[SQUEAKING] [RUSTLING] [CLICKING]

FRANK SCHILBACH: OK, so what happened last time? So we talked about the workhorse model of classical or classical economics, in terms of discounting-- the exponential discounted utility model. This is a very useful model for many settings. It's been very useful for explain and understand a bunch of behaviors.

However, this model has a few assumptions that create some predictions that are not necessarily borne out in the data. And we're trying to improve on those assumptions and make the model more realistic. What are these assumptions and what kinds of examples did I show you that perhaps were not consistent with the exponential discounting model?

I'll actually show them to you. But can somebody explain to you what they are and why those are not consistent with the exponential discounting model? Yes.

AUDIENCE: You shouldn't determine-- you shouldn't care about whether it's a [INAUDIBLE].

FRANK SCHILBACH: Right. So one key assumption in the exponential discounting model is that there's one parameter, delta, that essentially steers how much you care about different periods of time, the present and the future, or two future periods. The key assumption here is that how much weight do you put on different time periods, on early and a later date, is constant across time.

So from today versus tomorrow, there's the same difference compared to, like, when you look at a year from now and a year and a day from now, or 20 days and 21 days. The difference, if you care more about the present, that's constant across time from today's perspective, also from future perspective.

That's one assumption. And what evidence did we see or show that that's not consistent with?

AUDIENCE: Choosing a reward. Do you want the money now or is this today or in a year or [INAUDIBLE]?

FRANK SCHILBACH: Right. So we saw some evidence that people are impatient in the short run. So when it comes to today versus tomorrow or today versus in a week from now, people tend to be fairly impatient. They tend to discount the future a lot.

Now, there's sort of two issues with that. One is, if you do the same thing about asking people about future rewards, people tend to be more patient. So the discount factors that you elicit, we'll talk a little bit about that in a recitation, as well. The discount factors are just different for the same intervals of time.

Second, if you take the short run and patience seriously-- if you sort of say, OK, I take your choices from today versus tomorrow or from today versus next week and [INAUDIBLE] run that forward, you got sort of absurd implications as in, like, if you really seriously discount a week from now by a lot, then you shouldn't care about a year from now, five years from now. You shouldn't care about that at all.

However, we see in the real world, people have lots of-- people care about the future a lot. People save money. People get education and so on and so forth. So it can't be in some ways that they're really impatient in the long run, as well. So that's, like, number one.

What about point number two? Yes.

AUDIENCE: Your decision now for the future should be the same as [INAUDIBLE].

FRANK SCHILBACH: Exactly. That's called dynamic or time consistency. That's an assumption or implication of the exponential discounting model. That is to say, if I make some choices for the future, for future periods, about how I-- what I'm going to study next week or in a year from now, or whether I'm going to exercise and so on, for choices that involve time in the future, once that time comes, absent new information or absent circumstances change, I should follow through with my plans.

And so if that's not the case, then what we say is that preferences are dynamically inconsistent. That is to say, I have all sorts of plans of exercising next week and so on. The next week comes. Surprise, I'm not following through. That essentially means my preferences are inconsistent. And that's sort of evidence against the exponential discounting model.

OK. Number three, demand for commitment. Yes?

AUDIENCE: [INAUDIBLE]. There are [INAUDIBLE] people will often do that [INAUDIBLE] and that shows that people [INAUDIBLE].

FRANK SCHILBACH: Right. So one sort of key assumption or one sort of key insight from sort of economics in general, or classical economics, and it's like, more choices are good. Like, if you have more options available, that's great for you, because you can take advantage of them. You never know whether there's going to be uncertainty or certain circumstances change.

So you should never restrict your choices because you never know. Maybe in the future, maybe circumstances change. Maybe I got really busy, or whatever. Things happen. So more choice tends to be good. So in the neoclassical or classical economics model, there's no scope for or no reason to restrict your choices because more choices in the future are always good.

However, if you're worried about your future self misbehaving in certain ways, if you're worried about your future self not exercising, eating too many potato chips, and so on and so forth, you might want to engage in commitment devices, either by taking away certain options, like restricting your choice set to a smaller choice set, or, less drastically, making future choices more expensive.

That is to say, for example, if I say I'm going to submit or write a paper by the end of this week, if I'm going to tell you if I don't do that I'm going to give you $100, well, what I'm going to do essentially is I'm making not writing the paper more expensive.

And the neoclassical or the classical model had the exponential discounting model has essentially no scope for that. Why would I ever do that? Because who knows. Maybe I'll get sick or stuff happens. So there's no reason to do that. There's no benefit of doing that. And it's a potential cost.

So if I see myself or somebody else restricting their choices, that's a violation of the exponential discounting model. And the reason, of course, for that is number two. We know there's preference for it. If you think there are future preference reversals, if I think that in the future I might change my preferences in a way that I don't like, if I want my future self to behave in certain ways that I prefer right now, I might want to restrict my choice set. I might certain unpreferrable-- like, less preferable options more costly.

OK. Good. So then what we did next is to say, OK, now-- so I showed you a bunch of evidence of that. And then what we did is then we sort of presented a slight deviation from the exponential discounting model, which is the quasi-hyperbolic discounting model. That model is, in some sense, very similar to the exponential discounting model that I showed you, except for some small difference. There's an additional parameter that's called beta.

What does this beta do? It's like a short term discount factor that essentially guides all discounting or determines all discounting between the present period, right now, versus anything that's in the future. One broad question is like, what are the time periods? Is it daily? Is it yearly? Is it monthly or whatever?

There's actually quite a bit of research on this that people try to figure this out. Usually we think the beta or the time horizon of the beta is actually just a few days, even just a few hours. So for our purposes, we can say it's daily or today. And anything that you'll see in the problem sets and so on, it needs to be always defined.

But think of beta essentially anything that's like today or in the next couple of days. Anything outside of that scope is essentially outside of the scope of beta and sort of in the future. Now, what does this beta do? Well, it helps us to separate discounting for short horizons and long horizons, because if you think about discounting between the present and anything in the future, the data can steer that, right?

If beta is 0.7, 0.6, or 0.8 or whatever, I can now be impatient for the present. I can now prefer, say, $10 today versus $12 in the future. The beta will make sure that that happens. But when you look at all future periods, the beta is in all future periods.

It appears everywhere in the utility function in all future periods. So the beta does not affect how am I discounting any future time periods. So then the delta is steering that. We usually think that delta is close to 1. It's something like yearly 0.95 or 0.99 or the like. So delta is essentially reasonably close to 1, such that I can be fairly patient between 10 years from now and 11 years from now and 15 years from now.

It doesn't matter much to me. There's only going to be some slight discounting. I prefer probably things that are happening five years from now from things that are happening 30 years from now for the same sort of stuff happening. But people tend to be fairly patient. Do you have any questions about that?

OK. So let me now sort of spend, in fact, the rest of the lecture just talking about details of this model, in part sort of building intuition, sort of showing you the mechanics of how that works and how, essentially, having an additional discount parameter of the data affects people's choices.

And second, we're going to talk about something that's quite important, which is the question of sophistication versus naivete. What do I mean by that? Well, in the exponential discounting model, by assumption, the preferences are time consistent. What does that mean? If I make plans for the future, I will follow through with those plans.

So in some sense, it doesn't really matter what I think I'm going to do in the future, because essentially, I'm going to do anyway what I'm going to plan anyway. So the only thing I need to know is right now what I'm going to do. And then if I make that plan, unless circumstances change, I'm going to just follow through.

