[SQUEAKING]

[CLICKING]

[RUSTLING]

[CLICKING]

FRANK SCHILBACH: All right. I'm going to get started. This is lecture 19, of 14.13. I'm going to finish up some of what we discussed last time on gender discrimination and identity And then we're going to talk about nudges and default effects.

So let me sort of fairly quickly go through the remaining slides of lecture 18, and then move to the lecture slides. So remember, we talked last time about the study by Bertrand et al. about gender identity norms, where we talked about because-- so the authors argue there's this norm that some men perceive-- and perhaps even some women-- that men should be earning more than their wives.

There's this cliff-- this missing mass of households or couples in the US. We talked about different reasons for that. Missing couples that don't form in the first place. Such marriages and couples are less happy and stable, so they end up in divorce. And there's also, even for existing couples that form, whether the wife or the woman has more earnings potential. In some cases, at least, women tend to either work less few hours, and therefore perhaps that was promoted and the like.

And in addition, they might be involved in more nonmarket and childcare work. It's called the second shift. They have to do more work because the husband is very unhappy about her earning more. Which then, of course, holds potentially the wife back at her work.

Now, somebody was asking last time-- I believe Justin or others-- about the arrival of children and the gender gap. So there's quite a few studies like this one. So out of this here, which is the study by Kleven et al. in Denmark. And these studies essentially show there is a substantial what they call child penalty for women, which is when you look at men's and women's earnings.

On the left-hand side, there's a figure shows earnings. This is yearly data over time. On the right-hand side, you see hours worked. Similarly, you see productivity. But essentially, you tend to see is that for men-- this is the gray bar here on top-- you see essentially before the child is born, you match men and women to have sort of similar earnings, the trajectory would look very similar.

And then, for men, perhaps there's a slight decline or no further for the increase over time. But for women, there's a clear drop in earnings. And that drop persists for many, many years. So it's not just like there's a year after birth, there's lower earnings for, on average, women. But there's a persistent gap overall that has to be presumably with some women never returning back to work, some women remaining part time, just some women working a few hours when they have children.

Notice that that's not just about hours worked. So when you look at the earnings gap, it's about something like 20%. The long-run child penalty here is like 20%. It's also, for hours worked, it's something like 10%, which must mean that there's an additional 10% coming from productivity or wages conditional on hours worked. So it's not only that women are working fewer hours, but they're also paid less per hour, which presumably is once people start working few hours, take more breaks, or just take a break for a few years, that they're getting less promotions and so on. And that sort of carries through for a long time.

Now, interestingly, some of these child penalties are transmitted through generations-- meaning from parents to daughters-- suggesting that essentially, in the same family where there is child penalties from the parents-- as in the woman that is a childcare penalty for the mother, the daughter will also be likely to have a child penalty later on. Like, 20, 30 years once she's been actually working. So there seems to be something to do with the child environment and gender identity, which is the norm that the woman is supposed to take care of the children, versus the husband is not. And that seems to be sort of permitted through generations in the child environment.

Now, you might sort of say, well, then, couldn't we just do some policies that support women or support parents to help them deal with children better? And wouldn't that help? And there's some policies that people have done. In particular, in academia, you might say, well, if a woman has a child or a couple as a child in general, that's surely a hint to productivity. And so therefore, one could say, well, that's [INAUDIBLE].

That's in fact what people have been doing, is sort of, let's stop people's tenure clock which essentially is an assistant professor. People have something like 7, 8, 9 years to do this research until tenure. You might say, let's give people another year if they have a child.

Now, one proposal is to have gender neutral clock stopping policies. Now, these policies are intended to involve these gender neutral policies, meaning that both parents, regardless of who's giving-- or not just the person who's giving birth, but also the other person who's taking care of the child potentially is getting an additional year on the tenure clock. So the intention is here in particular to involve men more in child care.

But you know, those kinds of policies tend to be not enforced, and potentially they might even enhance gender inequality if it's the case that both men and women get another year. But the man essentially takes that as a year off where you can do additional work while the woman is back taking care of the household and the child. Then, such policies in some ways might make things worse. And this is exactly what some studies find, is that these gender neutral policies, in particular, in economics.

There's also some examples from other settings. But essentially, introducing such policies substantially reduce female tenure rates while increasing male tenure rates. So here's a policy that looks actually a good policy. It's gender neutral and so on. But in fact, it might not be, depending on who takes advantage of that and who doesn't.

So it's a very tricky question because some of these things are not enforceable. If instead you gave only these tenure exemptions to women, then you would get issues in households, that then women are sort of pushed out of work, which in part I think some of these neutral tenure clocks try to avoid in the first place. So it's not clear what to do here. But I think it's important to understand that some policies that look like they are in fact neutral or gender neutral, they might in fact not be.

Now, let me tell you what I [got about one more paper related to gender and the labor market or identity, which is very nice paper by Bursztyn et al., which asks the question whether women avoid career enhancing actions because these actions signal undesirable traits. So in particular, women might be-- ambitious women might be perceived as not attractive partners in the labor market, in particular, in some settings such as among MBA students.

And so, in particular, as you might be familiar with, there's lots of dating happening among MBA students, where people try to essentially find partners or somebody to get married to. And it appears to be the case that being ambitious or being perhaps overly ambitious is perceived as negatively as an undesirable trait on the dating market for MBAs.

Now, what Bursztyn et al. find-- and this is joint work with Amanda Pallais and Thomas Fujiwara they find that unmarried and unmarried female MBA students perform similarly when the performance is unobserved by classmates. So these are things like exams, problem sets. These are things that only the professor sees, but not necessarily their classmates. To here, male and-- sorry. Unmarried and married female MBA students do similarly well when we compare them. But unmarried women have lower participation rates.

So these are essentially things that, do you say something in class? And as you might know, in MBA classes, in some, at least, there's somebody sitting in the back actually recording each interaction, each participation very carefully so you know exactly who says how many things and how often, and so on. So it seems to me that unmarried women tend to not want to look too ambitious and too hard-working in class.

Now, that's of course only suggestive. So then, they also have a field experiment where they ask people about their desired salaries and willingness to travel and work long hours in real stakes placement questioning. So this essentially, at the end, there's this placement, there's these career fairs. But people ask these questions about, what salary are you expecting? Are you willing to travel really far? Are you willing to work long hours? And so on. And of course, these questions all signal ambition.

