think about it.

AUDIENCE: So what's e ? What do those things mean?

PROFESSOR: OK, the question is, are the equilibrium concentration of protein x is robust means it does not change in response to small changes in what quantities? So if I change the degradation rate, does it change equilibrium. If I change the β . And I'm asking about this case here, perfect negative autoregulation, just so we can try to establish our intuition here. K is this repression threshold. None means that it's not robust to any of these things. DK always means "don't know."

I'll give you an extra 30 seconds. This might be--

So this one's the production rate. This one's the degradation rate. This figure might be useful to you.

AUDIENCE: Can you define K again?

PROFESSOR: Yes, so K is the concentration of the protein x at which this super effective repression kicks in. So we're assuming perfect negative autoregulation. Beta is the rate of expression for low concentrations. The moment you get to concentration K, you get perfect repression and no more expression.

Do you need more time? Question.

AUDIENCE: By saying that x equilibrium is robust, so you mean that when you change these parameters, x equilibrium stays exactly the same, or will x equilibrium--

PROFESSOR: For now, what we'll mean right now is that a small change in this parameter leads to no change in x equilibrium. Now, for any real example, what we'll typically mean is 's going to be some sort of sensitivity analysis. For example, where you'll say oh, a 1% change in a parameter leads to a less than 1% change, for example. But in this case, there's going to be no change, I'll tell you, just so we can get the intuition clear here. All right, do you need more time? Let's go ahead and vote. Remember, you can vote for more than one thing if you like. Ready, three, two, one.

All right, some people are using our more than one. And of course, I can give you a hint. The reason that I'm letting you vote more than once is because more than one thing is going to be-- All right, so the majority of the group has got this, but not everyone. So let's discuss. Can somebody give an explanation for why both alpha and beta are going to work here?

AUDIENCE: So the equilibrium is basically when degradation involves production.

PROFESSOR: I want to make sure I'm okay. The equilibrium is when the production rate is equal to the degradation rate. So this is a very important thing to make sure we're on top of. And in this case, we have very sharp-- this production. So then what happens?

AUDIENCE: Well, [INAUDIBLE] is the intersection of--

PROFESSOR: Right, so in this case, what is the equilibrium concentration?

AUDIENCE: K.

PROFESSOR: It's equal to K. Now, I strongly encourage you, whenever possible, to draw things out. Because this is a problem that when you have the drawing. It's reasonable to do. And if you don't have the drawing, you're going to get yourself tied up into weird knots. And indeed, we can see that if we change alpha, what happens in this spot? Right, it changes the slope. And you can see that if we change the slope by small amounts, we get no change where this crossing point is. And even for a real system, if it came around, you'd see that it's going to end up being a less than proportional change in the equilibrium.

And what about what about beta? That just raises and lowers this. And again, that doesn't change the equilibrium. Of course, if we changed K, then we get a 1 to 1 change. So a 10% change in K leads to a 10% change in the equilibrium concentration of x.

So this is the sense in which the equilibrium concentration in negative autoregulation is robust to changes in both-- in the book, they say oh, the production rate, but it's actually also in principle the degradation rate over some range. And this could be useful, because there are lots of things that are going to affect the production rate of a protein, and also the degradation rate, for that matter. The division rate, it changes it.

Whereas it may be that K is subjected to less severe changes, because that's determined by, for example, in the kinetics of binding of this protein to this promoter. And that is perhaps less subject to changes. It can still change depending upon the pH and so forth of the interior of the cell. But at least it's probably not subject to the big changes that alpha and beta are going to be a good experience.

So the argument that Uri makes for why it is we see so much negative

autoregulation in the cell is because it both increases the rate that the cell can respond to changes, in the on direction, at least, but also that it makes the concentration of protein more robust to changes in several of the parameters that govern the equilibrium. Concentration And once again, you could argue about which one of these is more important, but I think they're both likely playing a significant role in different cases.

I'm going to want to move on, but I will tell you that only over some range of these -- alpha, beta, K-- will this thing be robust. So for example, if this comes up too high, we're going to lose this phenomenon of robustness. So I expect you to be able to tell me in some later date the conditions in which that might happen.

And I'm available for the next half hour after class, so if you do not know what I'm talking about right there, please hang out with me after, and I'll tell you the solution to that question on the exam. OK? All right.

But I do want to talk about positive autoregulation, because this is another interesting beast. So if negative autoregulation has those nice properties, then you can imagine that positive autoregulation will have some drawbacks in the same kind of ways. But it leads to some other very interesting, just qualitative features.

Positive autoregulation. So we have some x that is activating itself. And often we think about cases where it's activating its own expression in a cooperative fashion. In particular, we might assume that x dot is equal to, for example, some β_0 plus some β_1 of some cooperative thing here where N might be 2 3 4 and then again minus αx . Right? Now, if you just look at this, you might think oh, I don't know what this is going to do and so forth. But you've got to draw things out. Once you draw it, then you'll see that it's pretty straightforward.

So again, this is the production and the degradation rates. So that's production. And degradation, for example, might look like this. So this is the production. This is the degradation. So that's the αx term.

One question would be, how many fixed points does this system have? So a fixed

point means that if you started right there, and in the absence of any noise, you would stay right there. So it's clearly both stable and unstable at these points.