That's not true here anymore, as I'm going to show you. So now it's the case that if I make some plans right now, these plans might change in the future because of my preferences. Even in the absence of new information or new circumstances, once tomorrow comes or next week comes, I might sort of prefer different things. So now the preference are not dynamically consistent anymore. So what I prefer today might not be what I prefer next week.

So that now leads to a complication. Now I need to understand, well, what are my beliefs about what I'm going to do in the future is quite important. Do I have correct beliefs? Do I understand that I'm present biased? Do I understand that in the future I might be impatient, as well?

Or am I naive? Do I think right now I'm present biased, but tomorrow comes or next week I'm going to exercise, do the problem set early, and so on and so forth? And as we're going to see, that will lead to quite different-- starkly different conclusions or implications that are quite important for outcomes, welfare, and so on.

OK. So let me give you some examples first. So first, if you have a discount function for beta equals 1/2 and delta close to 1-- for our purposes, often I'm going to set delta to 1, not because I necessarily think that delta is 1. Delta is probably 0.99, 0.95, or the like. I'm just setting it to 1 to make the algebra or sort of any problem set exercises and so on being less messy.

Speaking of which, problem set exercises-- Pierre-Luc who sits in the very back and checks whether you are using your laptop only for classwork, Pierre-Luc sort of mostly in charge of doing the problem sets. So he also has office hours, to which you can come and ask specific questions about the problem sets. Of course, he's also very happy to answer questions about Piazza or in any other way. When are your office hours?

AUDIENCE: [INAUDIBLE].

FRANK SCHILBACH: 1:30 on Friday. And where are they?

AUDIENCE: [INAUDIBLE].

FRANK SCHILBACH: Perfect. So it's on the course website. So if you have specific questions about the problem set, again, in the recitation, we'll discuss quite a few issues related to that. But then if you have specific questions about the problem set, Pierre-Luc will be the best equipped to answer those.

OK. So if you have quasi-hyperbolic discounting preferences, with beta equals 1/2 and delta equals 1, your discounting function looks, essentially, as follows. The first period is sort of normalized to 1. [INAUDIBLE] like one about the current period.

In the future then it's, like, beta delta, beta delta square, beta delta to the power of 3, and so on. But if you sort of now use delta equals 1, you get, essentially-- you care about the present 1. And everything in the future is discounted by 1/2. OK, so relative to the present period, all future periods are worth less. So I care a lot about the present. That's why we call it sort of present bias or present focus. I care a lot more about the present than anything in the future.

Now, second, all discounting in this specific example takes place between the present and the immediate future. So I only care about-- in terms of discounting-- all future periods are the same for me. I don't care about whether things are in five periods or six or seven periods from now. That's all the same for me, because delta, in this case, is 1.

Instead, what I care about is that today versus tomorrow or today versus three days from now, and so on. Again, usually delta would not be one. So if you had a delta that's, like, 0.99 or the like, there would be some discounting in the future, as well. But not in this case. So that's sort of capturing the intuition that the long run-- in the long run, we are relatively patient. There is essentially no discounting happening anymore.

And so the utils a year from now are just as valuable as utils in two years from now. So then that also means that the timing of decisions or-- the decisions are sensitive to the timing of costs and benefits, when exactly did they occur. It matters a lot whether stuff happens today versus tomorrow or today versus three days from now. It matters much less about whether stuff happens in three days versus four days from now. I'm going to show you some examples of that.

But just to sort of give you some sense of, like, why is this called quasi-hyperbolic or hyperbolic discounting? So have your three functions. I'm not sure how well you can see this. But there's, like, a sort of pink line that's essentially like an exponential curve. That's essentially what exponential discounting looks like.

Essentially what you see is sort of a constant, if you want, sort of decline of like the discount function. And then I've also plotted to you hyperbolic discounting, which is essentially like a true hyperbolic curve. That's essentially the thing that you see on the left, the one that's sort of like much steeper at the beginning.

And then there is quasi-hyperbolic discounting that's essentially quite similar in shape, as you can see, to the hyperbolic discounting, except for the fact that it's not continuous. There's a jump from going one to beta at the very beginning. It's not a continuous function. Essentially, there's a discrete jump between today and the future. And then essentially it behaves very similarly to the exponential discounting curve.

But that's just to tell you it doesn't really-- you don't have to worry about the hyperbolic discounting and so on. We're not going to discuss that very much. But that's just for you to understand why is that the name? Well, it's because essentially it's quasi-hyperbolic. It's not really hyperbolic because it's not really hyperbola. But in fact, it looks quite similar, effectively.

So then when you think about goods that you might consume, there's different types of goods. There's leisure goods and there's investment goods. Let's sort of start with leisure goods. What's a leisure good? A leisure good is a good that has immediate rewards and delayed costs.

One example would be eating candy. So eating candy, of course, you get immediate utility benefits. It's pleasurable to eat candy. So that's positive. In this case, suppose the benefits are the pleasure of eating candy right now.

But often eating candy, it leads to delayed health costs or delayed costs at the dentist's office, which summarizes C health of 3 in the future. If you now say, well, if beta equals 1/2 and a delta equals 1, well, you can sort think about, would you eat candy today or candy in the future?

So first starting off with eating candy today-- well, if you sort of calculate the costs and benefits of eating candy, the pleasure that you get from candy is 2. If you discount the future by beta, essentially everything in the future is discounted by beta, it's like 1/2 times 3. So if you sort of do that calculation, you get something positive.

So if you think about eating candy right now, the answer is yes, because 2 minus 1/2 times 3 is positive and you're going to do that. If instead you think about, should I eat candy in the future? Well, you get different implications. The answer is no. Why is the answer no? Because if you're in the future, everything is discounted that's happening in the future, both the pleasure that comes from eating candy next week, but also the delayed costs are essentially discounted by beta.

So then you essentially just have, like, 1/2. This is the beta for everything in the future times two, which are the benefits, minus 3, which are the costs. And that's smaller than 0. So you're not going to do that. And that's sort of a general feature for leisure good, as we call them. These have immediate rewards and delayed benefits.

And so what you tend to do is you tend to overconsume them in the present relative to long run plans. Yeah?

AUDIENCE: Are you only considering [INAUDIBLE]

FRANK SCHILBACH: Yeah. So yeah, you can think of this as, like, a shorthand form of saying, for example, suppose you eat a lot of candy or, like, on one day you eat candy and suppose you know-- these are repeated choices, but I'm sort of aggregating it into one. You'll have, essentially, health costs in the future. Suppose you get dental pain in the future. Each period you have dental pain starting at age 30 or 40 or 50.

Now, you can sort of think of this as like each day you're going to have negative utility. What I'm doing now is I'm sort of collapsing that into one day in a sense of just making sure that-- to make it sort of tractable. So you can think of the delayed health costs as a summary of the health costs overall.

Instead you could also write down 0.1 for your 40, 0.1 for your 41, and so on and so forth, and sort of aggregating in that way. But that's just for simplicity here. Similarly, you could also say each day you'll have essentially a choice to eat candy. And these things sort of add up. And again, sort of simplifying in a sense of saying, like, should I eat candy overall, which is a number of different choices that each are between the present and the future. That's just like changing the units in some way overall. I think the structure of the problem is the same.

Yeah?

AUDIENCE: So obviously this is kind of comparing apples and oranges. But I the one kind of [INAUDIBLE] is your health is kind of the sum of a bunch of actions that you did over your lifetime that affect your health, versus candy is something that-- the pleasure that you get out of candy is exactly-- fully [INAUDIBLE]. So when people-- when you do this kind of [INAUDIBLE] discounting and you're trying to compare health in the future, like, pleasure now, the fact that the candy contributes such a small amount to your overall health, it seems kind of like an apples to oranges--

FRANK SCHILBACH: Yeah. To some degree, I think that's right. I think it's, in some sense, that's just an example. So I could show you many different things that are sort of similar in structure. I'm using candy just to make it concrete. You could do the same with problem sets and other pain, or exercising, and so on and so forth.