And then, the really interesting part of the experiment is, then, they varied whether this information is perceived to be private. So you just saying it to a recruiter but nobody else in your class, or anybody else might find out, versus whether there some expectation of some of their classmates actually learning about their preferences. And what they find, then, is that the single female students report lower desired salary, and willingness to travel, and work hours-- and work long hours when they think that there's an expectation that their classmates might find out about it. Which essentially is to say, they perceive it to be costly for their class, or their classmates find out about their ambition.

Now, importantly, two other things are also true. One is for other groups. So, other women that are not single or for men. Essentially, these responses are unaffected by peer observability. And second, it seems to be that what's driving the observability result is the observability by single male peers. So when other female students or other non-single male peers are potentially seeing these answers, there's no effect. But the effect's really coming from other single male peers potentially seeing it, which is precisely what you expect if you think having ambition is potentially a negative signal on the dating market.

So that's another sort of study showing that these norms or perceived norms are really important, not just in terms of punishing women and in some ways undeservedly for mistakes as we saw in Heather Sarson's paper, but also by holding women back by themselves. These are single women who are highly qualified MBA students who are deliberately not as ambitious and hard-working as perhaps you might think they would like to be-- which is what the private responses suggest-- because there's a penalty waiting for them, or receive penalty coming on the dating market.

Any questions about this? So here's a question about, can it be the opposite? Better ask this in person. It's a little bit hard for me to follow the chat. But can it be that more confident women tend to be married? Yes. In principle, yes. But I think what we see here in the experiment is, notice that what I was saying, I was careful to say these are about there's a perceived penalty of being overly ambitious. There may well be for some women, being very ambitious and so on. That's a very highly desirable trait. But it seems to be that the perception, at least, is that being highly ambitious and so on is not desired.

So I think there's another question by Lucy about married women compensating for discrimination. I think that's all. There might be other stuff going on. But I think the key part here is the variation between private and public among single female students. So when they are asking in private, they are saying essentially they want higher salaries. When they're asking in public for higher salaries and they're willing to work long hours, and so on.

When they're answering in public, they're essentially holding back. Which seems to suggest-- and that's consistent, I guess, also with the participation rates. And that's consistent with why there is at least a perception of lower desirability, in particular among their single male peers, which is what this results is driven by. It's not clear, by the way, that the study shows that that's actually true. That's about these women's perceptions.

The study also doesn't necessarily show that this translates at the end of the day in lower earnings, and so on, and so forth. The authors argue that that's the case. But these are only sort of questions-- real stakes questions and placement questionnaires. It's just sort of essentially more like a proof of concept arguing that this is an important force that might be at play, that women hold themselves back in some ways here and in the case of a placement questionnaire.

I had started this already last time, only briefly going through the Vesterlund study, which essentially is saying that there's two forces going on. There's demand [INAUDIBLE] in terms of tasks that women and men are perhaps asked differently. And beliefs could be quite important here, and mediating who has which tasks and when. So let me show you a little bit of what that study is about.

The study is essentially asking the question whether women are not saying no often enough In particular when it comes to work. And it's motivated by the fact those female faculty members tend to do a lot more what they call non-promotable tasks. These are tasks that essentially everybody wants somebody to do the task, but they don't want to do the tasks themselves. So that's like essentially University committees, and here, also undergraduate teaching, or undergraduate student advising.

Of course, everybody loves undergrads and wants to advise them. However, it's not something that's good for peoples' career. So at the end of the day, faculty are often evaluated by research output. So any time that's spent on committees, et cetera, these are not going to be helping for promotions eventually. And then, one first observation is that [? women ?] do that much more than men.

Now you might sort of ask, why is it that women spend their time differently? There's potentially issues of demand. Women are just asked a lot more to do these things. And there's potentially supply issues conditional on being asked. Women might be just much more likely to say yes. Now, what this paper argues is that demand and supply interact. Because women are more likely to say yes, they're also more likely to be asked in the first place.

OK. So why do we care? I think I already mentioned this. There's sort of a few broad perspectives. One is about, we kind of try to understand individual decision-making. and we want to understand how people make optimal decisions. Here's another reason where people report. Here is an example where people potentially mis-optimize. There's a managerial social planner perspective that if you want to sort of organize or you want organizations to be efficient, you want to make sure that everybody does equal work in some way. And so you don't want women who are very productive to be stuck in a bunch of other tasks that hold them back.

And then, there's a public policy perspective. It might essentially be that sex differences and the allocation of time. So women spending more time on these kind of tasks might explain why they're less likely to be promoted and so on. And then, understanding that and understanding these issues might sort of help with interventions to improve equity.

Now, OK. So what are these tasks? Again there's promotable tasks. That's doing research. And there's non-promotable tasks. So promotable tasks, meaning essentially, if you do these tasks, you're likely to be promoted or essentially get tenure eventually. A non-promotable task essentially are tasks that many people could do the task, but everybody wants somebody to do it. But often, nobody actually wants to do it themselves.

Now, the authors that first [INAUDIBLE] field study where they have faculty at a large public University where you send essentially emails from a chair, and they're requested to volunteer to join one of several university-wide faculty senate committees. And women are essentially, in these tasks, are way more likely to volunteer conditional on being asked. So that's the first piece of evidence that conditional on being asked women are more likely to be asked.

Then you might say, well, perhaps there's some personality characteristic [? it's ?] [? on. ?] Maybe women are nicer. Maybe women care more about the public good. Now, what the authors show is that that may well be true in some cases. But in other cases, it's also about expectation about who does it and who doesn't. Moreover, the ask about who is asked is quite different depending on whether you're male or female. And that's where the experiment comes in.

So what they do in an experiment is what's called the threshold public goods game. It's very simple, which is, a small group needs to find a volunteer for a task. These are groups of 3 people, and there's like 10 rounds where people are matched over time with each other. And so it's essentially set up to mimic real world situations where everyone prefers the task to be undertaken by someone, but not themselves. As you don't want to do it, you have to have somebody else to do it. You want that somebody does it, but you just don't want to do it yourself.

And so people get-- the group of 3 gets two minutes to decide whether to invest. And investing means essentially you have to [? truly ?] press a button here. Only one person can invest, and the round ends when someone invests. If no one invests, all the members get $1. If you are investing you get $1.25 and the remaining members get $2.

