

Probable Causation, Episode 38: CarlyWill Sloan

Jennifer [00:00:08] Hello and welcome to Probable Causation, a show about law, economics and crime. I'm your host, Jennifer Doleac of Texas A&M University, where I'm an Economics Professor and the Director of the Justice Tech Lab.

Jennifer [00:00:19] My guest this week is CarlyWill Sloan. CarlyWill is an Assistant Professor of Economics at Claremont Graduate University. CarlyWill, welcome to the show.

CarlyWill [00:00:27] Hello, thanks for having me.

Jennifer [00:00:29] Today, we're going to talk about your research on how race affects police use of force. But before we get into that, could you tell us about your research expertise and how you became interested in this topic?

CarlyWill [00:00:40] Sure. So most of my work focuses on crime and discrimination, and I've always been interested in writing a paper on police officers in general, you know, just because they're the criminal justice actor that you're probably the most likely to come in contact with. So once I kind of got it in my head that I wanted to work on a project about policing, I try to look at the current literature and develop what I thought was kind of the most important unanswered question. And so when I started this project, sadly, these things don't happen in only a few months and I think this is still true today, how much race matters for policing is a really important question. And it's a question I think, that many academics, police officers and just normal people would like to know the answer to. However, when I started the project, I didn't think that we had a really clear answer to this important question.

Jennifer [00:01:41] So your paper is titled "Does Race Matter for Police Use of Force? Evidence from 911 Calls" and it's coauthored with Mark Hoekstra. So give us a bit of context here to start out. What do we know about how often police use force and who they use it against.

CarlyWill [00:01:57] Sure. So I think kind of a recurring theme when talking about this project and just police use of force in general is really just how little data we have on this topic and therefore and often scenarios like how little we know. You know, it would be great to answer this question with some national statistics, for example, on police use of force. However, right now, there's no national database that I'm aware of on police use of force, although I do think that the FBI has just begun requesting data for this national use of force data collection.

CarlyWill [00:02:33] However, even their reporting is not going to be mandatory. So instead of giving just some national statistics, I'm going to talk about what we know based on data from some individual jurisdictions like Mark and I use in our paper some crowdsourced data and some surveys.

CarlyWill [00:02:52] So what do we know? The first thing that people should be aware of about use of force is that, you know, compared to say arrest, for example, use of force is a somewhat rare event, obviously very terrible, but still somewhat rare. So according to "The Washington Post" who has crowdsourced some cool data on police involved shootings, so what I mean by crowd here is that they've combed through different news articles and social media posts, for example, for their information and according to their data last year,

there are about a thousand deaths from officer involved shootings, I think nine hundred and ninety nine, to be exact. And in our setting, which is going to be in a large urban area, which I'll talk about a little bit more later on. Only about one in a thousand 911 police calls are going to end in a use of force and one in ten thousand 911 calls are going to end and a gun use of force.

CarlyWill [00:03:58] The other thing that you should know about use of force in general is that it seems that black Americans and minority Americans, we have data are disproportionately represented in use of force incidences. So in our data, very black neighborhoods, are going to experience 50 percent more use of force incidents than very white ones. And black neighborhoods are also going to be more likely to experience gun use of force incidences. This is reflected in "The Washington Post" data, so according to their data, black Americans and Hispanic Americans are killed at higher rates than white Americans. Finally, we can also see the similar pattern in the police public contact survey, which is a little bit different in that it's going to ask individuals to talk about their contact with police and here black civilians are again more likely to report that they've experienced some type of use of force.

Jennifer [00:04:58] I guess we should also define what exactly use of force is. My understanding is that it can range a lot and might vary across different police departments, depending on who you're getting data from, with anything from being shot and killed to just having the police putting their hands on you. So what do you mean when you say use of force?

CarlyWill [00:05:17] Great question. So in our setting, whenever we say use of force, we're going based off of the two police departments definition of use of force, so it's going to include those more severe use of force incidences where an officer discharges their weapon or an officer involved shooting or something like that, but it's it's going to include any time that a police officer strikes a civilian. So on the less severe side, that could be something like kicking or punching. I mean, obviously, if you're the one being kicked or punched, so pretty severe up to things like tasers or maybe a police officer might use their baton or nightstick that they're carrying. But like in "The Washington Post" data, for example, they've just focused on officer involved shootings, which would just be firing your gun.