I think at the end of the day, I think that's exactly right. For a lot of choices that people make-- be it brushing their teeth, be it exercising, be it eating healthily and so on, often the costs or the benefits are very salient in the present. They're very salient and concrete. Right now you have a cookie in front of you. You can eat it. You'll be happy, and that's for sure.

Sometime in 20 years from now, you might be sort of like unhealthy. You might get dental problems and so on and so forth. Not only is it sort of very far away, but it's also very diffuse. It's also maybe unlikely. Like, there's a chance of getting a heart attack at age 70. Maybe, maybe not. Maybe you'll be just fine.

So it's more complicated in a lot of health decisions. So there's other research saying, like, well, partially it's about uncertainty. Partially it's about the future being diffuse and very sort of unclear. One part of that people tend to think is time preferences. But I think what you're saying is the world is more complicated than that. And that's for sure true.

What I'm sort of saying is, like, thinking about time preferences can help you understand some choices that people make over time. Having said that, there's lots of other aspects here, which we're sort of grossly simplifying. But I think the key part here-- I think what's really important for a lot of health choices tends to be that people-- that exactly as you say, there's a lot of-- like, thousands of really small choices that sort of add up to something like-- what we call it unhealthy behavior in some way.

And each of them are concrete in some sense in terms of the costs and benefits. Like, brushing your teeth is like costly in some sense for five minutes or whatever. Eating candy is happy for a few minutes, as well, very concretely. And then the costs are often diffuse in the future and so on. That's really important. We're going to talk a little bit about this. But we sort of-- for now at least abstracting, essentially, that away. Yes?

AUDIENCE: [INAUDIBLE]. Why is it better than the [INAUDIBLE]. Why is it [INAUDIBLE] and what actually [INAUDIBLE] better than the exponential model?

FRANK SCHILBACH: So the exponential model would not predict you this type of behavior. So if you think about now we get to this-- like, if you sort of say you want to make choices for the future, the exponential discounter would give you the same choices for the future versus what actually-- when the future arrives.

So I'm going to show you a sort of-- more in the context of problem set, I'm going to show you this behavior that the exponential model cannot-- just not explain. Like, the exponential discounter would say, either I like candy and I don't care about-- or I just don't worry that much about the health costs and I eat candy, or I don't. I'm really worried about the health costs, and therefore I don't eat candy.

And that choice is the same today versus next week versus two weeks from now, and so on and so forth. So once I make a choice for eating candy-- so here notice that the person who will make the plans to not eat candy in the future, once the future arrives, that person will change their choice and then actually surprise themselves, potentially by eating candy.

AUDIENCE: [INAUDIBLE].

FRANK SCHILBACH: There's just behaviors in the world that the exponential discounting model cannot explain. That is one of them.

AUDIENCE: [INAUDIBLE].

FRANK SCHILBACH: So there's-- yeah. So I think there's a question on what other models can explain this. I think the hyperbolic model is, in fact, quite similar, as it sort of showed you in the curve here. You sort of see the discount functions are quite similar. There's research using the hyperbolic model.

In economics in the last sort of-- in behavioral economics in the last something like 20 years, the quasi-hyperbolic model is the model that everybody uses, partially for tractability and simplicity for no other good reasons in many settings. But in some sense, what we care about-- in some sense, I don't care actually that much about which model among sort of a class of model that are quite similar can explain things sort of-- if two models explain the same thing in similar ways, I don't care about which one we should use. In a sense, the simpler model is sort of easier to use, and that's why we sort of do it.

What I care about is there are some predictions from the exponential model that are sort of just not-- that are just false. And those we need to improve on. That could be either through a hyperbolic model or quasi-hyperbolic model.

To keep things simple, I'm using the quasi-hyperbolic model, in part because it nests very simply in the exponential model. And the reason why a lot of economists have used it is to say, OK, here's an exponential model. We know how to use that model. Now we're going to add one parameter and things look very similar. And beta equals 1 essentially collapses that model back to the exponential model. And therefore, people have used it. And that's essentially what we do.

OK. So now the flip side of those kinds of choices are investment goods. Investment goods are essentially very similar. These are goods that have immediate costs but delayed benefits. An example would be go to the gym. Like, you can go to the gym right now. Some people find that fun, but many people also find it costly. There's some effort cost of, say, in this case, 2. There's benefits, health benefits, of 3.

Again, that's really sort of like a stark example. But I'm just trying to make sort of a point to make this simple. Now, again, if we have beta equals 1/2 and delta equals 1, if you think about should you go to the gym today, the answer is no. And the reason is that you put a lot of weight on the present. You put sort of weight of 1 on the present, the cost of minus 1-- sorry. You put a weight of 1 on the cost of minus 2. So you have minus 2 costs.

The benefits are in the future. You discount them by beta. So you have 1/2 times 3, which is smaller than 0. Now, if you think about, then, are you planning to go to the gym next week, the answer is yes. Why is that? Well, now you're essentially discounting everything that's in the future, including the costs of going to the gym. So you can think about, should you go to the gym next week? The answer will be yes. The reason is because the health benefits overall are larger than the costs.

Now, again, what we're going to see then is people deviate from their long-run plans. Now they under-consume investment goods relative to long run plans. So we saw previously that people over consume leisure goods. They do too much fun stuff that has bad consequences in the future. And people tend to do too little of tedious stuff that will yield benefits in the future relative to their long run plans.

Again, you look at sort of timing consistency in a sense of, like, if I plan to go to the gym next week, I'm going to say yes. Now unless things change, once next week arrives, I'm going to be at the choice of, like, should I go to the gym today? I'm going to say no. I would like to not do that because now I'm essentially discounting the future but not the present. Any questions on that? Yes.

AUDIENCE: [INAUDIBLE] one more time why you discount the efforts as in the future and not--

FRANK SCHILBACH: Right. So think of these as different periods-- so the effort-- so if you think about, should you go until the costs are today-- think about these are different periods. Period right now. Health benefits are, like, say, a week later. OK.

In the choice about doing it today, the costs are right now-- that's the present period. So if you go back to the discount function, you put a weight of 1 on anything that's in the present. And anything that's in the future, a future period, you put in a weight of 1/2.

AUDIENCE: OK, so we can see every [INAUDIBLE] in the future because [INAUDIBLE].

FRANK SCHILBACH: Exactly. So the effort right now is in the present. So for like-- so start with a choice of, like, today. The effort is in the present. That's why I'm not sort of discounting that in the first choice for today. But if I think about doing it next week, then next week is also in the future. I'm discounting everything that's in the future, whether it's a week away or two weeks ago. I don't care. And therefore, everything is discounted by 1/2.

OK. So now one important distinction now is demand for commitment, or is there commitment available? So what we're going to do is think about like students who have to do problem sets with very simple beta delta preferences. So consider a student with beta equals 1/2, delta equals 1. There's three periods. And this is sort of getting pretty close to what you're going to do in the problem set, actually, yourself.

There are three periods. You have to do the problem set in exactly one of three periods. That's, of course, contrived and a little stylized. But that's just an example for now. So there's periods T equals 0, T equals 1, T equals 2. The instantaneous utility is of minus 1, three halves, and five halves. So the problem set becomes more and more painful the later we actually do it.