So if you're thinking about this task, what are you thinking about is now, A, is somebody else going to invest? And B, like conditional on your beliefs about, is somebody else going to invest any time soon? Do you want to invest or not? Clearly, it's the efficient thing that somebody invests. But you don't want to be yourself because then you get only $1.25 rather than the $2. And the clock ticks down until one person invests, or no investment is made 2 minutes.

OK. So now what do we find? Or what did they find in this experiment? Women are way more likely to invest in each of these rounds. This is sort of the red line here. The blue line here is for men. And so women are much more likely to invest than men. And that's very clear when you look at these figures. And that's also very persistent over time. So it's not just like, perhaps in the first round, and so on. This is even after 10 rounds. When people keep being matched with others over time, women are much more likely to invest.

Now, how do we interpret this? So how do we think about this? What can we learn about this? One explanation that I already mentioned is, women are much nicer. But can you infer this from this task? You observe if women are a lot nicer than men, well, then you should find that women are more cooperative, as Lucy suggests. If women are a lot more cooperative, you should see that in women-only groups, women are much more likely to invest than in male-only groups.

And that's exactly what they tried to do in the experiment. So they have essentially single sex sessions would have all-female and all-male groups. And what you find is, while the probability of investing is actually exactly the same for men and women, they also find that when trying to predict who's investing who is not, any measures of social preferences, risk preferences and so on that they've [? listed ?] previously actually not predictive of this behavior. That sort of suggests it's really not about altruism in a sense of being nicer or wanting.

So in some sense, investing is doing something nice for the others in your group. It doesn't seem to be that it's about that. It's also not so much about risk aversion. So you might say, well, I'm investing because I really like certainty. If I invest, I know exactly what I'm getting. If I'm not investing, I have a chance of like getting the $2 or essentially a lottery between $2 and $1, depending on, then, the probability is essentially about what you think is the probability of others investing. So your risk aversion might matter, but they also show that risk aversion is really not important.

So now they have essentially two experiments. One is the experiment, 1 that I showed you. Women are more likely to invest in mixed sex groups, and women and men are equally likely to invest in single sex groups. And so now, then, what they're doing is, then, they're trying to say, well, could it be that beliefs that women will invest are important here?

So that's essentially the idea. If you think women are more likely to invest than men, if women think this and men think this, then essentially you're going to get in a situation where if you have a mixed sex group, then you will have one woman saying to men, the woman thinks, oh, the men are not going to invest anyway. So either it's like we're investing or getting $1. So then, for the women, it's optimal to invest. And the men are going to think, well, the woman is going to invest. So why should they invest if the woman is investing anyway? For them, it's then optimal not to invest.

Now, experiment 3, then, get set this in some ways, which essentially is about now getting women and men to ask to invest. What they do is, essentially, they have what's called a-- they call it a photo ask, which they have 4 people per group. 3 people as 4 can invest. That's essentially the same as before. They call them green players. Not sure why, but they essentially [? do ?] [? their ?] green players.

And one person is unable to invest, but he can ask one of the three others to invest. It's not a binding request. Essentially, it's sort of a nudge. So if you want to, instead of essentially just suggest that they may exert some social pressure, that's the red player. And then, the red player is essentially incentivized. The red player wants the investment to happen.

So this is kind of what this looks like, is essentially three people that you see as the red player. And you can say, whom would you ask? Well, any of these three. You see the photos, so obviously you see the gender of the person. And so, what they then find is that women are much more likely to be asked. You can see here, the red lines are more to the right, or there's more red to the right, which is the total times asked to invest. So there's more asked to the right for women than for men.

Now, what did we learn from that? Well, absent a request, investment rates do not differ by gender. But when asked to invest, women are more likely to do so. So there's two things going on here. One is, when asked, women are more responsive. And [INAUDIBLE] says that the marginal increase of being asked is higher for women than for men.

And then, in addition, women are also more likely to be asked. Which is to say, when people have the choice about whom should you ask, well, they kind of know who is the person who is more likely to react, which is the woman in the group. So they correctly understand that if they ask the woman, that it's going to have more of an impact than asking men. And that's true. So when women are asked, they're also more likely to comply.

So now, that essentially sort of says these gender differences get amplified by increases for women to contribute. So not only is it the case that they're more likely to say yes, but on top of that, they also will get more requests. Which of course, then, is like, they do so much work overall.

And so, what's key here is about this. it's not about preferences. It's essentially just an equilibrium that at the end of the day will be sustained. Because again, if men think women are going to do it anyway, and if women think, well, nobody's going to do it unless I myself, [? when ?] [? I ?] [? am ?] in situations where it might be optimal for both women and men to do it, conditional on their [INAUDIBLE]. Any questions on this study?

OK, so then, let me just summarize a bit. So we saw large gender wage and earnings gaps that even persisted even until now. Substantial progress has been made due to the technological advances and other improvements, but sort of those persistent gender differences in many societies, including the US. Biased beliefs and identity concerns appear to play a particularly important role, in particular, in some issues such as educational differences and so on.

And given the technological advances, some of these technological issues are less important now. But really, it seems to be that biased beliefs and identity concerns are at the center of explaining things in addition to the child penalty. But even the child penalty, you might argue, is driven by beliefs, and identity, and so on. And then there are these feedback mechanisms between demand and supplies of tasks that women and men are doing. And in particular, it seems to be that, again, beliefs are important in the sense of, like, it matters what people think other people do. And then that seems to create, essentially, these inequalities across gender.

Now, if you want to sort of mitigate, or reduce, or eliminate the gender gap, it's really important to understand these issues. I think we have made quite a bit of progress in understanding what's going on, and identifying the issue. Now, I think the next steps are to try and understand what are policies that are in fact successful in helping us do that. So this lecture is about defaults, nudges, and frames. And originally, I had planned to also do mental accounting. But given that we didn't finish last time, I moved this to a later lecture, at least [INAUDIBLE].

OK, so where are we in the course? We talk a lot about preferences. We talked a lot about beliefs and the interaction of those, and sort of potential deviations from the classical model. Now, we're going to talk about non-standard decision-making. And one way to think about this is to say, even if you understand people's preferences very well, then you can understand these are non-standard preferences in some ways, and potentially biased beliefs.

