

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](https://ocw.mit.edu).

**PROFESSOR:** Today what we want to do is use the two papers that you read as kind of a backdrop to try to think something about the regulation of genes in response to changing environments. So there's the Mitchell paper that is talking about this idea of anticipatory regulation, whereby if the environmental changes have some typical pattern, then maybe the cells can take advantage of that and start preparing for environment number two when the cell sees environment number one.

But in other cases, it may be that the environment fluctuates in ways that are really fundamentally unpredictable. In that case, you can't use the sort of anticipatory regulation strategy, but instead, there may be a way that you can just stochastically switch between the different strategies and implement what's known as a bet hedging strategy. And this is modeled largely on the work that you just read about in the Kussell paper, *Science* maybe May 2005.

I do want to stress, however, that if you observe phenotypic heterogeneity in a clonal population, that does not necessarily mean that the cell or the population is implementing one of these bet hedging strategies. In particular, I would argue rather strongly that there are other possible evolutionary drivers for such phenotypic heterogeneity in the population.

First of all, just because you see some phenomena does not mean it is necessarily selected for. So it's possible that it's a side effect of something else. However, if you're looking for an evolutionary kind of explanation for something like this, bet hedging is not the only one. In particular, we'll talk about two other possible explanations, and they both have to do with kind of game dynamics that we might illuminate or talk about.

But I guess our game theory talk was-- that was a week and a half ago, almost two weeks ago now-- so maybe you've forgotten all the game theory that we discussed. But in particular, it could be that the phenotypic heterogeneity might be the implementation of a mixed strategy. Or possibly it could be an example of some sort of altruistic self-sacrifice. And we'll try to explain the theory behind each of these three things as well as possible biological examples of each of the three.

In this case, we may think about bet hedging as maybe an explanation for antibiotic persistence, this idea that cells can switch into these slow-growing persister states in which they're resistant to antibiotics and other stresses. Mixed strategies could be-- well, we're going to argue it could be implemented in the context of mixed sugar environments. And altruistic self-sacrifice may be the explanation behind colicin production in bacteria, so it's a toxin.

So I want to start by thinking about this thing about adaptive prediction of environmental changes by Mitchell. I think this is a very interesting paper in a number of different ways. One is, I think that it's sort of a big idea that can be explored in these simple experiments. I think that it's an exceptionally clear paper in some ways, and that they really say, oh, we're going to propose that this strategy should be characterized by these three things. And then they go and they try to show you the three things.

The figures, I think, are also very nice, in the sense that in many of the cases you could have shown the data just in the context of a table, where you said, oh, for each of these strains, this is the up regulation or so. But if you had done that, it would've been much, I think, less compelling, even though of course it's the same data.

So I think this is a neat paper, in my opinion, both from the standpoint of the ideas that are being explored, but also because it highlights some of the things that you should be thinking about when you're writing your own papers. You want to try to make the ideas as clear as possible. You want to lay the groundwork so that what you're about to show is going to be-- the reader's going to feel is a really important

thing. And then you want to take nice advantage of color and some legends that are there. So we'll kind of talk about all these issues as we go.

But before we get started, can somebody-- there's a very real sense that in this field of decision making and systems biology that a lot of this research program is kind of driven by following classic ideas from other fields. And what would be the corresponding classic idea this paper is exploring?

**AUDIENCE:** Conditioning.

**PROFESSOR:** Conditioning. Right. And whose name do we associate with conditioning typically?

**AUDIENCE:** Pavlov.

**PROFESSOR:** Pavlov. All right. So I have not ever read these studies. Pavlovian conditioning. And I think that this is just-- it's good to highlight. This is something that you might have learned in your high school psychology class. And it's, again, a big idea, but it's not the kind of thing-- we've all heard-- well, we're going to talk about this in a moment.

Many of us, I think, have heard of this. But this is an example of how you take something that you learned in high school and you make it useful to your daily life, or pseudo daily life. Because I think you'll see that there are many kind of examples of this throughout this literature, where someone takes an idea that is, in some ways, you open up a random textbook in introductory psychology, and you can just march through and try to see to what degree the ideas that were developed in the context of humans or animals, to what degree might they be relevant in the context of cell decision making?

Can somebody just say what is this Pavlovian conditioning idea? Did you guys take high school psychology? I'm sure somebody did. No. All right. Well, incidentally, I recommend everybody should take a solid introductory psychology class if you have not done so. We offer a class 900. I'm sure that it's good and interesting. Yes.

**AUDIENCE:** So there's two different events or stimuli, and in Pavlovian conditioning you give one and then always give the other one following.

**PROFESSOR:** That's right. So there's some sense that what we might call-- there are two events of some sort. One is following the other one, so it's A and then followed by B. And people have done many different examples of this, but what's the one that we typically-- what's the classic experiment that Pavlov did?

**AUDIENCE:** You ring a bell, and then you get food.

**PROFESSOR:** Yeah. Right. So you ring a bell and then you give-- and this is dogs, at least in the version I remember. So you have some dogs, you do this thing where you ring the bell and then you give them the food. And what's the response that you're supposed to get?

**AUDIENCE:** Drooling.

**PROFESSOR:** Yeah, right. So the dog is supposed to start salivating. OK. So you might think your experiments are gross, but-- right. So ring a bell, and the idea is that you can train the dog to start salivating in response to the bell, rather than, of course, if you give the food they're going to start salivating. But here you can train the dog to start salivating in response to the food.

So now, with training, this results in salivation. And did Pavlov show there was a fitness benefit associated with the salivation? Not that I'm not I'm aware of. But this is the classic story that we learned in introductory psychology. And the question is, can we apply a similar idea in the case of microbial decision making?

Now was this-- can somebody say what they think was the author's contribution to the literature in the sense of-- via a vis this? How much of what-- well, we can be maybe more concrete. In the case of E. coli, what was known-- or maybe we should describe the experiment once more and then try to figure out what they did that was new new. At least can somebody kind of summarize the basic idea with E. coli?