Jennifer [00:06:13] OK, so in this paper, you're considering whether police officers of different races use different amounts of force on average, as well as how police officers of different races respond to civilians of different races. So that is how the race of the officer and the civilian interact. So why might we expect the race of the officer to matter in this context?

CarlyWill [00:06:33] Sure. So the first thing is that there is a long history of police mistreatment and discrimination against minorities in the United States. Just one example of that is that police officers and pretty much only white police officers were integral in enforcing the property rights of slave owners and Jim Crow laws. And then more recently, you can think of these high profile killings. So that's going to be an example of an extreme use of force of oftentimes unarmed African-Americans. There's also perceptions of bias and excessive use of force among black Americans. Okay. So, for example, there's a survey by the Pew Research Center that says that only 33% of black Americans believe police and their communities use the right amount of force. And just more generally, race has been shown to be a really important factor in other settings like the labor market and education. And this just kind of reflects a general concern that individuals often favor their

own group or people that might look similar to them. And like you mentioned, in this paper, we're getting at two different research questions. But I think the examples I give here have illustrated we have reason to believe that police officer race might matter, but civilian race is also important, too. So we want to make sure that we're including both in our results. And that's why we have this research design in our paper that's going to focus on the interaction of the race of the civilian and the police officer.

Jennifer [00:08:17] So you mentioned that when you first started this project, you looked at what the literature had to say and didn't find a whole lot. But people have been trying to work on this question for a while. Like you said, it's very important. So before you started this paper, what did we know about how race affects police behavior?

CarlyWill [00:08:32] Yeah. So before this paper, there's definitely was a sizable amount of work done to try and understand racial bias in traffic stops and vehicle searches. You can think of papers by Anwar and Fang, Persico and Todd, Antonovics and Knight, and there's definitely others that, you know, they're trying to address this issue of officer selection and to interactions with civilians. What I mean by that is as an economist, we might like that officers would be randomly dropped throughout a city, but most of the time they're choosing who to stop and who to interact with. So they're going to test for racial bias by modeling police behavior and implementing tests based on hit rates relative to some odds that you might encounter a black or white civilian. There's also related literature that looks at trying to address the difficulty of assessing this benchmark and counter rate by exploiting changes in ambient light to test for racial profiling in traffic stops. So this is papers by like Grogger and Ridgeway, Horrace and Rohlin. These are cool papers, right, because they're exploiting the idea that during the night time it's much harder to tell the race of the driver that you encounter.

CarlyWill [00:09:53] So my work with Mark is going to complement this literature in that we're going to take a pretty different approach to solving the problem of who officers choose to engage with. And there is a paper on traffic stops that uses a similar approach, and that's a paper by Jeremy West. And he has this nice project where he also uses 911 calls and this conditionally random assignment of officers to traffic accidents to look at racial bias and traffic citations, and spoiler, he does find some evidence of racial bias. Right, and, you know, the advantage of Jeremy's paper is that he has this really great objective information on whether certain citations were correct or merited, which, you know, is obviously much more difficult to do in the use of force setting. I don't want to say it will never be possible because maybe someone has a really ingenious idea for getting around this. Right. The idea being that we don't actually know which scenarios truly justify the use of force. You know, the advantage of what we're doing here is that in Jeremy's paper and the other work on traffic stops is that we think we're looking at an outcome that's a bit more severe and potentially more important when we focus on police use of force.

CarlyWill [00:11:19] But there's probably another group of papers and specifically a paper that's forthcoming by Roland Fryer that are most similar to what we're doing in that they're going to try to focus on police use of force. So in Fryer's paper, he has amassed a pretty impressive range of data sets. One of those data sets is going to be this stop and frisk data from New York City. So what's great about this, right, is that he's going to see a lot more incidences, even if they don't necessarily end in an arrest. So there he's going to use a really rigorous set of controls and he's going to conclude that blacks and Hispanics are more likely to experience non-lethal force after controlling for a bunch of stuff. Right. So the non-lethal force here, you can think of potentially the police officer hits the civilian or uses a taser, but they're not going to use a gun. And additionally, he also has data on