Even conditional on that, people might behave in some ways where they deviate from the neoclassical model. That is hard to explain just by deviations and preferences [INAUDIBLE]. So there are additional deviations from the neoclassical model that are beyond changes in preferences and beliefs. In particular, I'm going to talk now about frames, defaults, and nudges. Mental accounting is a similar case of that. We'll talk about this later.

OK, and so the most famous paper in this space, or one of the most famous paper [INAUDIBLE] is about 401(k) savings. So companies offer these 401(k) savings. That's right. So it's in fact the most common voluntary savings vehicle in the US. What you do is you set aside some money for retirement, and you choose what's called the contribution rate. It's like what fraction of the income you want to contribute, and you choose the asset allocation.

That is offered by companies-- by employers. Now, why are employers doing this? Because employers want to-- it's seen as sort of a nice thing that employers can do for their employees. It's potentially a way in which you can create some form of solidarity with your employees. In a way, you want your workers to do well in retirement beyond their working for you, and so on.

And then, sort of overall, what companies often do is they offer matching contributions up to a threshold. This often incentivizes that, in fact, fairly strongly. There's also a penalty for early withdrawal. Why is there a penalty for early withdrawal? Yeah, exactly. It's like a soft commitment device in some ways, in the sense of-- or it's a hard commitment device in the sense that it's dominated.

But essentially, it's like you want-- it's sort of a mixture. You could have if like a hard commitment, or a fully hard commitment, or very hard commitment device. But you could say, you cannot withdraw this until age 60, or something, or 55, or 65, whenever you retire. That's not what it is. Instead, there's a penalty that's often like 10%. And the penalty is there because you want to get people-- avoid that people-- what's called leakage, that people essentially withdraw money early, presumably because of self-control problems, present bias, and so on. There's lots of leakage, in fact, even though there's a 10% withdrawal.

There's a 10% withdrawal as opposed to shutting it down entirely because sometimes people have shocks. They have health problems and so on. They need a bunch of money. And in those cases, of course, you want to allow this, or the company wants to allow that to people. So a 10% penalty tends to be sort of compromise between that.

There's tax deferral, which is you pay, usually, a lower marginal tax rate during retirement, right? When people are 30, 40 they earn a lot of money compared to when they're 70 during retirement. So people can pay taxes when they receive the money at retirement, so to the extent that your marginal tax rate is lower then, you're going to save some money from that. So that's another reason why you might want to use your 401(k) savings.

Now, what are some patterns 401(k) investments? 2/3 of employees believe that they're saving too little when you just ask them. And 1/4 of these intends to raise their savings in the next two months. So that's sort of the typical thing that sort of smells like self-control problems, where you ask people, like, are you happy with your savings? Should you be saving more or less? Then people tend to say, oh, no, no, I want to really save more. Next two months, I'm going to do it. This is sort of the same as when you ask people about the gym. I [INAUDIBLE] exercise more every month. But it happens.

Almost nobody follows through. In fact, the reported undersavers have low savings rates. You find similar patterns also in other surveys. If you're interested in more of the Choi et al survey, it does a very nice job at summarizing these patterns.

Now, given these patterns, what are standard economics tools to increase savings? Suppose you are the social planner, or a company want to increase people's savings. What would standard economics 1401, 14-A, or whatever tell you what should we do to increase people's savings? How would you do that?

First, the matching contribution is very much like an incentive, so, like, a very standard, straightforward incentive, Econ 101 would tell you people react to incentives. So let's give people incentives. We provide matching contributions, and therefore people-- essentially increases the rate of return on savings. And these matching contributions-- I'll show you in a second-- tend to be very high, or fairly high. And so people should be doing it more if you make it more attractive [INAUDIBLE] under usual circumstances. What else could we do?

And maybe that's not necessarily standard economics. Of course, one option will be, like, the default. But what other options could have or we do? What else could we do [INAUDIBLE] people? Right? So you could-- so the matching contribution is a way of doing that, essentially because I'm telling you I'm matching however much you're contributing by adding so much money in your account, and you're not supposed to withdraw it. So that's implicitly giving you a return.

Another thing I think that you're already saying is what I have here as well, which is provide additional choices or better choices in some ways, or improving what's called our financial education. That's to say, maybe people are not investing in the right thing. And so maybe I could either give them better investment funds or better options that have higher returns. And by doing so, people will be inclined to save more. Or maybe I can explain to them better, for example, compound interest and so on, and sort of [INAUDIBLE] that people don't quite understand, perhaps because of lower education.

Now, that's not necessarily standard economics in the sense of, like, people are confused. That's sort of not what standard economics would say. But it might well be that people are-- if you explain to them better what is compound interest, and if you can save now 5% of your income, how much will that be in 30 years from now, people are often surprised because that's very high. You could sort of get them-- provide them with additional choices, maybe make it more attractive for them.

You could even sort of subsidize it and say, here's a fund, and I'm giving you more money. Notice that that would be quite expensive for the company to do over time. So companies tend to not do that.

Now, it turns out, none of these tools tend to be very effective. So people have tried this, despite the fact that-- they have been matching contributions despite the effect of financial education, sort of information sessions, providing more choices, and so on. Notice that it actually sometimes backfires, because people get confused. There's too many choices, or what's called choice overload. So these things tend to not work particularly well.

Now, one question again, why are companies actually trying to increase people's savings? So in some sense, you might say, well, we're offering-- here's your 401(k) savings. We'll offer it to you, but whoever wants to actually do it can use it. And otherwise, we really don't care that much. So why are companies still interested in doing this? Why are people paying this match? We had some answer already before, but that's actually not quite the answer why this is [? going. ?]

So to be clear, there's two questions. One question that we talked about a little bit previously was, like, why might the company offer savings options anyway? And the second one is conditional on offering that-- why is the company really interested in increasing take-up? Because they say, well, whatever. We offer it. Some people like it, some people don't. If they don't like it, why do we care?

So I always thought, you know-- when I first read about this, I thought companies are nice and they want to sort of be good to their employers, and they want the best for their workers. And you're already a lot more cynical in a sense of, like, you're kind of saying, how do we get workers to work a lot more? It could be. I mean, I haven't heard of what you were saying.

But I think in principle, that could be a reason where-- in a way, I think companies like workers to be happy. And to the extent that if they consume more right now, that probably makes them happier and perhaps more productive, so then, in some ways, probably you have the incentives to have people not save too much. Because in some ways, you don't care that much about whether somebody who is 60 is particularly happy or has money.