**AUDIENCE:** We know that they have a certain metabolic life cycle and that they exist in different environments.

**PROFESSOR:** OK. So that E. coli existent in different environments. And what environments are

you referring to?

**AUDIENCE:** Gut and outside.

**PROFESSOR:** Gut and outside. So if this is true-- then what is going to be your next statement based on this?

**AUDIENCE:** Well, going through mammalian guts there's always a certain sequence of events. At the beginning, there's some acid, and then less acid and more food, and stuff like that.

**PROFESSOR:** Right. So in the context of *E. coli*, there's this idea that part of the life cycle of *E. coli* and some other of these microbes is to live inside the gut of mammals. And we're going to focus in particular on this aspect, the part of the life cycle that is in the gut.

So let's just imagine, so *E. coli* in gut or entering gut, going through the gut, they might encounter some sort of typical environmental cues, or maybe orders of environmental exposure. And actually, even before this paper, there was another paper that was published by Saeed Tavazoie-- hard name to pronounce-- who was at Princeton at the time. So he had published a *Science* paper just the year before-- well, couple years before-- demonstrating a related phenomenon. What was it that this other paper had demonstrated? Does anybody remember from the introduction?

So it wasn't the experiment that these guys did. But in the introduction they spent half of a paragraph describing this other experiment. Another kind of environmental cue that *E. coli* get when they enter the gut. So what's-- yeah.

**AUDIENCE:** Isn't there a temperature increase and oxygen decrease?

**PROFESSOR:** Temperature increase and oxygen decrease. So temperature goes up and oxygen goes down. And this makes sense. The inside of our bodies typically warmer than outside our body. Also, it's a relatively anaerobic environment inside the gut as compared to the ready availability of oxygen outside.

So what is it that this other paper then showed? Yes.

**AUDIENCE:** So I think they were trying to show a symmetric response [INAUDIBLE]. If adaptation in one situation occurred, then-- if it received one stimulus, it would adapt to both.

**PROFESSOR:** Yes. OK. And the word adapt, we always have to be very careful, because it could mean two very different things. And in some cases, it might be that both-- possibly it might be true, but when you say that it receives one stimulus-- so this we might call S1, and this is S2-- so these are stimuli that the cell senses-- are you saying that it might sense one, and then it adapts to both. When you say adapt, are you referring to evolutionary adaptation or something else?

**AUDIENCE:** I'm referring to a single cell implementing the appropriate metabolic--

**PROFESSOR:** That sort of cell level adaptation. The sort of default thing is you might, say, well, if you get the stimuli one and increase the temperature, you should implement the program one or the response one that corresponds to some set of genes that are required to grow in high temperature.

Now similarly with stimulus 2-- well, we'll say this is R2. So these could, in principal, have different gene network responses to these two environmental changes. But what they found is that actually, if you take E. coli and you expose it to increase in temperature, then not only does it activate the genes required to handle high temperature, but also activates the genes required handle low oxygen, and vice versa.

Now there's a question, though. If you see this, it could just be the case that you've mislabeled what really is R1, R2. Maybe it's just a single R, single response that allows the population to adapt to both of those environmental changes.

Does anybody remember what the authors did to show that that was maybe not what was going on?

You can look at your sheet of paper while I--I just want to write down the reference in case anybody wants to look at this more. This is Tagkopoulos, *Science* 2008.

What would you do if you wanted to show that this was a real regulatory strategy, this so-called symmetric anticipatory regulation?

Yes.

**AUDIENCE:** Would it be you could measure genes' pressure somehow, and then you could get the function of each--

**PROFESSOR:** Right. So you could expose to this stimulus, measure gene expression, expose to this stimulus, measures gene expression, and try to figure out, maybe, based on the difference which genes were more or less-- but the argument was somehow that in response to stimulus 1, there were some genes that were really unnecessary.

So I guess you could go and you could do knockouts of each of these 100 genes, and then try to measure the effects. And that would be a nice direct-- and it's related to something that they did in this paper with the yeast. They went and looked at the gene expression profiles.

What they did in this paper by Saeed was that they took the E. coli that displayed this, what they're arguing is a symmetric-- and based on some previous studies, I think, saying that, oh, this had been annotated to be good for temperature but not for oxygen and vice versa. So there were already hints that this kind of made sense. But what they did is they evolved the populations in just one of these environments and showed that they could decouple them.

So that's at least making you feel that it's a little bit-- that maybe this was not due to just some sense that it's the same genes that are required for protection against these two environments, but rather that because they had evolved in this environment where, when they see an increase in temperature, they often see a decrease in oxygen.

**AUDIENCE:** So when you say decoupled, you mean you evolve them in high-temperature environments?

**PROFESSOR:** That's right.

**AUDIENCE:** You expose them to decreased oxygen, and then they die?

**PROFESSOR:** No, no, no, no. So just, they showed that if you evolve the population so that it receives this stimulus without this stimulus, then it removes this arrow, that it stops activating these genes. And then I think they also did look at the functional consequences. But the first order thing is just that they saw that they could remove one of these arrows, which is related to something that they do here as well.

Now I think that it's worth thinking about the question of whether this study, published in *Science* in 2008, do you think that it's a complete scoop of the paper that you just read, which is Amir Mitchell's paper published in *Nature* in 2009? No. OK. And on one level you say, well, they still got it published in *Nature*, so therefore it was not a scoop.

But I can tell you that Amir was very worried, because he had been working on this paper for many years. And they really studied both *E. coli* and yeasts. It was a major thing. And then this paper came out, and I'm sure that he was despondent. And there will be times where you feel despondent as a result of something like this coming out, because it really is a very similar idea. The *E. coli* stuff is very similar, because it's same organism-- the idea is, in the gut.

So in what ways do you think that their paper is different, and also, what strategies did they follow in order to differentiate their paper? Yeah.