officer involved shootings in a couple of different cities and counties, which she's going to benchmark using arrests. So what I mean here by benchmarking, right, is that when he uses these officer involved shooting data sets, he's just going to have a lot of rich information about events that actually led to a shooting. Right. So the problem there is that you need to try to figure out what happened in the counterfactual world, right, in events that didn't lead to a shooting. So he's going to compile a data set of kind of more severe arrests to stand in as a counterfactual. So using this data, again, he's going to do a selection on observables approach to try to control for all the contextual factors that he can observe in the data. And he finds that black and Hispanics are not more likely to experience an officer involved shooting.

Jennifer [00:13:16] So two things. So, one, there's an interview with Jeremy West that I'll post a link to in the show notes as he was on the podcast sometime late last year. And on the Fryer paper, I think the main caveat to that selection on observables approach, which is basically just controlling for stuff, right, that people have brought up, is that you have - he basically controls for whatever the police wrote down in their reports of what happened - I think is what's going on there. Is that - is that right? Am I remembering that, right?

CarlyWill [00:13:44] That's exactly correct. And so, you know, there's two things that are going on when we just use a selection on observables approach that might make us worry that we're not getting it right. You know, like the first thing is like imagine that you go to a group of police officers and you say, like, hey, primarily, hey, guys, primarily guys, you know, whenever you're in black neighborhoods, you're like 50% more likely to use force compared to white neighborhoods. And I think what they would say is that, well, yeah, that's not necessarily about civilian race. It could be that these incidences in black neighborhoods are fundamentally different than the ones in white neighborhoods. Right. So maybe the civilians are more likely to have weapons or, you know, there's just something different about the geographic layout of those neighborhoods. So it's not necessarily about civilian race. And then the other thing about using controls from a police record is that some of the things that he observes might actually be affected by the interaction of the police officer and the civilian themselves. So it's not exact - I'm not exactly sure that we should be controlling for those type of things - like the way that the civilian responds, for example, might be part of kind of this treatment effect of observing a police officer.

CarlyWill [00:15:15] Yeah. So I think that's what's really nice about our data, is that a lot of the characteristics that we observe come from the 911 dispatcher. So it's going to be these characteristics that are determined before an officer even shows up on a scene. And whenever we look at our results for civilians and police officers, we're able to control for just baseline differences between black and white areas. So that means we're controlling for the things that we do observe and also other characteristics that we can't observe or at least are not observable to the researcher whenever I'm trying to do some type of analysis.

Jennifer [00:15:58] So we're recording this in mid-July. So the country's been talking about these issues for a solid two months now. And of course, it comes up periodically whenever there's a viral video of some civilian being shot or killed by an officer, apparently unnecessarily. And so this is just such an important and policy relevant topic. And you've already been alluding to some of the challenges here, but why don't we know more than we do? Why is this so difficult to nail down?

CarlyWill [00:16:29] Yeah, so I think there are two big challenges to understanding how race is going to matter for for policing. The first is what I've been talking about is this

benchmarking problem. Right. So again, what's problematic here is that oftentimes we don't observe anything about incidences that don't lead to an arrest or don't lead to a use of force. Just to kind of give you a sense of the scale of the issue, if you're looking at our 911 data and say like we only had data on the calls that ended in a use of force, right, you would be missing 99.9% of civilian and police officer interactions. Like think about how crazy that would be in another context. Like imagine you want to look at the effects of years of education on earnings, but you're missing data on 99.9% of students. And you know, the 0.1% of earnings data you observe might only be for those that completed college with a 4.0. So in the policing context, it's not just that we're missing data, but we're also missing it for these less severe and potentially fundamentally different interactions.

CarlyWill [00:17:43] And then the other problem is the problem of selection. So most of the time, police officers choose which civilians they're going to interact with. So think about the statistic from The Washington Post. Black men are almost two and a half times as likely to be killed by a police officer compared to white men. But it's difficult to get a sense of how much race matters just from that stat alone. We'd probably want to know more information like how often do black civilians come in contact with police officers? And further, like, are the characteristics of those interactions with black civilians, like the severity of the interaction similar to the interactions that police officers are having with white civilians? So I see those as kind of being the two big challenges that researchers have have tried to deal with in the past.