OK. And there's only one day in which you can actually do it. So now the first thing that I think about is suppose there's commitment available. What does that mean? Suppose a student can essentially just pick, when is she going to do the problem set, and actually stick to that. Suppose they could just sort of dictate what she's going to do in the future, what would she choose?

Well-- and then she actually has to sort of stick with that. So now we can sort think about, from the perspective of period zero, remember the person is a hyperbolic discounter with beta equals 1/2. From the perspective of right now, I'm going to say, well, the discounted utility today is minus 1. There's no discounting happening for today if I do it today.

For tomorrow, I'm going to discount the future by 1/2. So I'm going to get minus 3/4. And in the period 2 in 2 days from now, I'm going to do 1/2 minus-- times minus five halves, which is five quarters. So if I could sort of predict-- if I could sort of say what I would like to do, I would say, well, I'm going to do it in [INAUDIBLE] equals 1. I'm going to do it tomorrow if I could follow through with that. Is that clear?

OK. So that's if commitment is available. That is if I could essentially just force myself to do stuff in the future and not deviate from that at all. Now if no commitment is available, how does that thinking change? Yes.

AUDIENCE: T equals 1 [INAUDIBLE].

FRANK SCHILBACH: Yes, exactly. Once T equals 1 arrives, I might actually prefer not to do it in T equals 1 but rather say I'd rather do it in T equals 2. And this is exactly what you're going to see here. This is the exact same problem that I showed you before. Now, suppose the student has no [INAUDIBLE] so the commitment technology is just to say, what if the student can just choose in every period, what is she going to do, would she actually do it in period one?

Well, if period one arrives, if the problem set is not done, if she does it at period one, from the perspective of period one, the discounted costs are minus three halves. That's just because that's in the current period from the perspective of period one. If she instead sort of were to do it in period two, she now discounts, again, period two from the perspective of period one by 1/2.

So 1/2 times minus 5/2 hops is 5/4. Well, 5/4 is less bad than 3/2. So she is going to say, I'm going to do it in period two. So that's to say that students' preferences are dynamically inconsistent. Again, that's sort of something that the exponential discounting model would not predict or cannot explain and the quasi-hyperbolic model is one way of simply creating such timing consistency. Yes.

AUDIENCE: I was wondering how we should think about the fact that, for example, you know that you're going to experience [INAUDIBLE] on the second period. So should then your second period activity function be the one that matters? Shouldn't you, as a sophisticated [INAUDIBLE] minus 3/2 that it's worse from your perspective [INAUDIBLE]. That utility function should be the one that matters and is used [INAUDIBLE].

FRANK SCHILBACH: Right so you're sort of saying-- and sounds if I understood right-- is to say-- and let's see whether that's true. You're saying, well, if no commitment is available, I should never let it get to that point. If I'm in period zero, I should know I'm not going to do it actually in period one. Therefore, I'd rather do it right now to avoid that all sort of like procrastinated from period one to period two. And if I'm not in period 0, I might say I do it rather right now, because otherwise I'm going to just do it in period two.

And from the perspective of period zero, I'd rather do it right now, compared to period two. I would prefer to do it in period one, but that's not feasible because my period one self will not behave in the absence of commitment.

AUDIENCE: I was more thinking in a way that, why would you want to commit your previous one [INAUDIBLE] that your [INAUDIBLE]? Because you have a different utility function in that period. [INAUDIBLE].

FRANK SCHILBACH: It's just a-- I mean, something that's an assumption in the sense of from the perspective of period zero, you make certain choices. You're going to say, I could do it right now. I could do it in period one. Or I could do it in period two.

Now, you have certain preferences. And you're saying maybe your preferences should be different. And that's obviously fine. I sort of specified the preferences for the student. I said that beta is 1/2. So she cares about 1/2 about anything that happens in the future in period one and two compared to period zero.

And then sort of just sort of like-- you just look at what the utilities are and it sort of turns out that from her perspective, from period zero, she prefers period one. Now, you could sort of specify the utility function differently. But for that specific example, I guess that's what sort of the algebra tells us. Yeah.

So now the key question now is to say, OK, now the preferences are dynamically inconsistent. So now what I said-- well, what's really important now-- and in some sense, maybe you were asking about that, as well-- is, well, it's the question on, like, is the student aware of this time of consistency? So you could say, well, when deciding in period zero, you might sort of say, well, I'd like to do it in period one. If I believe that in period one, I'm time consistent, or if I believe that I'm an exponential discounter in the future, I will be very virtuous in the future, well then, I'm very happy to wait until period one. And then I'm going to follow through with my plans.

If instead I know that in period one I'm going to be present bias, well then what's going to happen is then I can't trust myself in period one to actually do it. I know I'm going to procrastinate further until period two. So I'd rather do it at period zero to avoid sort of having to deal with it in period two.

Let me sort of walk you through that. So I guess we need sort of additional parameters to start with which is called beta hat. Beta hat does essentially the sophistication or naivete parameter. Beta hat is your belief about what you think your beta will be in the future. So beta is like what the actual beta is. That's the true value of beta in the present and in the future.

And beta hat is like what you think your beta is in the future, what you believe your true or your beta is going forward. It's not what you believe-- you know your beta today. You know that essentially your present bias today. Everybody understands that well. That's an assumption, as well.

But people do not know or may not know what their beta is in the future. And that's what the parameter of beta hat measures. You're going to have three cases. I'm going to talk about two cases today. There's going to be a third case starting on Tuesday.

And so the two cases that we are discussing is full naivete. That's to say my beta hat equals 1. That's to say I think-- essentially, I know I'm present biased today. But I think I'm an exponential discounter in the future. OK, so that's to say I know I have self control problems right now. But next week surely will be different.

Like this is like in the beginning of the semester, you know you kind of like screwed up in the past. But this semester will be very different. You'll be virtuous and so on. That happens kind of every semester. So that's beta hat equals 1.

So the person does not realize that she will change her mind. She thinks that she's going to follow through with her plan. Her plan is to do it in T equals 1. She thinks she's going to follow through. There will be surprises about present bias because you think you're going to do it in period one. Period one arrives, and surprise, she's present biased. She misestimated what she's going to do and then ends up sort of not doing it and has to do it then in period two.

There's a sort of like false optimism, if you want, about future plans and sort of the attitude of this time is different. From tomorrow onwards, things will be very different. We'll be all virtuous in the future. And then the second part is sort of like-- call it like perfect sophistication. That's bets hat equals beta. That is essentially rational expectation in a sense-- this is a person who perfectly understands their preferences.

I know I have seen this movie before. I know tomorrow I'm going to procrastinate, or I'm going to be present biased. I perfectly understand what I'm going to do. And therefore I take into account that when making choices right now.

And so she understands perfectly that she will change her mind. So she understands that she has plans to do it in period one. But she understands that the period one self has different preferences. The period one self prefers to do it in period two rather than in period one.

So she's kind of like, in some sense, sort of correctly pessimistic about her future self. And sort of then in the case of these kinds of investment goods, that's a good thing, because she doesn't sort of leave stuff out. She doesn't delay things a lot for the future. She doesn't let things get bad because she knows if she procrastinates further, she'd rather do it right now, anticipating that the future self will not follow through.

And here in the perfect sophistication case, there will be no surprises. The person always follows through with their plans. The third case I'm going to talk about next week is when beta hat equals-- is between beta and one, which is we call a partial sophistication or partial naivety. So it's a little more complicated, but in fact, quite similar. I think there was a question. Yeah.

AUDIENCE: [INAUDIBLE]. But the future self actually can [INAUDIBLE].

FRANK SCHILBACH: Sorry, say that again?

AUDIENCE: So the person does not realize that she will change her mind over the course of the semester, but she [INAUDIBLE] her plans.