**AUDIENCE:** Maybe potentially trying to find differing regulation mechanisms, so they said, well, there's a symmetric regulation mechanism that someone talked about last year, and we've got the asymmetric one.

**PROFESSOR:** That's right. And I think that this is-- of course, we cannot replay history. But I think that they do a very nice thing in figure one of laying the groundwork to make it very clear that their work is different from Saeed's work. And the thing is that if this had not been published, they maybe would have framed their work differently. They might not have had this difference between a symmetrical anticipatory regulation



and the asymmetrical.

But by laying things out that way, it clarifies that, oh yeah, these are different things. And they're going to explore this one, whereas Saeed's group had explored a rather different mechanism. But then, of course, depending on how you look at it, they're either more or less similar. But I think that they do a very nice job of kind of right at the beginning setting out how their work is going to be different.

So they explained that there's the simple direct sensing mechanisms. You'd say that you just have different stimuli that lead to different responses. That's fine. They also have the stochastic switching, which we're going to talk about more.

The idea here is that even though we might be getting different stimuli at some rates, we would still switch into these different responses. And we'll talk about various reasons that this might happen. Incidentally, these things don't have to happen at the same rates. So it could be that stimulus one it primarily does R1, whereas stimulus 2 primarily does R2, for example. If you'd like, these arrows could have different strengths.

But then what they do is they point out that there really are-- well, this symmetric anticipatory regulation, where it's crossed, maybe I'll just write out the asymmetric one.

Can somebody say why it is that we might want to do the asymmetrical as compared to the symmetrical? Yes.

**AUDIENCE:** If both of the stimuli are appearing at the same time, then you might want to have a symmetric anticipatory regulation?

**PROFESSOR:** Great. So the important thin is here that this is, because it's at the same time, the two stimuli kind of come together, or very similar times, whereas this, if there's a clear temporal order. But you could think about this in some other way as well, maybe. Can somebody--

**AUDIENCE:** It's along the same lines, but there's a cost to having S2 also performing. Then if

you ever use that--

**PROFESSOR:** Right. So depending on the costs, it could change what's sort of optimal, perhaps, right? And I think it's also perhaps worth just saying that it may also depend upon how frequent these different environmental perturbations are, because you can also imagine situations where it's very common that the temperature goes up, but it's very rare that oxygen goes down. But there could be a sense in which this is, basically, based on the logic I just said, this would be then less predictive of this than this would be of this. Do you guys see what I'm saying?

Because it could be the case that any time that oxygen goes down, then you really know that that means that you're entering somebody's gut and you also better activate this, whereas it could also be the case that if you sit out in the sun, you get hot, it doesn't mean that oxygen's going to go down. So it could be that there is more or less information, or that each of the stimuli have more or less information about another stimuli. So even if there's no temporal order, it could be the case that something like that could lead to a situation where this might be optimal.

**AUDIENCE:** So it seems people try to make formal statements about information in this context, just like they did in the context of stochastic theory.

**PROFESSOR:** So there's some of it. I haven't seen anything that's quite as clear and formal as what Edo Kussell did in that *Science* paper you guys just read. But for example, Amir did do a paper, he wrote a PNAS paper basically laying out this model that's sort of sketched at the end of this paper a little bit in more detail in terms of explaining, depending on the probability that something occurs and the time, so kind of exploring this a bit more, but not as formal in terms of information as--

**AUDIENCE:** So it's not quite clear to me why you would have a symmetric anticipatory regulation comparing direct sensing. If you have direct sensing, you can do the same thing.

**PROFESSOR:** That's right. I think it's true. But it could also be that because these things are both fluctuating and they are measured with some noise and the environment, there could be a sense that-- because you can imagine a world in which these two things

are both uncertain predictors of whether you're in the gut. But if they both go, then that really means something, in which case then it would make sense to really look at both of them and activate both more strongly as a result of seeing both.

**AUDIENCE:** But it could also make you more susceptible to trigger a response that you don't necessarily want.

**PROFESSOR:** That's true. But if you say, oh, well this happens at a low rate and this happens at a low rate, then the probability that both of them happen at the same time would be very small. So your rate of false positive, in that sense, would be-- you can really reduce the rate of false positives by integrating both pieces of information.

And this paper, just to remind ourselves, this paper is primarily exploring the idea of asymmetrical anticipatory regulation. And their argument is because there's a typical temporal order to how E. Coli that are ingested encounter different sugar. And the claim in the paper is that bacteria will typically see what before what?

**AUDIENCE:** Lactose before maltose.

**PROFESSOR:** Correct. So the claim is that in the gut, the bacteria will see lactose before they see maltose. Now in order to evolve this asymmetrical anticipatory regulation, does it have to be that you kind of always see lactose before you see maltose?

So let's imagine that only 10% of the time that the bacteria see lactose-- or sorry, only 10% of the time that the bacteria encounter maltose that they actually saw lactose before. Since 10% is much less than 1, should that totally scotch this whole mechanism? No. And can you say why not?

**AUDIENCE:** Every time they see lactose they see maltose.

**PROFESSOR:** That's right. In some ways, that's not the right way to think about it, but in terms of a cost-benefit type argument, why is it that maybe it's OK in principle that maltose is frequently observed without lactose?

**AUDIENCE:** There's no difference between the two strains if you see maltose that way.

**PROFESSOR:** Right. And I think this is a kind of an interesting, important point, which is that if you imagine a competition between the strain that just does, we can imagine, direct sensing versus we'll call that AAR. So direct sensing versus asymmetrical anticipatory regulation.

Now if you imagine the situation where you encounter maltose without encountering lactose, then these two strategies behave essentially identically, because they both just see the maltose, they're both surprised. They both activate the mal genes, everything's fine.

However, it's in this other direction that you have to think about it. So the real question is, when you encounter lactose, you want to think about, well, what fraction of the time and how long a delay is there before you encounter maltose? Because if it's the case that you encounter lactose a lot, and only infrequently do you then come on to maltose, that's going to be a problem for evolving this into asymmetrical anticipatory regulation, or it would not necessarily be optimal to do so, because in that case you're going to be activating all those mal genes in advance, and frequently it won't be useful.