Because essentially, you can't really engage or you-- age discrimination is illegal. So you can't fire, like, a 57-year-old person who is doing their job reasonably well, but maybe is not as productive as you'd like them to be. And so to them, if that person has saved a lot of money, they probably are more likely to retire.

You guys are more cynical than I am. So the reported reason for a lot of these policies is, in fact, the reason that there is IRS nondiscrimination rules. What does that mean? Essentially, by the IRS rules, it cannot be that people who are paid a lot can benefit from certain retirement and other schemes way more than people who are paid very little. Now it turns out that low income or low wage workers are much less likely to save in 401(k) and other savings schemes. Therefore, they're much less likely to receive any of these 401(k) matching and other contributions.

And therefore, essentially, their companies tend to have way more money spent on sort of high paid employees than low paid employees. And there are some nondiscrimination rules by the IRS that says that's illegal. You can't do that. And therefore, companies are very keen on getting lower-- if you are like a CEO or something, you want your workers to earn very little to invest in 401(k) and other types of schemes that they get rewarded for, because that allows them, essentially, to pay the high paid employees higher bonuses, in these kinds of schemes as well.

That's essentially one of the reasons why companies are very, very interested in increasing their 401(k) savings, in particular among workers with low wages. And that's exactly what the default to default policy's tend to do. Chad, did you have another question, or was that just your hand that's still up?

OK, so let me move on. So why would you participate in 401(k) savings schemes? We already talked about this. What are the costs of nonparticipation? There are essentially three main reasons from the worker's perspective. There's foregone tax benefits. And tax benefits are, again, your marginal tax rate, usually while people are working, tends to be higher than the marginal tax rate when people are retiring. So essentially, by investing in 401(k) savings, they can defer taxes and save some taxes.

Second, people forgo the employer match. I'll tell you about this in the second. And number three is-- that's a little more subtle, which is foregone consumption smoothing, which, essentially, is your marginal-- the marginal value of a dollar for you, while you're working at age 30 or 40, tends to be lower than the marginal value of a dollar once you're 70 or 80. The reason being that, essentially, you're richer. You have more money available when you're 30 or 40, so you actually want to smooth consumption to later in time.

Of course, there's a discounting motive going the other way. But it shouldn't be that, essentially, at age 59 or 64, you earn a lot of money and consume a lot, and then the year after, once you retire, your consumption drops a lot. Economists would say you should smooth this out by, essentially, saving some money to the future. And 401(k) savings [INAUDIBLE] might help you with that.

Now, why do companies care? I already said that. Essentially-- whoops, sorry, going back. Non-highly compensated employees don't save enough, or don't save very much, in particular in new schemes. And the IRS nondiscrimination tests-- has essentially nondiscrimination tests in pension plans that precisely try to prevent that. So really, what you want to do is you want to get your low income workers to save.

OK, so now here comes the paper by Madrian and Shea. That's a very successful and one of the most successful papers in behavioral economics. The reason being that, A, standard economics has sort of failed in some ways, in the sense of saying, like, we tried matching contributions, we tried giving people education or information, and so on, and so forth, and none of these things have worked very well. And then instead, something that shouldn't matter-- which is what I'm going to show you in a second-- the defaults mattered a lot, and had huge effects, which is very hard for standard economics to explain.

And then on top of that, this is in a domain that we really care about a lot. Savings is one of the most important choices that economists have considered. And so that's a very good example of an early example of an aspect of the economics, in this case default, that really matters in some ways that's hard to explain for standard economics. And the standard tools of economics seem to not have been working very well.

OK, so what does the paper actually do? It's a large publicly-traded, fortune 500 healthcare company. You can enroll in 401(k) savings plans any day by essentially filling out an enrollment form, or calling the 401(k) record keeper. What does that mean? Essentially, you call somebody and say, I want you to invest in 401(k) savings when you're in this company. And you're going to say, I'm going to-- I want to save, say, 5% of my earnings. I want to save every month. Or you could fill out a form.

Now, these days, a lot of this is online. So if at MIT, you want to change your 401(k) savings, you essentially have to fill out an online form, and change your online allocation [? potential. ?] So there's potentially small transaction costs involved. But notice that the transaction cost tends to be very minor, in a sense, like, it's going to take you maybe 10 minutes, maybe 20 minutes-- maybe if you think about it for a while, it takes you an hour-- but really, it's a very minor cost of time overall that this costs you compared to the potentially huge benefits of employer matches and so on.

Now, what are these employer matches? In this particular company, it's a 50% matching contribution for the first 6%. So again, what are these percentages? The percentages are essentially the fraction of your income that you'd like to contribute to these savings. For example, 5% or 3%, or 7% is what you're saving. And now the company says they're going to match 50% for the first 6%.

For example, if an employee chooses 4%, the company pays an additional 2% into that savings account right away. So that's actually quite a bit of money. If an employee chooses 10%, the company pays an additional 3%, because that employee is above the cap of 6%. So you cannot get more than 3% from the company.

Other companies these days have even, like, 100% matching contribution for 5% or the like. That tends to be fairly common. So for example, MIT has 100% matching contribution for the first, I think, 5%. In this specific company, the employees are first eligible after one year of employment before the change happened of the policy that I'm going to talk about in a second.

Now, notice that this is a terrible policy from a behavioral economics perspective, in terms of if you wanted people to sign up and go do this, that's not what you would do. If I'm telling you, like, here's a thing that I want you to sign up for, in a year from now please come back and do it, well, in a year from now, people might forget. People might procrastinate. People might be busy with other things. So it's not a policy that's great to do to start with to get people involved. And in some sense, there's no wonder that participation early on was kind of low.

Now, the company does that because turnover tends to be high. And you try to avoid paying people early on that might leave anyway. And you have to deal with the paperwork, and so on and so forth. Notice that that might be exactly the wrong thing to do, because once people have paid into their 401(k) savings, maybe that creates some loyalty and makes workers to stay more at the company anyway. But for what it's worth, in this particular case, employees were only eligible and one year after employment.

Now, then there's a discontinuity of the 401(k) plan default based on the date of hire. So here's the old cohort-- the window cohort, as I call it-- and the new cohort. So the old cohort are people who are hired between 1st of April, '96 and 31st of March, '97. So they are essentially people who have been previously hired. This is what I just told you. You can, after one year-- after being hired after one year, you can-- you're not automatically enrolled, but you can essentially just call somebody and say, I would like to start saving in my 401(k) savings. Or you could fill out this form.