FRANK SCHILBACH: So you would not sort of-- so you wouldn't really carry-- so if you change your mind in the future, you would not carry through with your plans.

AUDIENCE: [INAUDIBLE].

FRANK SCHILBACH: So either she will change her mind in the future or not. So if she doesn't change her mind in the future, there's no problem, because then there is nothing to be aware. That's kind of like the case of the exponential discounter.

So if beta hat-- sorry, if beta is actually one to start with, which is kind of like exponential discounting, then there's no problem because in the future she will just want the same as in the present. So then sort of by construction or by sort of assumption, then beta hat equals beta equals 1.

And now if she changes her mind in the future about her preferences-- essentially if her beta is smaller than 1, then the question arises, what's the beta hat? Is it like one? That's the perfect sophistication case. Or is it beta, which is-- beta hat equals beta. That's the perfect sophistication case.

Or if beta hat equals 1, that's the full naivete case. And then there's sort of cases in between, where she sort of-- she does understand that the beta is smaller than 1 in the future, but doesn't fully understand that and underestimates-- overestimates her beta, but at least only to some extent. Any other questions?

OK. So now when we sort of look at now, what does a naive student actually do? So what does the naive student do at T equals zero? Well, from above, we sort of know that the self zero prefers to do the problem set at T equals 1. So since she's naive, she thinks she's going to follow through with those plans. So she believes that she will actually do it at T equals 1. So she doesn't do it at T equals 0.

Now, of course, what does she actually do at period T equals 1? We already said that. She does actually not want to do the period in T equals 1, the p set at period T equals 1. So then hence the person sort of surprises herself and actually ends up doing it in period two. Right.

And so what's sort of the summary of that? And this is a more general phenomenon. If you read the paper for class today, that's sort of much more general in that paper and you could do that in a much more complicated way. So essentially, there is the belief in period zero-- she thinks, well, I can not do it right now. I can delay it.

The cost of delaying it will not be that large because I'm going to just do it in period one. But it turns out the cost of delaying is actually much larger, because in period one, she doesn't do it. She'll do it in period two. Now, if you had more and more periods, you could sort of have a whole cascade of doing it every day-- every day you say you're going to do it tomorrow, but you're never actually going to do it.

And so that can lead to actually very large costs of procrastination, because you can essentially delay things infinitely by every time thinking you're going to do it in the future. And so that kind of behavior might persist for a long time. And it's sort of an example, as we call it, like naive procrastination. Every day you think-- or every week you think you're going to do it tomorrow or next week and so on. But you're never going to actually do it.

And then that can lead to really large welfare costs, because there's stuff that's actually not that hard to do. But you just don't do it every day and it gets more and more costly over time. And you never end up doing it. And it's way worse to do it in 50 periods from now. You could have just done it early on if you had understood that you were not going to do it any time soon. Any questions on that? Yeah.

AUDIENCE: [INAUDIBLE] what's the interplay between learning and sophistication [INAUDIBLE]? What type of people [INAUDIBLE].

FRANK SCHILBACH: Yes. So one big question is, how is this even possible that we have here? Here's a person who keeps surprising herself every period. I mean, like, OK, actually, I thought I was going to do it. But surprise, I'm actually not going to do it. But I'll tomorrow I'll do it.

In a way, that's sort of surprising, in the sense of, like, there are these behaviors-- like, for example, when people think about-- when do people go to bed, for example? And people think people systematically under invest in sleep. And in some sense, it's an odd sort of behavior in the sense of, you go to bed every night. You kind of know that you're going to be tired tomorrow.

Yet, every day you sort of have this belief that tomorrow you'll be just fine. And from tomorrow, things will not change. So there are some behaviors that people tend to not learn. And we don't quite understand why that is. In general, sort of like learning about beta hat is sort of-- kind of like at the frontier of research right now.

People have started estimating beta hat. They have found people tend to be fairly naive in various situations. But then the question is, how is it possible that people are so naive in situations where, like-- we're going to the gym. Every week you're going to have the choice of going to the gym. You should have learned by now what your type is, in some sense, what your beta is. And you should have sort of then updated your beliefs accordingly and sort of changed your decisions.

But somehow people have not done that. And people tend to be naive in situations where really they have lots of situations or lots of occasions to learn. Similarly, if you think about problem sets that you're doing, in some ways you know at the beginning of the semester kind of what's going to happen, in some ways. We should now. Yet, people tend to be overoptimistic in various cases.

I was an undergrad advisor for a while. And I was sort of advising students on course choices. Every semester it was the same thing. I would say, like, four classes are grades, and so on. Students would take, like, six or seven. And then I would say, well, you know, that tends to not end up very well. Students are like, no, no, no. This semester is different. I'm going to work really hard. It's going to be all great.

And then a month later, we would meet. And then we would go down to five classes. And I'd be like, well, maybe four classes are also fine. They're like, no, no, no, no. Five classes are great. And then at the end, we end up at four or three.

So in a way, there are a bunch of behaviors that repeat over time. And in some sense, you would still think, why are people not learning? One explanation why people don't learn is in some sense, they don't want to learn. In the sense of like, I like to be a person. I like to think of myself as a person who is virtuous, who is hardworking, who is going to be really doing great at classes and so on.

So I want to be that person. I feel better about myself by wanting that. And therefore, I might not learn as much as I could, even though that's costly for me to do. And we're going to talk a little about this in terms of overconfidence and so on, where essentially people have motivated beliefs and people want to be the person with beta equals 1. But in fact, they might not be.

OK. So now what does the sophisticated student do? We already discussed that. Essentially, the person already knows. The person has rational expectations. So the person knows that if she doesn't do it in period zero, she's not going to do it in period one, because the period one self will change their minds and not follow through with her plans. So she's going to do it in period two.

Now anticipating that for a choice is essentially, effectively, between doing it right now versus in period two. That's the effective choice that she has. And we can now just make the choice. Does she prefer doing it period zero versus period two? Well, she prefers doing it in period zero, because period two is too costly for her, even if that's in the future. And therefore, she does the problem set at period T equals zero.

OK. Any questions? OK. So what's the summary of that? Well, the sophisticated student, for a case of investment goods, the person knows that if she delays, she'll delay even more. So once you start procrastinating, you know this is going to unravel. I'm never going to actually do it. And anticipating that, she's going to say, well, I'd rather do it right now, because I know this is going to unravel or this is going to derail in the future. I'd rather not have to do it very late, late at night, and so on and so forth. So let me do it right now.

So that student in that specific case does better than the naive student. And for the case of investment goods, sophisticated procrastination does not cause large welfare costs. It's not possible for a perfectly sophisticated person to procrastinate things for a long time and sort of lead to a very inefficient or costly outcomes, the reason being if you know that's going to be costly in the future, if it's really that bad, you would rather do it right now.

And so therefore, they cannot be-- there can be some costs of delaying. You might do it tomorrow-- you might do it a little bit later than optimally. But it cannot be that you procrastinated for a long time and never do it and therefore get like really large costs of that.

OK. So one question that also came up in the online forum that you also saw in the readings that we're going to also talk a little bit about in the problem set is, well, is sophistication always good? And it's very intuitive to think that sophistication always helps. After all, knowing things better, having correct beliefs about yourself, that seems to be the correct thing. So intuitively, it feels like that should only help, should only make things better. If I understood my biases better, I should be doing better.

Now, it turns out that's true for investment goods. So for investment goods, these are the goods where you have to do something in the present, but the costs-- sorry, the benefits are in the future. For that, sophistication is always better, or weakly better, the reason being that, again, you can sort of understand your future misbehavior. And by that understanding of future misbehavior, you can sort of avoid bad things happening in the future by doing it right now and earlier.