And this is relevant, perhaps, just because what fraction of the time do you think that *E. coli* going to the gut will encounter lactose?

We should be able to make a reasonable estimate of this, is my claim. All right. Where does lactose come from? Lactose comes from milk. Which animals drink milk? Well, right, OK. So we're thinking about mammals. Just take a random mammal of your choice. You might want to not choose humans just because we're a little bit atypical in this regard, although even there I think the numbers end up being kind of the same.

But imagine a dog or a bear or your favorite mammal. Which of them are drinking milk? The young, right? OK. So there's some period of one's life where you drink milk, and then as an adult you typically don't. I guess all I'm saying is that if you imagine the bacteria entering into a mammalian gut, it might be 20%. I don't know what, but it's basically proportional to the years. The kids, they may eat more dirty

things. So you have to wait or whatever.

The point is that they're going to encounter lactose if they enter the gut of a baby or young child, whereas if you enter an adult, then you won't see lactose, and you'll just maybe then directly see the maltose.

But there's no real cost associated with missing that signal vis a vis the so-called direct strategy.

I think that they do lay out some other kind of conditions that they think would be reasonable to demonstrate if you want to do the asymmetrical anticipatory regulation. And once again, I think that this is both very clear and also just lays the groundwork for the reader in the sense that they tell you what they would need to show in order for you to be convinced. And then they go and they do those things.

Now in principle, you could argue about one or the other things. But I think they do such a nice job of being clear about what would be needed that you don't even think to object about any of it. And of course they're not going to bring up the three things that they would need to demonstrate unless they actually were able to demonstrate those three things.

So the three things that they claim is that for this to have been an evolved trait that pre-exposure to S1 increases the fitness under S2 or R2-- I guess in this case-- that's right. So getting S1 helps you in S2. So seeing lactose helps you in maltose. Indeed they're going to show that.

Second, that there's some cost associated with up regulating that, in this case up regulating the mall genes. And third, there's this idea of specificity. Now the claim is that it's really S1 that is activating this R2 response, and that various other S's, in this case other sugar sources, maybe should not pre-activate.

Now in figure 1, they basically are exploring the difference between this arrow and this one. So if you expose the cells to maltose you activate the mall genes at a high level. Is it required that exposure to lactose up regulates the mall genes to the same level as maltose? No. Was it at the same level as maltose? No. But the idea is it up

regulates it some, as compared to other carbon sources.

Now they also made a comment about whether this arrow, the so-called anticipatory regulation, whether it was specific to their E. coli strain, and their claim was that it was or was not? Would it be good for it to be specific to their strain, or bad?

Probably bad, although the fact that-- but then they say, oh, it's not specific to our strain.

And then they cite another paper, which-- and this is, again, a nice way of phrasing it because they're adding strength to their paper because they're saying, oh, it's not just our strain. It's in other strains as well. So you read that sentence and say yes, that's good. Of course you could read that sentence and say, oh, but somebody else already showed this? But this is, again-- and I don't have any real complaint about that, because I think what they've done is they've really brought a set of measurements together in one conceptual framework that I think is rather compelling.

Now of course, you can always argue about whether each of these examples really did evolve for this purpose or not and so forth. But just because somebody else had seen that there was some cross-regulation of the mal genes in response to lactose. Doesn't mean that this paper is worthless. But it is interesting to note that somebody else had already seen this.

Moreover, in figure one they compare this asymmetrical anticipatory regulation for their wild type strain as compared to another strain. And what other strain did they use? Yes.

**AUDIENCE:** You didn't expose the cells to maltose after the lactose [INAUDIBLE].

**PROFESSOR:** Right. They evolved the cells in lactose for 500 generations. And then what did they see?

**AUDIENCE:** The arrow.

**PROFESSOR:** That's right. When they saw was that evolution-- and this is in lactose-- removed

that. Now 500 generations, how long do you think this is going to take them?

**AUDIENCE:** A few months.

**PROFESSOR:** A few months. Have you guys ever heard of any other paper where they evolved E. coli in lactose for 500 generations in your extensive reading of the literature that you've engaged in over this semester?

**AUDIENCE:** The cost-benefit paper.

**PROFESSOR:** The cost-benefit paper. So it's fascinating to note that this sentence, they said, "We have examined laboratory-evolved strains of E. coli. Did they say that they did the evolution? No, they didn't. And indeed, if you follow Citation 14 over here, you can say, oh, right. So it's another paper by Uri and company.

So I don't know if it was the exact same experiment, but the idea is that-- so I think what they did is, they took the strains that were evolved for another purpose, and that's fine. If somebody else already did the daily dilutions for three months that is the exact same that you were going to do, then you might as well just analyze their evolved strains, right? There's nothing that says you have to go and repeat the work. You would get the same thing, maybe, presumably.

So this is just worth highlighting. But it makes the paper much stronger that they were able to do that. So it's good to keep in mind that in some cases, what you need to do to make your point has already been done or has already been evolved and so forth.

So now what we have is a situation where somebody else, maybe, has already seen that lactose activates maltose. Somebody else already did this evolution on lactose. They didn't yet show that this regulation was removed. So that was something that these guys had to do, and indeed, it looks like it's removed.

But once again, I think there's a sense that they're really looking at this problem from kind of a new, maybe comprehensive way that allows them to see the connections between all these different experiments. And then they are doing new

measurements. But the measurements themselves are perhaps relatively simple, but the trick is knowing what measurements you need to do in order to say something that's very interesting.

So in figure three, what they went and they did is they asked about this question of whether there really is a fitness benefit associated with it. Now I think this is another case where you could have made this figure in the form of a table, and it just would not be very compelling.