Can you read that? I'll give you 15 seconds to count them.

Ready, three, two, one. All right, it seems like we have pretty good agreement.

There are indeed 3 fixed points. Once again, the fixed point is where these curves cross. So we have one right here, one here, and one here,

Now, how many are stable? We're going to do this verbally. Ready, three, two, one.

AUDIENCE: 2.

PROFESSOR: 2. Let's try that again. Ready, three, two, one.

AUDIENCE: 2.

PROFESSOR: 2. Yeah, you get so used to the card, it's hard to speak. So they're the ones on the ends of the stable ones. And you see here that around this point, the production rate over here is more than the degradation rate. That means that if you leave that fixed point, you're going to get pushed away.

So it's very nice to draw these little arrows here to make one happy. So this thing here is stable, unstable, and again stable.

Now, the reason we call this bistability is because there are 2 stable fixed points. This is important because this phenomenon is the basic dynamical system's origin of memory. Now it's, not obvious how memory comes from this.

So memory is a generalist idea that the g-network or the cell can retain a memory of its past state. And we're going to see examples of this over the next few weeks. But just to be clear, if, for example, we imagine a situation where the alpha changes. And it could be division rate, for example, high-food, low-food environments.

What we do is we can plot-- Often, you can plot, for example, the equilibrium, but that's a little bit trickier. So I'm just going to plot the production rate as a function of

alpha. Now, the question is, if we change alpha, what's going to happen? Now, for a fixed alpha, you can see already that there are two different production rates that are stable in this case.

But what happens if we increase alpha? So we increase the growth rate so it goes like this. Can that change the number of fixed points?

And indeed, what we can see is that as this line gets steeper here, eventually you only have a single fixed point, and it's stable. And that's known as a bifurcation of the dynamics of the system. So this is for large alpha, you end up-- And just to be clear, this is beta 0 down here. And then up here is the beta 1.

So what we do is we know that beta 0 is where we get for large alpha. Now, for small alpha, do we end up getting another-- We get another bifurcation. So actually, there's only again one stable point up here at small alpha. And what we're going to get is what's known as a full bifurcation where solid lines denote stable points, stable fixed points. Dashed lines represent unstable fixed points. So stable, and the dash is unstable. There are some regions of alpha conditions where the system is bistable. But then outside of that, it's just monostable.

Can somebody explain why this thing-- why I might make the argument that this thing displays memory? Well, one of those two is fine, but any new people want to explain my thought process? No. All right, maybe you.

AUDIENCE: All right, well depending on whether we had high degradation or low degradation rates in the past, we'll be on the lower or the upper range of that if we return to normal.

PROFESSOR: That's right. So the argument here is that-- let's say that this is some normal condition. This is where you are right now, for example. Now, depending upon whether you're sitting here or here, that's perhaps giving some information about the past state of the cell. Because if you were here, that means oh, maybe in the past you were out at high degradation rates, whereas if you're here, maybe were at low.

In particular, you could reset things. If you start here, then you can reset this memory module by coming over here. Once you get to this point here, that's the bifurcation dynamics, the full bifurcation. Then you come up here, and now you'll retain this state. In principle, until you get over here. Of course, there could be stochastic switching dynamics. We're going to talk a lot about that in the coming weeks. But at least in the limit of a low rates of stochastic switching, then this represents some sort of memory module, the simplest version of it.

I'd say that in the cell, most examples of such memory modules involve not just positive feedback of one protein activating itself, although this happens, but often through a whole network, where the one protein activates another, activates another, and then you come back. Or it could be repressing, repressing. 2 0's, is a po-- two negatives is a positive. Just like two lefts is a right.

Right, so are there any questions about the sense in which this thing can serve as a basic memory module?

And this is maybe not the most interesting example of it, because alpha is such a global parameter. But you can also get similar dynamics as a function of, for example, the galactose in the concentration of some sugar in your media.

So given that different small molecules such as food sources can act as inputs into these g-networks, you can also get these sorts of dynamics as a function of what you might call really some simple, external molecule, which is nice, because that means that you can have memory modules that are really independent of all the other memory modules that are going on in your cell. Whereas if you had a vary alpha, then this changes everything. Whereas if it's just a concentration of some sugar outside, then you can imagine that that could be very useful to retain a memory of what the cell has encountered in the past.

So today, what we've been able to do is analyze something about a possible evolutionary explanation for why autoregulation is as commonly observed as is. So negative autoregulation is the one that's observed perhaps most frequently. And that, I think, has some very clear purposes.

And this idea of the concentration being robust to other biochemical parameters I think is a big idea. We're going to see this idea of robustness crop up multiple times over the course of this semester. And I think that it's nice to think about robustness in this case, because it's perhaps the simplest example of how robustness as an approach can be useful as a way of thinking about a problem.

We're later going to be thinking about robustness in the context of perfect adaptation in chemotaxis, where bacteria try to find food. And there, I think everything's more subtle, because already the base phenomenon that is robust is a form of robustness. And so it kind of gets you mixed up. So I think that it's good to be very clear about what robustness means here, so that we can use that to think about robustness in other biological functions.

With that, have a good weekend. Good luck on the problem set, and I'll see you on Tuesday.