So in that case, sophistication is always better. Now, it turns out for leisure goods, that's actually not the case. So leisure goods, again, are goods where you can enjoy something in the present and there may be some costs delayed in the future. And so let's look at-- and this is very close to now the paper that you read.

Suppose there's a student who has beta equals 1/2 and delta equals 1, as before. Suppose the student can go to only one movie in exactly one of the four periods. There's only one movie ticket that she could use in period 0, 1, 2, and 3. By the way, Matthew Rabin is a huge Johnny Depp fan, in case you haven't noticed. So he uses Johnny Depp in various of his papers, and so on.

He also had a huge Johnny Depp poster in his office. So the instantaneous utilities are such that they're 1, 3/4, 9/4, and 27/8, so they're sort of increasing over time. So if you-- so there's essentially one day where it's like the most fun to watch the movie, perhaps because you can go with your friends or whatever-- like, it's a better day to do it. Maybe you're a bit better rested or whatever it is. But sort of delaying would be good.

Now, the problem, of course, is that the person is impatient. The person might not want to wait because the person is present biased. He might want to watch it earlier. The question now is, when does she actually go watch that movie?

So now there's a way in which you can sort of write this down in a very simple form, which is a table of discounted utilities. That's a simple tool for solving these kinds of problems. This is supposed to come sequentially, but it isn't today, it's on something that's already filled out for you. So that's easier for you.

So what I have filled out for you here is the instantaneous utility. So the first row here is essentially to say, these are sort of the assumptions that we have. You have the instantaneous utilities-- U0, U1, UT, U3, and so on. So this is essentially what you see on the first row of this table.

Second then I filled out for you sort of the utilities from the perspective of period zero, period one, and period two. So from the perspective of T equals 0, the first period T equals 0 equals just the same. That's just 1, because she puts weight one on that. But then everything in the future is discounted by 1/2, because the beta is 1/2. Yes?

And then from the perspective of T equals 1, of course T equals 0 is sort of irrelevant because that was in the past. By definition, we only care about stuff that happens in the future, because we can't change the past. Notice that's also an assumption, but it's a common assumption made in economics.

So there 3/2 now is in the present. So that gets weight 1. And everything that's in the future gets weight 1/2. So you have 9/8 and 27/16. And now for period 2 again-- period now 0 and 1 are in the past. Now she cares about the present [INAUDIBLE] T equals 2 with the weight of 1. That's 9/4.

And now T equals 3 is in the future. That's discounted by 1/2. OK. And now what I also have done, I've written down the ranking of these different periods from the perspective of-- from the sort of just looking at the instantaneous utilities, three's better than two is better than one is better than zero. From the perspective of T equals 0, you prefer three over two over zero over one. From T equals 1, you prefer three over one over two. And from the perspective of T equals 2, you prefer two over three. Does this all make sense what I've written down here? I'll let you look at this for a second.

OK. So what does the naive person do with beta hat equals 1? Yes.

AUDIENCE: Oh, wait, hold on.

FRANK SCHILBACH: We can do this step by step.

AUDIENCE: OK, so [INAUDIBLE] zero. And then from the perspective of T equals 0, you would most prefer to do something at time three.

FRANK SCHILBACH: Yes.

AUDIENCE: So I guess that means that you don't do it in period one and then move on to T equals 1. And then again you want to do it in period three so you don't do it in period one here. And then you end up at T equals 2. And then you'd rather do it in period two than period three, so you do it in period two.

FRANK SCHILBACH: Yes. Correct. So let me do that in slow motion. So in period T, she makes some plans. You wrote down what is her preferred plan in period T equals zero. We have our ranking at the very right in the right-- in the last column on the right. She prefers to do it in period three. She thinks that she's going to follow through with her plans. So she's not going to do it in period 0 because she thinks she prefers period three. So let's just wait.

So she doesn't go in in periods T equals 0. And period T equals 1 is the same thing. And now she essentially says, well, she prefers three over one over two. It's better to do it in period three so let's just wait until period three. Now, of course, period two comes. Once period two comes, it turns out we cooked up the number such as like prefers a period two over period three. So she surprises herself by actually watching the movie in period two.

OK. Any questions? OK. So now instead, what does a sophisticated student do? So this was the perfectly naive student doing it and watching the movie in period two. No, if you're sophisticated, what do you do? Yes.

AUDIENCE: She starts thinking that if she waits, [INAUDIBLE] equals 2, she's going to at two-- wait.

FRANK SCHILBACH: Correct.

AUDIENCE: So then she knows that at T equals 1, she's not going wait until two. She's going to--

FRANK SCHILBACH: Any help? Yes, OK.

AUDIENCE: So first student, at T equals 2, she would go. So if she were at T equals 1, she knows that if she waits, then she would go at two instead of three. But she likes one. So her choice is not between one and three, but between one and two. So if she were at time period one, then she would go-- choose to go at one.

And then if she's at time T equals 0, her choice is going to be between zero and one, not zero and three. So that if she were at time period zero, she would just end up going at time period zero.

FRANK SCHILBACH: Correct. So let me also say that a little slower again. So if we look at perfectly naive people, they're going to start from the beginning and go forwards. We have a naive person who has plans going forward. And then we just look at how does this plan evolve? So you start at how does this person person decide at period T equals zero, thinking she's going to follow through? Then we go to period one thinking she's going to follow through. Then we go to period two thinking she's going to follow through. And that's how we sort of solve the problem.

Now, for the sophisticated person, you would do things differently. We essentially do this backwards deduction, essentially to say let's start at the very end. The reason that we can use backwards induction is because people have rational expectations. The person knows exactly what she's going to do, which is she's going to start from the end and say, OK, how is this going to play out? What am I going to do in the end?

And sort of anticipating that perfectly, we're going to then solve backwards. OK, so as you just said, in period two, we know that she's going to go, if she hasn't, because she prefers two over three. That's the last row that we have. She likes two better than three once she comes to period two. Now, in period T equals 1, she knows that she's not going to wait until period three. So period three is actually not an option for her. She knows that's not credible. She's not going to follow through.

So then it's just a choice between one and two. She prefers one over two. So she's going to-- going in period one. Now then, in period one two equals 0, now we know-- she realized that she won't wait until two or three. Two or three are not an option for her. She would actually prefer both two and three. But that's not possible because her future self will not stick to her choice.

So then the choice really just becomes between zero and one, and she prefers zero over one. Now, essentially, she goes to that movie, and gets a trilogy of one. But actually, she would prefer any of these other outcomes-- actually, any of her other self would prefer her to not do that. But she actually picks period one.

So that's an example. And there's several examples of that where sophistication can hurt. Now, why is sophistication hurting here? What's the key part, and why is it worse for the student to be sophisticated? In some sense, if you think about it from the perspective of T equals 0, she would prefer a three over two over zero over one.

Or overall, if you look at the instantaneous utilities, period two is better. So from the perspective of T equals 0, period two is better than period zero. Yet, he ends up going in period zero. So a sophisticated person actually would be better off if she were naive. If she could sort of commit to something and say, I would prefer to go in period two, if she could say-- if I only could be naive, that would sort of help her to actually do it in period two.

Of course, you can't sort of choose that. But here she's worse off from being sophisticated. But what's going on here? Like, why is that happening? Yes.

AUDIENCE: So she almost mistrusts her future self and so preempts that by taking the previous decision. But then also preempting that, she would preempt her future self. She would take the previous decision again. And so that makes it so that she does the action at the very start instead. For us, the naive person always blunders through and manages to take the more-- the action with the higher utility.