Jennifer [00:18:41] Right. And so what you'd really love to see, like the dream data, is something like you'd want to see the pool of all people who are actually committing crime or something like that or who are actually behaving in a dangerous way. But that's probably impossible. We'd have like cameras everywhere and learn how to use that footage, you know, in a way that researchers could use it, which maybe isn't as farfetched as I think, but that, you know, you're always just going to have this pool where the officers decided to stop and talk to the person and then the scenario either escalated or it didn't. But you don't know what would have happened if they'd stop the person they didn't stop. It's that missing data in that problem that is kind of the heart of this. Am I getting that right?

CarlyWill [00:19:25] Yeah, that's exactly right.

Jennifer [00:19:27] So walk us through what you're going to do in this paper. What's the natural experiment that you and Mark used to separate the effect of officer race from other factors that might affect whether force is used in the events that they're dealing with?

CarlyWill [00:19:41] Sure. So in our paper, we're going to overcome both of these two challenges, both the benchmarking and selection that I just talked about by first using a setting where police officers have no discretion in which events they respond to. What I mean by that is that we're going to be using two different cities with a 911 dispatch procedure that says, okay, you know, it's essentially going to be a coin flip, whether a civilian is going to have their call assigned to a white or black officer. And then because we're using this 911 data, we're going to observe a bunch of characteristics and information about all the police and civilian interactions even when there's not an arrest or use of force made. So we're going to use this set up to answer two different research questions. Like I mentioned, we're going to be using data from two different cities. The first one is primarily white and black, and the other one is primarily Hispanic and black. I'm going to kind of walk you through the thought experiment that you should have in your head for our paper. But I'm just going to be using a white and black example for simplicity. But if you want to read the paper, you can hear more about the Hispanic and white city.

CarlyWill [00:21:01] So first, we're going to get at do black and white officers use force at different rates? So we can answer this question by just comparing the probability of use of force across officer race within a specific beat and time. So police beat, you can just think of it as a small geographic area where police officers are going to be assigned to work. Okay, so right, again, we're able to do this because in our setting, officers don't get to choose which calls that they're working on. They're just going to be assigned them. So it's going to be a coin flip whether a black or white officer is going to respond to every call. So next, we also want to think about whether civilian race matters. So, of course, like I've explained, this is pretty difficult to do because it could be possible that calls and interactions in black neighborhoods are just much more or much less dangerous than those that occur in white neighborhoods. This is also the point where I like to give that example that I gave before about like imagine if you go to a police department and you were to try to, you know, tell them that, okay, you're using force a lot more in these really black neighborhoods, like, is this due to civilian race? And I think the first thing the officers would say is there's a lot of things in those interactions that, like you're not observing as a researcher. So like how dangerous the scenario was, if the civilian has a weapon. So that summary statistic alone isn't going to tell us if civilian race is really what's mattering there. But it also means that I'm not super comfortable just using a selection on observables approach because I'm not observing everything about an incident as a researcher.

CarlyWill [00:22:54] So rather than just asking if police use force more in black versus white neighborhoods, we can get at this question, which is, does the interaction of police officer and civilian race matter? So specifically, we're going to think about do officers - do black officers scale up their use of force as much as white officers whenever they travel to black neighborhoods? So if black and white officers both scale up or substantially increase their use of force by the same amount in neighborhoods with black civilians, perhaps race is not as important in that setting. But if we observe that whenever white officers travel to black neighborhoods, they scale up their force much more than black officers, it's going to show us that civilian race does matter.

Jennifer [00:23:50] At least with the interaction there. So, as you said before, a challenge in this setting, as in many settings, is you don't know the right answer. You don't know exactly how much force should have been used in any given case. And so you're really going to rely on the black officer as a reference group of sorts to get a sense of what a different type of officer might have done in the same scenario.

CarlyWill [00:24:15] Yeah, that's exactly right. Yeah. One of the challenges with most policing work, aside from that ticketing example I gave in Jeremy West's paper. Sadly as the researcher, I don't know the truth, so I don't know which incident is really where use of force was truly justified. So all that I'm assuming here is that in a world with no bias, the difference in use of force rates between black and white officers is going to be the same no matter the race of the neighborhood that they're in.