So you'd look at it and you just wouldn't be very excited by it, whereas in this case, they have the figure, which are these fitness advantages associated with whether you take the strains, you first expose them to lactose and then go to maltose, or you go from maltose to lactose, or in this case GAL to maltose, sucrose to maltose.

What they showed was it was a selective benefit there, that the cells really only gain the benefit of pre-exposure to a sugar if it's lactose. And they show that for the evolved strain, not only does it no longer have that arrow in the sense of the response, but it also no longer has the benefit. It's good to do both of those measurements because it's possible you might have missed something. Yeah.

**AUDIENCE:** So but like, it's for one protocol going from lactose to maltose.

**PROFESSOR:** That's right. Yep. And this is the nature of science, or in particular this kind of science, I think, that you can collect data that is consistent with this interpretation, but it does not at all prove that this is not gut, either. And so you could have imagined experiments where you maybe mixed two strains that follow this direct versus the AAR, and you mix them together 50-50 in the gut. And then you go and you do sequencing later.

And actually, people do such kinds of experiments, and in some cases, amazingly, sometimes something comes out the other end that is interesting. I couldn't help myself. Sorry about that.

So yes. One other point on this paper, or this figure, that they even color code the data points. So they have their main thing in this black, but then the other ones are



red, yellow, blue. it's telling us that this thing is directional. It's S1, S2, S2, S1, specificity, extinction. So they are even kind of telling us what each of those data points mean.

And you might say, oh, that's kind of overkill. But given that it's just so important that you're keeping track of why we're doing each of these measurements and what it's supposed to tell us about the big story, I think that it's really just wonderful to be reminded, OK, what is this controlling for, or what not. Yes.

**AUDIENCE:** I like the colors too, but they don't stand out very well in black and white.

**PROFESSOR:** Yes, well, OK. So we can lead entire discussions about both how to write a paper and how to read a paper. How to read a paper is that you should just get a color printer, OK? So that's the advice. In terms of how to write a paper, I would say, though, that it's always good-- you should take advantage of colors, but you should choose the color scheme/scale such that it is legible or compelling in both color and black and white.

And actually there's a very nice set of articles written by Wong, who's a graphic artist that works at the Broad. He wrote maybe 15 different kind of one- or two-page articles for *Nature Methods*, basically just about how to write figures. And he talks about these different color scales. You know, there's-- what do the words mean when there's brightness? Color, hue, saturation. I can't remember what all these things mean anymore. But that's why you should read the articles. And it's fascinating stuff, actually.

So this is the basic argument they made in this paper, that because *E. coli* are exposed to a typical order of carbon sources in the mammalian gut, that that's the possible explanation for why this thing evolved that they can measure experimentally, which is that when, these *E. coli*, they see lactose, they start activating the mall genes but not vice versa.

Now, like always, you can't prove that this is what's going on. I think very interesting set of ideas. Some of the people that study gut microbes just think that it's

implausible that this would actually work out in the sense that the fitness cost benefits, et cetera, would be relevant. Do we know that the time scales even work out? I don't know how-- so there are a bunch of things that you could worry about. And I would say I don't know enough about what goes on in the gut to actually have an educated opinion on that point. But I think it's a very kind of clear exposition of the ideas. Yeah.

**AUDIENCE:** Do you have a sense for what fraction of E. coli bacteria are in the gut of animals? Do you see what I mean? Because it's just a very small fraction.

**PROFESSOR:** That's right. Yeah, and I guess-- so I myself definitely do not know what fraction of E.coli are in the gut as compared to in the soil or in the-- it's not that E. coli are the dominant bacteria in our gut, either. It's just that they do enter the gut at some rate.

And then the question is, if we knew the answer, I don't even know if it would tell us-- I mean, it gives some guidance. But there's also a question of where is it that cells are actually dividing, because it could be the case that the life cycle is that they go in the gut, and that's where they're dividing and happy. And then they end up in some soil somewhere, and they're in some stationary phase, and maybe most of them are dying at some small rate. Eventually some of them get uptaken again and then they divide, divide, divide.

So it could be the cases that only 0.1% of all E. coli, or 01% of the time. It could be a very small amount of time or something, and they're in the gut, but it could still be the dominant event in their lives. So I just don't know.

I think I'm maybe going to skip the second half of the talk, or second half of the paper, where they discuss this idea of the typical order of events in which yeast in the context of fermentation and making wine might encounter or experience different environmental assaults. And it's a similar kind of idea of cross-protection against different assaults. But in some ways, this is really sort of two mini-papers that tie together in this general thing of anticipatory regulation.

Are there any questions about this before we move on? Yes.

**AUDIENCE:** I have a question about when the baby males grow up. Do they need this mechanism? Do E. coli?

**PROFESSOR:** OK. So your question is whether E. coli in the gut of adults have this. I don't know, but my feeling is that these things are probably fairly transient in the sense that you probably don't have-- well, certainly, from the standpoint of entering the gut, the bacteria don't have any choice about whose gut they enter.

So it probably is going to be the case that they have to have a basic strategy that is going to work across both adults and the kids, although it's true that the gut communities in kids are different from adults. But it's hard to know what to make of that. Yeah.

**AUDIENCE:** So here they decouple the two responses. Are there any examples of researchers coupling them, like Pavlov's?

**PROFESSOR:** Yeah. That's a good question. I know that I've heard people talking about it, but I can't point to a specific example of where it's been convincing. That doesn't mean it's not out there, though. It's a very important, interesting question.

So what we want to now is kind of switch gears a little bit and switch to this topic of-- it's going to be kind of partly about phenotypic heterogeneity, but partly again about fluctuating environments. And the way that those two topics are related is via bet hedging, but I just want to highlight that bet hedging is a way of using phenotypic heterogeneity in order to survive fluctuating environments, whereas it's not the only explanation for this phenotypic heterogeneity.