FRANK SCHILBACH: Right. So in some ways, the problem is set up in a way that, in a way, realistic pessimism-- the sophisticated person is, like, realistic in their pessimism. She knows, essentially, that in the future she's not going to behave. So she would love to wait. But what she needs to do to wait, she needs to have a sufficiently high benefit of waiting.

Now, how high is the benefit of waiting? It depends on what your future self is going to do. If you only wait a little bit, the person will say, if I go-- if the choice is between zero and one, it's not worth waiting. I might as well go right now.

So if I can wrongly sort of have the illusion to myself that I say, well, I'm going to not go in one, or in two, for that matter, then in some sense the benefits of waiting-- the perceived benefits of waiting are higher than they actually are. And in this case, it mitigates the self-control problem or the present bias.

The present bias person wants to do it right now and needs a sufficiently high reward for waiting. So if you can sort of perceive-- if you can sort of deceive yourself into thinking that the rewards of waiting are high, you're going to actually wait. And that's good for you in this case.

If, instead, you're realistic about your future self, you know that waiting is not really helping you very much. You know that essentially, if you wait for one period, you're going to go and go in the next period, anyway. So then the reward-- and that's in part because the numbers are set up that way-- but that's now essentially sort of telling you, well, waiting is not really doing very much, because your future self is also misbehaving. It sort of makes it impossible to actually wait all the way until the very end.

And therefore, it sort of makes the self-control problem, the present bias, worse, because now it's like it's not worth-- why even bother to wait a little bit, because that's literally not worth doing? Really what you would need is wait all the way through. Yeah.

AUDIENCE: So I think this is related to what I was trying to ask before. If you're sophisticated, you know you have [INAUDIBLE]. If you are in period zero, you'll know that on the next period, you're [INAUDIBLE]. So why would you prefer to experience [INAUDIBLE] one now [INAUDIBLE] experience [INAUDIBLE].

FRANK SCHILBACH: No, but I think it's-- in a way, there is a preference over-- and something like that's, again, an assumption how people behave. People care a lot about the present. And that's sort of like, in some sense, in a bunch of experiments and so on, people really want stuff right now, even if they know that they will prefer stuff in two days from now or three days from now and so on.

You just-- you are impatient right now. And that's what the utility function is like. And so you put less weight on stuff that's in the future. Now, if you sort think about your future preferences, there's a question on-- that's a philosophical question in some ways-- how do you think about the fact that your future self wants something different?

That is to say, for example, if I want to exercise a lot and I also like to sit on the couch and watch TV a lot, now my current self wants to exercise a lot in the future. My future self wants to-- and I know that-- wants to sit on the couch. Now the question is which-- and when you think about welfare and so on, which self should we respect?

And one assumption-- and this is an assumption that's built in here-- is to say my current self has certain plans for the future. And those preferences are the preferences that I'm using to maximize. I know that my future self wants other things. But in some sense, I'm assuming I know that I know better. This is the utility function that counts is the current one, not the future one.

And so then I respect that one and I'm trying to maximize that sort of subject to my future preferences being different. But in some sense, you're exactly right. If you think about, like, should we tax potato chips, for example-- talk about welfare and so on-- should we tax potato chips? You could say, well, on the one hand, I might say I would like to tax potato chips, because I know in the future I'm going to eat too many potato chips. That's good for me.

But there's a future self-- actually, myself-- who would love to eat potato chips. That self would be really unhappy. Now you get very tricky questions of welfare and a sense of saying, like, well, who are we to say that the current self is different or should get priority over the future self? The future self will really be unhappy.

So there are different views on this. Some people, including Matthew Rabin including David Laibson and so on, would sort of say, what counts is the current self, or a self that chooses for the future. That's sort of the virtue of self and so on. We should use that for welfare evaluation.

That's an assumption. There's other people who would argue that no, like, there's all sorts of self who want different things. So if you think about, is it good to increase taxes for potato chips? Maybe, because some selves are better off and some selves are worse off. And how do we ever aggregate across those different selves?

So you get very tricky questions. But here for our purposes what we're going to do is we're going to say the current self is essentially disrespecting future utility functions and preferences in the own maximization and the maximization problem. That is to say, I have a utility function right now, and sort of discounting function that tells me how I want to aggregate my future utility over time. That is what I want.

And then I have essentially like a-- sort of like constraints which are coming sort of in the-- when I think about sort of like when you write down a utility maximization problem, there's an objective function, which is my maximization problem. That's the utility function that I want right now.

And then there's constraints in your maximization problem, which are in the future, my preferences will change. So some options might not be feasible because of these preference reversals, and so on. But that's not to say-- I'm respecting your future preferences to the extent that I need to have feasible plans. I'm not respecting them in the sense of, like, I know right now. I know better what's good for myself.

Did that answer your question? OK. Any other-- yeah.

AUDIENCE: [INAUDIBLE].

FRANK SCHILBACH: Sorry, what exactly is the question?

AUDIENCE: So does that mean that if you go see the movie in T equals 2, do you evaluate that as-- sorry, [INAUDIBLE].

FRANK SCHILBACH: So from the perspective of period T equals 2, you evaluate it with 9/4. From the perspective of period T equals 1 or T equals 0, that is in the future. So you multiply the 9/4 times 1/2, which is the beta because it's in the future.

AUDIENCE: So then when you say the student does better [INAUDIBLE].

FRANK SCHILBACH: I'm using sort of from that perspective. So in some sense, that's a bit of a question philosophically, what's better and worse, because you have different perspective. You can take from the perspective of T equals 0-- for example, a sophisticated person would much prefer period two over period zero, one over zero.

So I just showed you that the sophisticated person goes in period zero. But in fact, the person would prefer a three and two. So the outcome of the naive person would be better for the sophisticated person from the perspective of T equals 0.

And so then there's a bit of a question. From the perspective of t equals 2, if you had gone already in the previous period, probably prefer if you hadn't gone previously. But that's a bit silly, in some sense. But you can essentially-- you can write down for each self, what would each self prefer? And I think as it is in this example, it's essentially, the sophisticated essentially does worse for all selves compared to a-- weakly worse compared to the naive person.

OK, but happy to talk after. OK, so I think we have all of that. Sophistication can make things worse. So now overall let me-- so more generally, the lessons are-- so the question is kind of like whether future misbehavior raises or lowers the cost of current misbehavior. So essentially, whether future and current misbehavior are compliments or substitutes-- so if future misbehavior raises the costs of current misbehavior, then sophistication helps in overcoming short run impatience.

That tends to be true for investment goods. That was the example that I showed you first. If I know I'm going to misbehave in the future, I'm not going to like-- I'm going to procrastinate a lot if I don't do it right now. Then sophistication is good because I can sort of prevent that by just doing stuff right now.

So those are the investment goods with immediate cost. There are some exceptions to that. But generally, that tends to be true. Now, if instead future misbehavior lowers the costs of current misbehavior, because essentially the benefits of waiting in this case that I showed you for the movie are just lower now, so the costs of misbehaving now have become lower, then I say, well, I might as well do it right now. I might as well eat the cake right now, or whatever. There's no point in waiting.

So then sophistication hurts in overcoming short run impatience. And that tends to be the case for leisure goods. So avoid immediate rewards. Now-- yes?

AUDIENCE: [INAUDIBLE].

FRANK SCHILBACH: It has to-- it's a little tricky. It's kind of complicated. In some sense, it has to do with specific cases. But the key question is essentially is like, in the future, does sophistication or misbehavior raise or lower the costs of future misbehavior? Does it raise or lower the costs of current misbehavior?