Then, there is the new cohort, which essentially is automatic enrollment, which is everything is exactly the same, but by the date of hire. So when you're hired immediately, you're automatically enrolled. So they're not asking you. They're giving you some information, but they say, unless you change, unless you want to make a change, you're going to be automatically enrolled immediately.

And when you automatically enrolled somebody, then you need to also specify where the money goes to. So they have an automatic enrollment. You need to choose the contribution rate, which is the default contribution rate. And they need to have a default fund allocation, which is-- so they have a default contribution rate which is 3%. And you need to put the money somewhere, right? So if I'm automatically doing something for you with your money, I need to put the money somewhere. What the company is doing is the company puts the money into the money market fund.

Is that a good idea, a bad idea? Why is the company putting the money into the money market fund? Is that what you're supposed to be doing with your savings? [INAUDIBLE] Or maybe first explain what is a money market fund, and then tell is it a good or bad idea. Somebody says they saw dot come bubble coming, potentially. But I think if you save for retirement, what should you do? Here's free investment advice for everybody. What should you do? What are people telling you, like, if you go to somebody who does finance?

[INAUDIBLE] but essentially, what people would argue you should be doing is you should invest in stocks or some forms of diversified portfolio of stocks. You should buy, essentially, the stock market, S&P 500, or the world stock market, or the like, which essentially is risky, but also has a much higher return. Now, given that if you started investing at age 25 or 30, if you do that for 30, 40 years, that has a very likely very high long-run return, and a much higher long-run return than you will have if you invest in the money market.

Now, there's a question of what is the optimal risk and so on. But investing, nobody-- no investment advisor in their right mind, maybe perhaps partially because they want to make money, but I think including people who want the best for you will tell you you should invest everything and the money market. Now why are the companies doing that?

Essentially, the company is worried about liability, and about getting sued, or getting, at least, in trouble with their employees. If they default-- so essentially, if I'm not offering any retirement savings, then I'm like, of course, that's not a big risk to take. If I'm in a sense saying, I'm hiring you, and I'm putting your money without really getting explicit consent for you, without you explicitly saying you want it, and putting your money into some investment fund, potentially a risky investment fund, and that thing essentially crashes, and people lose a lot of money, they're going to get really upset with you, potentially. And that's, from the company's perspective, just not worth doing.

Essentially, the company wants to help people. Again, remember a part of the motive of the company was just to get people to contribute in the first place, to essentially get around the IRS nondiscrimination rules. And so the company is really not necessarily that interested in high returns for their workers. And in particular, companies worry potentially about, essentially, the stock market crashing, and then sort of people getting really upset with them, which would be quite costly for them. And then they would be, like, why did you put the money in this risky thing while in fact, you could have put it in the money market, and so on?

And the reason being is essentially because there's no explicit consent, because it's essentially a default option where the default is such that you're automatically enrolled. So unless you tell me otherwise, you're going to be-- the money's going to put there. And so there's never something where the person says, yes, I want the money to be invested in x or y.

So then, this here is the main result of the paper, which is the participation rates in 401(k) savings accounts are dramatically higher with automatic enrollment compared to the previous regime. So you see here, this is the old cohort. This is the window cohort. And these are previous cohorts. So you see that prior to automatic enrollment, the participation increases-- or increased with the tenure. That is to say, the longer you're at the company, the higher the fraction of people who participate in 401(k) savings.

But even people who have been in the company for 20 plus years have only, like, a contribution rate of about 83%. Now, if you do automatic enrollment, which is getting people enrolled-- but people are free to choose whatever they want. If they don't want to invest money into the 401(k) savings account. They can opt out at any point in time. But when you do that, people get essentially way more likely to be enrolled. The participation rate is as high as 86%.

So there are essentially dramatic differences in what you would argue is a very important choice, which is, if you think about what are important choices in your life, savings choice is one of the most important ones. So there's dramatic differences. And when you do automatic enrollment, people are way more likely to enroll than if not.

Now, in addition-- sorry. And then this is here is, I guess, the part about the low income workers. When you look at-- this is-- when you look at the first two columns of this table, here you see these effects. So first you see, this is the window cohort here. This is the cohort that was essentially before the change was done. And here is the new cohort, which is the cohort that received automatic enrollment. The difference is dramatic from like 37% to 86%. And this is now by compensation levels of workers.

And what you see here is, in the window cohort, you see this huge gradient here by compensation. So workers who are-- and this is before the automatic enrollment. Workers who received fewer than $20,000 per year, only 12.5% of them actually were participating in 401(k) savings to start with. And among people who earned a lot already anyway, this thing was, like, 68%. There's a huge discrepancy here. And this is precisely where the IRS discrimination rules were [? fighting. ?]

Now, but if you look at after the enrollment, there's still-- notice that there's still some gradient. There's still, essentially, like, a higher participation rate among the high paid employees. But the grading is much flatter, because essentially, so many among the low income workers or low wage workers are now enrolled as well. So there's huge effect, and the effective is disproportionately large among low income workers. What are potential reasons for that?

So they have more likely-- essentially, their credit constraints are more likely to bind. Or they have a high marginal utility of money, and therefore, they might be less inclined to save. Notice that low income workers also have a high marginal utility of income in the future.

So in some ways, it's not obvious, actually, that when you think about-- [? because when you ?] write down a model, you think about should you save right now versus consume money-- so consume versus save. You will get, essentially, what this-- what this depends on is the ratio of your marginal utility right now versus in the future. And the marginal utility of poor workers in the future will also be high. So it's not obvious that they shouldn't be saving. But I think it's right that we're saying they have more emergencies, and more reasons to essentially not want the money right now.

So one thing could be just it's mechanical in the sense of, in some ways, it could just be that for whatever reason, the poor are not saving to start with. Maybe that's in part due to tenure. Maybe it's due to some other reasons that we don't quite understand. But it might also be that the poor might be more prone to default effects, which could be in part what Carmen was saying. They're sort of-- in a way, perhaps financial sophistication, or education, or information-- just what's the right thing to do-- might be lower. And you might, as in, like, I don't understand what's going on. So I'm just doing what the company does for me and just hope for the best. And then you might just go with whatever that is, and you're less confident in making different-- or overriding such choices.