So yeah, I think maybe we'll-- so first of all, on the topic of phenotypic heterogeneity, when we say that we are often referring to the fact that even clonal populations can be remarkably heterogeneous in a given environment. So we'll say phenotypic heterogeneity in clonal populations, because of course, if there's a lot of genetic diversity in the population, then it's not, perhaps, a surprise there would be phenotypic heterogeneity.

What we're going to do is go through these three different possible explanations

and discuss, maybe, canonical biological examples of each. And the examples that we're going to give-- well, the classic example is in the context of seed germination. But more recently there's this question of the persister cells, or the persistence phenotype.

Now in both of these cases, it's really a way in which a clonal population can send a small population of either seeds or cells into a protected state that's protected against some assault. So seed germination, this is a case where the seed, if it doesn't germinate, it's protected against draughts. Draught. And this one, maybe, is antibiotics.

What I want to do is just make sure that we understand the seed germination idea. And so this is, from the standpoint of ecology, these ideas came out in the '70s, at least. So what we can imagine is a situation where you are some, maybe, desert annual. And so you send out these seeds.

Now most of the time, you get enough rain, the seeds are fine. They can germinate, sprout the little seedlings, and they'll be able to grow up just fine. But at some rate there might be a severe drought that would kill all the cells that actually germinated.

So the question is, what might be the optimal strategy? So for example, let's just imagine that we live in a world in which 75% of the time there's rain, 25% of the time there's no rain. Now if you germinate-- so a particular seed, if it germinates, it pops out-- then it will expect to send two seeds to the following year, whereas as it germinates and there's no rain, then it only has a 10% probability of surviving, or that it could on average send 1/10 of a seed to the next generation. Now of course, you don't have to germinate. If you don't germinate, then what happens is that you just send one seed to the next year. If you'd like, you can include some probability that the seed doesn't survive, but for simplicity, we can do that.

So the idea is that you don't know whether it's going to rain or not. But what you can do is you can either germinate or not germinate. And then this tells us how many seeds make it to the following year. Do you understand the framework?

So the question is, which one of these strategies has a higher long term growth rate in terms of the number of seeds that you have in the population?

I'm going to give you 30 seconds just to kind of think about this, play with the numbers a little bit. And the question is, should you just have a strategy where you germinate or don't germinate? And then what we're going to find, maybe, is that the optimal strategy is to do this probabilistically. But it's useful to just play with such numbers for a little bit.

The physics office, unfortunately, is trying to steal our cards from us, so we don't have them. But what we're going to instead do is, if you think you want to germinate you raise your left hand. If you don't want to germinate, raise your right hand.

OK. Ready, 3, 2, 1. So we have a fair distribution of-- it really is kind of 50-50. That's interesting. All right, let's go ahead and-- yeah, let's spend a minute. Turn to your neighbor and try to convince them that it's either better to germinate or not.

[SIDE CONVERSATIONS]

You can start with 100 of one, 100 of another, and then just--

**AUDIENCE:** Start with 100 and then you just go--

[SIDE CONVERSATIONS]

**PROFESSOR:** Let's say 10 to the 4. So we're trying to not think about this--

[SIDE CONVERSATIONS]

So there's another question about short term versus long term. And we're actually trying to not worry about that. So you can think about large populations, so it's really--

**AUDIENCE:** This is after one year. You will maximize your expected value.

**PROFESSOR:** That's why I said the long term growth of the population. So if we just compare-- so why don't we reconvene. Various people have been talking about the question of

how many seeds you're starting with, what the population size is, because you have to worry about extinction. And that's another kind of interesting question of short term versus long term payouts of fitness. That's actually not the effect that we're trying to get at right now.

So we can even think about this in the limit of just large population sizes. And we want to compare the two strategies. And for any strategy you choose, including the probabilistic strategies, there's going to be a long term growth rate of the population. So it will grow or shrink exponentially over time. And what you want to do is you want to find the strategy, in the context of bet hedging, you want to find the strategy that maximizes this long term growth rate of the population.

And this is essentially what Edo talks about in his paper, but in a more mathematical, general-- because he's talking about  $n$  different phenotypes and environments and arbitrary switching rates. And it's beautiful and general and he makes nice connections to information theory. But it's possible to read that paper and then not actually understand what bet hedging, or what any of these things are, quite.

So it's good to take the simple thing, and then it's also good to know that somebody like Edo has thought about this very deeply and figured out all these beautiful things. And I know that paper was hard, so don't-- but it's really like a two-page paper if you look at the text. So it shouldn't take too much time just to get a sense of what he's thinking about. But then it's good, also, to be concrete in a simple situation.

So what we want to know is that, over the long term-- then what you can do is you can just say, well, the ratios of rain, no rain over the long term will be 3 to 1, because it's over a long term. It doesn't matter what order they come in because over the long run it's all going to average out. We're assuming we're a large population size, so we don't have to worry about stochastic extinction.

The question is, since we don't have to worry about stochastic extinction, then can we just say that we just calculate, for example, the mean weighted by these

fractions of these numbers? And actually, the subtle possible thing is that it's the geometric mean that is relevant, not the normal mean. And that's somehow tricky when given these numbers.

But it's sort of obvious if you just switch this to a 0. Because if this is a 0, then is germinate going to be a good strategy? No, right? Because what that would mean is that if everybody always germinates, it may be that this number could be 1,000, and it's wonderful except that if, at some low rate, there's no rain and this is a 0, then the population's dead. And you'll always encounter that eventually. So that just highlights that it's the geometric mean you want to be focusing on.

And in particular, we know that on average, over the long term, there are going to be three of these years for each of these years. So we can just multiply the numbers. So it's really going to be there's going to be three times where you double for each one time that you divide by 10.

What you see is that over the long run, this population shrinks, whereas in this case it's 1 times 1 times 1 times 1. So the long term growth rate/shrinking rate of these populations is that in the long run this population does not change in size, whereas over the long run, this population shrinks.

**AUDIENCE:** So in the geometric mean the expected change in the log population, which is also the expected percentage change in the population. Is that right?