And then depending on whether that's the case, sophistication or naivete go in different directions. So essentially there's a true answer to that. And then essentially if you're naive to this, it sort of flips the direction, depending on the numbers. But it's more complicated to that. It sort of-- these are sort of fairly general statements. So it's a little complicated.

But you'll see some examples again in problem sets and so on. Yes.

AUDIENCE: Does anybody model [INAUDIBLE]?

FRANK SCHILBACH: Not so much. So usually people tend to-- so that's an interesting question. So usually people think beta hat or any preference parameters are often independent of the costs and choices that you see. So the idea often is you are born with a beta and a beta hat and a delta and so on that fell somehow from the sky.

There's some research on how environmental factors and so on affect preferences. For example, we'll show you some research on how contact with different types of people affect your social preferences. For example, if you are exposed to many poor people as a rich person in your life, you might become a nicer person and things like that. I'll show you some research about that.

So there's some research on how exposure to environments affects people's preferences. In the case of time preferences, usually people just think by assumption, somehow you came with a certain discount factor, and that's how it is, or beta delta and so on. That applies to all sorts of settings. Yeah.

AUDIENCE: [INAUDIBLE]. But if it's less costly, you should be able to be more sophisticated. So I mean--

FRANK SCHILBACH: Right. That's an interesting question. So that's not what people have done and thought about. In a way, there's a bit of a question whether your cost and benefit structure can explain your behavior for a given beta hat. But it might also be that for some case-- in some situations, you might say you want to be-- for example, what I was saying earlier, it's like, why are people not learning?

For some behaviors, maybe you actually don't care if your beta is 0.7 or whatever. It's like, fine. But some people really care about being a good and virtuous person. And therefore these might be important choices in life that you would say, I'm a serious person. I do this well. My beta hat-- you'd like to think of yourself as your beta being high.

That's not what people have done very much. So there's a little bit of research of good, specific discounting, like your discounting for certain goods might be different than for others, depending on what kinds of goods they are. So some goods are temptation goods and others are not.

Even that is not really in the mainstream of economics very much, let alone sort of saying your beta hat varies by circumstance. I think that's a very nice idea. I don't know any research that has done that so far. OK.

So now one question you might ask-- well, is it beneficial to be sophisticated? Couldn't we just all be naive and everybody's better off? So in theory-- at least in principle-- it's sort of unclear whether sophistication is good. As I told you-- I showed you just an example-- in some cases it's good and other cases it's bad.

If you think overall many important decisions or choices in life tend to be decisions that involve one term-- one time efforts that yield future benefits, these are investment goods. So in many cases in life, you have sort of investment goods-- like finishing papers, presentations, reports, finding good investment options for retirement, quitting bad habits, quitting smoking, finding a job, and so on and so forth.

So lots of choices are essentially immediate effort, or at one time or several times putting in effort right now where you get future benefits. So a lot of important things in life tend to be investment goods. And in that case, I guess sophistication tends to be good. So essentially it helps you.

Now, in addition, sophisticates take advantage of commitment devices. So I offer you a commitment device that I was discussing earlier. The sophisticated person will say, great. The commitment device might help me. Great. I'm going to demand it and might be better off because of that.

I'm going to show you next week commitment devices don't always work. Maybe sometimes they could actually make things worse. But in principle, at least, you think commitment devices could improve things, and sort of in that sense, being sophisticated is good.

So I think overall, if I had to choose if I'm sophisticated or naive, I would probably choose to be sophisticated, at least in this case. Whether that's actually true, I'm not so sure. There's one somewhat more subtle thing here is to say, like, is it sort of-- is the issue here really impatience or time inconsistency-- in a sense of, is it just about people not caring sufficiently much for the future? Or is it about how important is time inconsistency?

And as I said, some of the behavior that I'm going to show-- that I showed you is, impatience is not really the key issue here. It's not about people not necessarily caring about the future, but rather, people switching their choices and choosing dominated options in the sense of overall dominated. If you sort of commit, there are some options that you prefer that you end up not choosing because you change your choices over time.

Impatience cannot create this type of behavior. What I showed you, where the sophisticated person says, I would like these options but I'm not choosing them because I know my future self will change their minds, that is not possible. You cannot generate this behavior coming from just impatience. A delta being 0.5 or whatever would not generate you this behavior.

What you need is time inconsistency. You need, essentially, preference reversals over time. OK, so then in some ways I think-- perhaps more most practical for your purposes, in part, for the problem set-- how do we solve sort of problems with quasi-hyperbolic discounting?

So as I told you before, for the naive person, we can start from the beginning. And why do we start from the beginning as a naive person? Why does this work? So I sort of showed you already how to solve it. Now I'm sort of generalizing that, in some sense. Why is that a useful thing to do? Yeah.

AUDIENCE: [INAUDIBLE].

FRANK SCHILBACH: Yeah. So the naive person thinks, I'm going to make plans now for the future. And by assumption, the naive person thinks they're going to follow through on those plans. So what you can essentially do is you can start with some plans, see what the person is choosing thinking they're going to follow through. You don't have to worry about what's going to happen in the future, because you can explain perfectly what the person is doing in the present by making plans for-- sorry, by making your choice for the future periods.

You don't have to actually worry about whether the person is going to follow through, because the person doesn't even think about that. So you say, I make a plan for the future, assuming I'm going to stick through with that, or stick to this. And then you can just sort of roll forwards and say, OK, I start with period T equals 0, make the plan for that, think I'm going to follow through.

Surprise, maybe period one I'm going to do something different. But then I'm going to have the same assumption, and so on. I can just go from T equals 0 going forward under the assumption that I'm going to stick to those plans. For the sophisticated person, that doesn't work, because a sophisticated person knows that if I make some plans for the future, those might not be feasible because the future self will misbehave and change their plans.

I know essentially that some plans are not feasible. So what I need to do is essentially just start from the end and see, OK, if certain options are available, what I'm going to do, and then walk backwards to solve it. And that's essentially rational expectations where essentially you perfectly understand your future behavior and sort of take into account your future choices. You're going to make them current choices at the end.

What about the exponential person? How does that person solve the problem?

AUDIENCE: I just had a question about what you just said.

FRANK SCHILBACH: Go ahead.

AUDIENCE: So to summarize a sophisticated person, where they're going from the end to the beginning, as they move backwards, are you just removing choices that you know you won't ever--

FRANK SCHILBACH: Correct. Exactly. So what we did in the movie choice, if you go through this example in more detail, you'll notice that essentially, if I know that, for example, suppose I want to go in period three but I know I'm going to go in period two if it comes down to that, then I'm going to remove period three for my choice set or for my options when optimizing in period one or period zero, anticipating what I'm going to do in the future.

So what about the exponential discounter? Yes?

AUDIENCE: [INAUDIBLE].

FRANK SCHILBACH: Yes. So there's two options, in fact, for the exponential discounter. It actually doesn't matter because you follow through with your plans anyway. So you could start from the beginning and just walk forward. In fact, what you can do is you can start in period zero make a plan for everything and just follow through. You don't even have to go to period 1, 2, and 3 because you know you're going to stick to that plan.

So you can essentially just do that. You could actually also do backward transactions. And I'm pretty sure you get the exact same answer. It doesn't matter. The much easier thing to do, of course, is to solve the problem in period 0 and then just assuming you're going to follow through.

So sorry, this should say readings for Tuesday. The class is on Tuesday, not on Monday. Please read the paper by Ariely and Wortenbroch. Read the entire article. We're going to discuss applications from smoking to drinking, setting deadlines, and so on, and apply these models in real-world settings. Thank you very much.