Another explanation-- and I'm going to talk about this when we talk about poverty a bit-- which is the idea that poverty taxes people's bandwidth and attention. That's to say, like, essentially, being poor is really hard in various ways. You have to struggle with so many different things about paying your bills, and so on and so forth, which is precisely kind of what Jose [INAUDIBLE] was saying in part. Your marginal utility of money might be high, and you're really concerned with that. And you worried about feeding your children, and so on and so forth. And it's just really hard to find a time to think about stuff carefully.

And therefore, you might, again, sort of just go with the default. Because you just don't have the bandwidth-- the mental bandwidth to think about what's best for you and override that. So that's an open question, in fact, whether the poor are more likely to be prone to default effects. But I think what's almost surely the case, or what's been shown in various studies, is that these effects tend to have larger effects among poor, low compensation workers, which is precisely one of the reasons why they were so popular overall.

OK, so then in addition, not only are there large effects on the contribution decision overall. It's like, do you want your the participation decision overall? Do you want to participate versus not? Like, do you want to contribute any money versus not?

But also there seem to be large effects on the contribution rate. That is to say, I think I told you-- remember, I told you the default contribution rate was 3%. So if you were automatically enrolled, the company would just give you a 3% contribution standard. You could change that to 4%, 2%, 7%, whatever you want to do. But the default would be 3%. And you see, essentially, the majority of people-- the fraction of employees in the new cohort, which is the automatic enrollment, chooses 3%.

There's another thing that's quite interesting here. What you see is that among people who arguably would have chosen-- if you just look at the window cohort, there's some people, for example, who chose 6%. 11% of people in the window cohort-- that's the cohort that was just before the automatic enrollment-- 11% of people chose 6%. That fraction actually goes down to 7%. That is to say, some people would have chosen more than 3% in the absence of automatic enrollment, now they are moving down to 3%.

So there is some effect-- there's a bunch of mass coming from 0% moving to 3%. But there's also some mass coming from here that moves from, essentially, positive amounts down to 3%. So there's sort of very heterogeneous effect. Some people, or quite a few people, increased their savings from very low rates, or 0 rates, to 3%. But there's always another other effect that, essentially, some people who would have sat down and done the choice themselves, they would have probably chosen something like 6%, 7%, whatever, even 15%, and some of them have moved down to 3% by the automatic enrollment. So there's some negative effect on savings coming from that, potentially.

In addition, the asset allocation, here, people essentially predominantly just choose the default. Here you see, essentially, the assets or stocks, bonds, and money market. Think about money market essentially as like a 0%-- or a close to 0%-- very safe option. And in the new cohort, 80% of people-- I'm not sure how well you can see this-- but 80% of people choose the money market, which essentially is exactly what the company is choosing for them.

Well, if you have-- in the other cohorts, this is all conditional on participating. The fraction who chooses stocks and bonds is way, way higher. And as we discussed, the long-run return to investing in the money market is way, way lower. So that's potentially, again, another reason why-- in fact, why defaults aren't increasing the faction of people that are choosing to save overall, or to participate on the savings overall. The long-run savings might not necessarily be high.

Now, let me summarize this. I think I said, already, most of this. So 40% to 50% of individuals are passive savers, as you might want to call them. They follow the default plan. They essentially-- both in terms of the participation, but also on the contribution rate and the asset allocation-- they essentially do whatever the company chooses for them.

The suggestive choice is not very attractive unless it's the default. Notice-- remember, the window cohort. The window cohort were essentially the cohort that now could just do whatever they wanted to do. They were told on the 1st of April, I think, '97 or '98, now you can invest immediately. You can do whatever you want. We're not choosing for you, but you are now eligible, even for people who have been at the company for something like, say, six months, three months, and so on. The company just made the participation available to anybody.

Since they had been hired previously, they didn't default people in. They just made it, like, you can change your choice now if you wanted. But among those people, the window cohort is very much looking like the old cohort. They're not sort of following the company [INAUDIBLE].

You might say, well, the company has a new default now. And the new default, that's a recommendation, so I should just follow that default because I don't know what to do. That's not what people seem to do. So it's not like that-- so there doesn't seem to be that the default is a perceived choice that's suggest [INAUDIBLE] company, because the window cohort really does not follow this default or the suggested default. The window cohort really looks like the old cohort.

Now, I only showed you the results for one company for Madrian and Shea. But it's a very robust result that, if you look at the paper by Choi et al. that looks at this in more detail, that finds this over and over again in different companies as well, default effects tend to be very strong in those companies. Now one question you might ask is, what's explaining-- what's the underlying reason why these default effects are so strong? There are potential candidates which are awareness, implicit endorsement by the company, inattention/memory, or potentially present bias.

And what researchers found, it's sort of not necessarily conclusive. But I think the suggestive evidence is very much that-- and particularly from Blumenstock et al., that we think that present bias might be a quite important reason, in addition to the cognitive costs of thinking through different savings scenarios. So people essentially tend to push this off and say, oh, I'm going to change it in the future. At some point, I'll get to it. I'll get to it when it's important.

Notice that it requires some naiveté, because as we discussed, if you were sophisticated, you might of push it off maybe a period or two. But you would never push it off by 10 years, or five years, because you know you're going to lose a lot of money. But if you have some present bias plus some naiveté, you might say, well, I'm going to do it in the future and not right now.

And that seems to fit the data the best, in part because papers such as the one by Blumenstock et al. have ruled out some other explanation, for example, the inattention/memory stories. They set reminders. That didn't seem to have an effect. They also could rule out the implicit endorsement, because they essentially explicitly randomized the people who were defaulted in versus defaulted out. That also had effects. So it doesn't seem to be that.

They also made people very clearly aware of the different retirement savings options. And again, that didn't seem to make a big effect. So while we don't have conclusive evidence that present bias is sort of causing these default defects, in the retirement savings choice scheme, perhaps if I had to guess, or if I had to say which I think is the most likely explanation, we think that present bias and naiveté seem to be the best explanations that we have so far.

So now, one question you might sort of say is, well, is automatic enrollment optimal? I showed you, essentially, that once we do automatic enrollment, people are more likely to participate. However, while they're more likely to participate, also these defaults seem to make some people save less, right? We saw some people going essentially from 6% participation-- contribution rate to 3% contribution rate. They also have a much more conservative asset allocation. So in the long run, it's not obvious, in the case of this company, whether people are actually saving more.