**PROFESSOR:** Yes. So I think that if you take the logs, it ends up being equivalent. But you have to make sure you just do it all. And I guess what I want to highlight in this case is that, A, you have to be careful of taking averages, but B, to get the intuition behind the situations where bet hedging, I think, is really valuable, it's situations where this thing gets very small and this might become rather large.

So you can imagine that this could be 200 and this could be 10 to the minus 9. Or if you'd like, you could say that's 0. And it's clear that if this is 0, then you can't germinate all the time, because then the population's definitely going to go extinct over the long run. But this is also a situation where there's real tension between

long term and short term situations if this probability is very small.

But it's in these situations that it pays, in terms of the long run growth rate of the population, to do this some fraction of the time. And it could be very small. It could be that only one in  $10^4$  of the cells enter this persister state that is resistant to antibiotics. But it can maximize the long term growth rate of the population.

There are various models that you can write down for either seed germination or this or that. But the relevant thing is, for example, in the persister case, it might be that the rate of persister formation might only be one cell in  $10^5$  or  $10^6$ . So a very small sub-population of the cells are in this so-called persister state.

But those cells, they have a couple of properties. One is that they are dividing slowly. That means there really is a cost to entering that persister state, but if it's only one of  $10^5$ , that means that it's really a rather small cost in terms of decreasing the growth of the population. And those cells also could be resistant to antibiotics. So you can select for them in the sense that you add a bunch of antibiotics and maybe only the persister cells are left after 12 hours.

But even though I say you selected for them, it's not a true antibiotic resistance kind of mutation, because if you take that persister population, you let it grow back up, what you find is that it is still sensitive to the antibiotic.

Now it's also possible that at some low rate, maybe one in  $10^8$ , you evolve resistance to an antibiotic. That's genetic resistance that stays with you, whereas the persister kind of resistance, this is something that is a phenotypic switch. And when it grows back up, it kind of reverts back to the original phenotype, which is being sensitive to the antibiotic. So that's the classic way that you distinguish between genetic resistance and this persister-based or phenotypic resistance. Yes.

**AUDIENCE:** What would be the mechanism of something like that?

**PROFESSOR:** That's right. So people argue a fair amount about this. And there's a thought that there are these toxin, anti-toxin modules in bacteria that are normally suppressed but can be triggered, and then can lead to these states entering, although this ends



up kind of a muddy debate/subject somehow.

But the idea is that there's something in the cell that is normally repressed, but then at some rate activates and causes the cell to-- and of course, you can always argue, OK, did it evolve for this purpose or not? Because also, if you're dividing slowly, then you're going to be resistant to many things.

But then there's also arguments that there are somehow more specific persistertype states against particular assaults or particular antibiotics. So I think that there are many subtleties to this whole debate.

Any other questions about this idea of bet hedging? So I think that bet hedging is indeed one possible way in which evolution could select for phenotypic heterogeneity in a population. So this is a population of seeds that are all, in principle, identical. You put them out and you plant them all and they see the same environment. They don't all germinate. But it's not that all those seeds that didn't germinate are dead, because actually, the following year, they still go back and they can follow some stochastic strategy.

So I just want to stress that right now we're considering the two extremes, but for bet hedging, you say probability  $p$  germinate,  $1 - p$  don't germinate. And those are, in general, going to be some probability  $p$  that can be non-0 that maximizes long term growth rate of the population.

But I do want to talk about two other possible explanations, which is this idea of mixed strategies and altruism. Now mixed strategies, in some ways, is an obvious one, because we just got done talking about mixed strategies in game theory, where there is this notion that in a given game, sometimes this Nash equilibrium was a pure strategy, or the evolutionary stable state. But sometimes at the equilibrium there was coexistence of these two strategies, and that could be either implemented by two different genotypes, or by one genotype implementing both phenotypes.

So this was the classic game that people talk about in this area is this hawk dove

game. We talked about this just a little bit in class. So this is a model of animal conflict. And the idea is the animals can either be-- so the idea is that there are two animals and they come across some resource. And the resource could be food or it could be a mate or something else. And the question is, should those animals fight for the resource or should they back down?

Now the assumption in all this is that the hawk is the strategy that fights and the dove is the strategy that backs down. Now the question is, how do we assign some sort of costs and benefits to these various strategies? Again, where this is the situation where we take the two-player game and then we'll kind of generalize it to the population. Well, if two doves encounter each other, then they don't need to fight at all, and they just split the benefit. So there's some benefit to the resource, and they each get half of it.

Now of course, if a hawk meets a dove, the hawk actually does better, because then the hawk gets the entire benefit, doesn't have to pay the cost. So already we can see from what I've said so far, is dove a Nash equilibrium? No, because you can already see that if your opponent is a dove, it's better for you. If you start in the situation where everybody's a dove, so everybody's getting  $b$  over  $2$ , then one individual has the incentive to change strategies and get the full benefit. That hawk never has to fight anybody because they will all back down.

Now of course this dove that fought that hawk doesn't get any of the resource, but also doesn't have to fight. However, question is, what happens if two hawks meet? And that's a situation where they really go at it. And what you might say is, then, there's a 50-50 shot, that each individual has a 50% chance of getting the benefit, but then also has to pay this cost. And then you can decide whether you want to have-- so I think it's normally parameterized to  $b$  minus  $c$  over  $2$ .

The idea here is that you have a 50% chance of getting the benefit, but you have a 50% chance of getting beaten up. OK. Very simple model. Now the way that we normally parameterize this is that-- let's say that we want the hawk-- wait, hold on a second. Under which circumstances will the hawk strategy be a Nash equilibrium?

**AUDIENCE:** If  $b$  is larger than  $c$ .

**PROFESSOR:** That's right. If  $b$  is larger than  $c$ , then the hawk would be a Nash equilibrium, which is actually what we don't want in this game. So what we typically assume is that the benefit is less than the cost. So in that sense, it's somehow not worth fighting for. And that means that neither strategy is a Nash equilibrium. So neither the hawk nor the dove is a Nash equilibrium. And indeed you can calculate what the Nash equilibrium is, and the Nash equilibrium is to follow a mixed strategy with some probability  $p^*$  that is, I think,  $b$  over  $c$ .