Now, so how would we answer this question? We're going to get back to this when we talk about policy. But how should we set the default? What should we do? What are the considerations that we have? And how can we decide what's good for people?

You can essentially force people to make a choice. There were some companies where, essentially, people were very, very strongly encouraged to-- essentially, here's a form. Fill it out. Tell me what you want. And was sort of part of the hiring package, and essentially was sort of arguing-- I don't think that's actually legal. But it was essentially sort of implying that one [INAUDIBLE] to doing this form before you can start working, and otherwise, you can't start working at this company.

And now, we can look at, like, what are people actually choosing when everybody chooses actively? And that's presumably when people are paying attention, or maybe people do the best they can for their choices, and they're not automatically enrolled either in 0%, or 3%, or other percentages. And so this paper for by Carroll et al. looks attractive choices. And when people choose actively, they look a lot like the new cohort in Madrian and Shea. So essentially, that's sort of seems to say that the default did in fact-- the default into automatic enrollment seemed to have alleviated undersavings.

Because when people choose actively, they choose a lot-- look a lot like the new cohort compared to the old cohort. It seems to be that people, when they're defaulted in the no savings option, they kind of maybe think, oh, I'll do it in the future, and so on and so forth. But then they never sort of get to it. Now, when you actually ask them actually what you want, it seems like people are mostly wanting to enroll.

Now, one thing that is quite interesting here is that it seems like the default effects seem to mostly disappear after three years. That's true for quite a few studies. However, even if that's the case, well, there's no catch-up in levels, right? So if you start contributing after three years, while you miss out on three years of savings, you're never going to catch up overall.

Moreover, people tend to change their employers very frequently, particularly in the US. So if it takes you only-- if it takes you three years to sign up for retirement savings every time you switch your employer, you're going to not save very much. Chetty et al had an example where they find, in fact, very clear effects on long-run savings of these types of default [INAUDIBLE].

Now, let me tell you a little bit now about-- I'm going to sort of move the frames and nudges until next time, which I think should be fine because there's not that much left anyway. But let me tell you a little bit about a cautionary tale about defaults. This is a paper by-- a very short paper by Cronqvist and Thaler that looks at the privatization of social security in Sweden.

So they had 456 funds. And there was one default fund chosen by the government. And in 2000, this default fund was discouraged with a massive marketing campaign. Essentially people said, like, don't choose the default fund. Just pick for yourself. Make the choice that's the best for you. And then in 2003, this marketing campaign ended. And now essentially, the majority, again, sort of chose the default account or fund.

Now, as it happens, the portfolios actively chosen in 2000 did a lot worse than the default. So people, essentially, were choosing on their own. And in part, I think because they're just unlucky, because it happened to be at the wrong time and the wrong place, chose sort of very risky assets. And the default tended to be somewhat more conservative. So people lost a lot of money.

So now, that's the more general issue is that active choice is less attractive if consumers are less financially sophisticated. Again, in the Cronqvist and Thaler case, maybe they were just unlucky. But in general, if you want people to choose actively and they don't know what they're doing, well, then they're going to be potentially worse.

There is similar evidence, for example, for health care insurance choices, where people just don't know what health insurance choice is best for them. So you can give them more options, they might be potentially doing worse, because they just don't know what to do. And then they choose options that are worse, and they're more likely, for example, to be exploited. And then, having a default choice that is perhaps not optimal for everybody, but it's pretty conservative and pretty good for most people, might be, in fact, quite valuable.

And so now, when you think about, like, optimal decision regimes, when you think of an active choice versus default, consumer heterogeneity makes active choice more attractive. If you think people want very different things, and they know about what's best for them, and we don't have that information-- we don't know what's good for some people versus others-- well, it's best to let people choose on their own. Let's let them pick what's good for them, and then they'll be happy.

But active choice only improves the outcome if people actually know what's good for them. If everybody is confused and doesn't know what they're doing, well then, letting people choose actively will potentially make things worse for quite a few people. And so you want to make sure that defaults don't make people worse off. This is what Jose was saying. In particular, people might sort of oversave if you default them into high savings. And then they have credit card debt or other sort of problems arising from that. So we want to be quite careful.

One option one can try and do is try to provide information sessions and make sure people know, here's-- or trying to sort of provide a curated set of three, four options that we think are good options for different types of people, provide information for them, make sure they understand these options, and that then people choose on their own. Another popular alternative-- I'm going to talk about this a little bit when we talk about policy-- which is called the smart plan, or, like, auto-escalation plans, which are essentially more like automatic raises in the future, where you essentially can commit to automatically raising your savings in the future.

You say, essentially, every time I get a pay raise, 1%, or some fraction of that pay raise, or one percentage point of my earnings of that pay raise, will go through my savings. So then, you don't have to reduce your consumption or your salary right now that you receive, but in the future. That's good. In particular, it addresses present bias and loss aversion. Because people are really averse to reducing their paycheck right now, and for present bias and loss aversion reasons.

When I was asking about information versus financial education, I mean kind of both of those. You can find information about what are these options, what's good about this investment option versus another, and so providing financial education, but in particular, financial education plus active choice. Essentially, by active choice, I mean pushing people really to choose, and saying, look, what do you want? Do you want option A, which is no retirement savings, or option B, which is like 3%, or 6%, or whatever.

The current format is more-- essentially, is, like, no information, and the default is essentially 0 savings. And then, you have to actively choose on your own as an employer. And many employees-- sorry, as an employee. Many employees just don't know what to do, and get confused. And then they get frazzled, and then procrastinate it, and then never do it, perhaps in part because of confusion.

OK, and then other settings have enormous before the fact. This is, for example, true for organ donations. I talked about this already in the first or second lecture. There's some other settings such as voter registration, green energy options, and so on. So default effects can be really, really powerful. But one wants to be quite careful in using them. I'm going to talk about this a little bit at the end of this lecture next time.

Let me stop here. Let me just tell you what to read, which is-- one second-- which is the paper by Ariely et al., sections one through four. You can read more if you want, but section one through four already gets you pretty far. This is really a fun paper that questions a lot of assumptions of neoclassical economics, makes you wonder where preferences come from and what they really mean. So it's a fun read. I'll talk about this in the remaining part of this lecture-- which is a bit on frames and nudges-- more generally next time on Wednesday. Let me know if you have any questions. Thank you so much.