So if there's a bigger benefit,  $p^*$  is the probability of following the hawk strategy.

But as the costs grow, then this Nash equilibrium probability decreases. So this is a case where it's a mixed strategy that is this Nash equilibrium. And so this is a situation that displays what we call negative frequency dependent selection. You can see that the rare strategies do better than common strategies. The key way to think about this is that if you have a population of doves, then a rare hawk does better. However, if you have a population of hawks, then it's a rare dove that does better.

And it's that kind of negative frequency of dependence that pushes the population towards some equilibrium between the phenotypes, or it could be equilibrium between genotypes as well, because you could have pure strategists of a hawk and a dove, so different genotypes that just are different. And the differences in fitness will naturally lead to evolution to an equilibrium hawk kind of frequency in the population that is equal to this.

So this tells us either about the equilibrium frequency of the genotype of hawk in the population. Or it could be that it's a single genotype that implements these two strategies with some frequency that is  $b$  over  $c$ . So I'd say that this is a simple idea from the context of game theory, and the question is, what circumstances might this actually arise in kind of real populations.

Oh, before I say anything else, I do want to stress that all of you have played mixed

strategies before, because in the Rock Paper Scissors game, which we're going to be talking about more in the coming week or so-- you guys all know this, that rock beats-- not paper-- but rock beats scissors, scissors beats paper, paper beats rock-- so the Nash equilibrium of that game is to do  $1/3$   $1/3$ ,  $1/3$ , because if everybody's following that strategy, nobody has the incentive to change strategies. Definition of the Nash equilibrium. Is that the optimal strategy against any opponent? No.

And this is very important to point out, because when you think about these Nash equilibrium and everything in abstract terms, it's easy to get confused about-- when we say that it's-- sometimes we say optimal, or it's the solution to the game, that doesn't mean that it's the best response to any given strategy.

And you know that, because if you know that your little brother is going to play rock, then you know-- well, depending on whether you want to make him happy or not, you decide what to do, right? So it's clear that this  $1/3$ ,  $1/3$ ,  $1/3$  thing is, it can be the solution to the game of the Nash equilibrium without being the best response to any given strategy.

So one of things that we're doing in my group, actually, is we're arguing that in the case of yeasts that are exposed to mixed sugar environments-- so yeast in environments that contain a little bit of glucose and a little bit of galactose-- what we and actually other people have found is that there's a bimodal response. So if you look at the number of cells as a function of the GAL operon expression, so GAL is the expression of the genes required to break down galactose, what you see is that some of the cells activate the GAL genes highly, some don't activate it at all.

So the question is, what might be an explanation for this phenotypic heterogeneity observed in a clonal population? And one way to think about this is that it could be the implementation of some sort of foraging game. So you can imagine a population of animals that want to go and eat some berries. So there are two different food sources-- say there are blueberries and there are red berries.

You can imagine this is the kind of situation that would lead to negative frequency dependence, because if everybody else goes and eats-- I guess this was the

blueberry-- so if everybody else goes and eats the blueberries, then it's better for you to eat the red berry because you don't have to share with anyone, whereas if everyone else goes and eats the red berries, then it's better for you to eat the blueberries.

So that's the kind of situation that would naturally lead to the so-called negative frequency dependent selection, in which coexistence of the phenotypes is somehow the equilibrium condition, and that can be implemented either by different genotypes or it can be implemented by a single genotype that follows both phenotypes.

Incidentally, is this equilibrium-- does it maximize the fitness of the population? No. So what is it that would maximize the fitness of the population? If everybody were a dove, because-- and you can go ahead and do this calculation, find out what the fitness of everybody is, but what you'll see is that  $p^*$  does not maximize fitness, and does not maximize payoff. So in these games, the equilibrium is not necessarily fitness maximizing.

So so far what we've done is we've given two examples of different evolutionary drivers for phenotypic heterogeneity in a population. And I just want to highlight what I think is maybe the last big one, which would be altruism. And sometimes people call it division of labor, or sacrifice or so. And a nice example of this is colicin production.

The idea is that in many bacteria, they express these colicin, which are toxins. And we're going to read more about this later. So colicin is a toxin. And the perhaps surprising thing is that in many gram-negative bacteria, like *E. coli*, when they make this colicin, the only way it's released is cell lysis.

Now the question is, why would I possibly do this? And the answer is, perhaps, that if there's a plasmid that encodes the genes to make this toxin, it also encodes an immunity protein. What that means is that if one cell lyses-- so it bursts open and releases the colicin-- it will inhibit other bacteria, but not the bacteria that are the clone mates that are related, so not the bacteria that also carry that plasmid.

So it's important to note there that now when we talk about kin selection or relatedness, we're talking about the other cells that carry that gene. So what's relevant is that the plasma that encodes this toxin will inhibit other bacteria that do not encode the toxin.

Now you can see that this kind of lysis behavior has to be stochastic, so there has to be phenotypic heterogeneity. And why is that?

**AUDIENCE:** Everyone always lyses.

**PROFESSOR:** That's right. If this plasmid always makes everybody lyse, then it's not going to get very far, right? But you can kind of imagine situations where, if you just 1% of the population lyses but then inhibits the cells that don't care that plasmid, then the plasmid can spread in the population. So this has to be stochastic.

So I think that this is another general explanation for why there might be phenotypic heterogeneity in the population. All of these three explanations are both conceptually different and they're experimentally different. They make different falsifiable predictions about what's driving that phenotypic heterogeneity.

In this field, I think often when people see phenotypic heterogeneity, they immediately assume that it's bet hedging, but these other two explanations, I think, are just as general and may be just as common. So we really have to go and make measurements to try to elucidate what's going on in each case.