Jennifer [00:24:47] And then to go back to the thought experiment we have in mind here. So you describe it as the coin flip, but let's say a little bit more just about how the 911 dispatchers are deciding who gets called to the scene. Is it basically just the closest officer?

CarlyWill [00:25:00] Yes. So it's a little bit different in our two cities. But the way that you should be thinking about it in general is that a 911 dispatcher is going to get a call. They're going to record a bunch of important information about the call, like the severity of the call

and kind of a standardized description. Then they're going to look on their computer screen and they're going to see, okay, for this really severe and urgent call that I have, who's going to be the closest police officer? And they're always going to dispatch the closest guy. They're also able to observe on that computer screen, right, if one police officer in the area is responding to another call, they're not going to assign that officer. Or if the police officer is on a lunch break or something like that, they're also going to know that that's going on and just send the closest available police officer, which gives us that nice kind of random variation or a coin flip, in officer race.

Jennifer [00:26:00] Right. And then you'll be able to control - because these officers are constantly moving around and responding to all kinds of calls - you'll be able to actually control for the individual officer even. And so you can even look like within an officer for the same person responding to this call versus this call. You can see how their responses change as they go to different neighborhoods or as different officers respond to crimes in exactly the same neighborhoods because there's so much movement around. And so, yes, you're able to kind of control for all the other stuff that might be going on there. But it's - basically the coin flip aspect of this is really just people are moving around all the time, so the closest available person gets it and the dispatcher doesn't care if that person is black or white.

CarlyWill [00:26:45] Yeah, that's exactly right. I mean, I think one of the thoughts that we kind of have going into projects is like maybe some of our effect could be explained by strategic assignment of police officers to specific areas, for example. But whenever we include this individual officer fixed effect, right, so when we just look, say, if I was a police officer, we'd just look at the scenarios where I encounter a black versus a white civilian, our estimate is not very different than whenever we don't include that officer fixed effect, which I think is - provides evidence that there doesn't seem to be, at least in our setting, a lot of kind of strategic positioning of police officers. Obviously in different cities, this could be a unique story, but for us, that seems to be the case.

Jennifer [00:27:38] Okay, so as you said, you're doing this in the context of two specific cities, which you are not able to name in the paper due to your data use agreements. But tell us a little bit about these cities. What do they look like?

CarlyWill [00:27:51] Yeah, so we're using data based on all the 911 calls from two different cities, which we're going to link to a police officer race and use of force. The first city, again, is going to be primarily white and black civilians, has a population of over 240,000 people, so it's pretty sizable, and it has a homicide rate that ranks it in the top 20 among the nation's largest one hundred cities. So it's definitely a city that has some crime. And, you know, it is not the same city that Roland Fryer is using in his paper, which could be part of the reason that we find different results. Of course, the other reason is differences in methods, which I talked about before. And then in the second city it's going to be primarily Hispanic and white civilians, and it has a population of over 150,000 civilians.

Jennifer [00:28:58] Okay, so tell us a little bit more about these these data sets. What information were you able to get from these two cities?

CarlyWill [00:29:03] First of all, we have this nice 911 data that comes from both cities, and it's going to include important information about these police interactions. So in our first city, there's about 1,300 police use of force and 94 incidences of police use of force with a gun. The data is also going to include the time and date of the call, the priority

assigned to the call, which you should think of as a measure of how severe and urgent the call is, and dispatchers have a standardized procedure for determining that. It's also going to include a standardized description of the call. So that might be something like, you know, this is a robbery that's in action or it could be something less severe, like, hey, could you patrol when possible? Then we're also going to observe the first officer or officers dispatched to the scene and of course, whether or not force was used. And like I mentioned, we're going to see the first officer sent. There's - it's definitely the case that other guys might show up to the call, but we're not going to see them. So we're kind of thinking of this as an intent to treat effect.

CarlyWill [00:30:22] We also see the address where the call came from, okay, which we're going to geocode, and that allows us to link into census block group data to assign civilian race. Right. Like I mentioned before, obviously a huge issue is like all the interactions wherever there's not an arrest or not a use of force, we don't typically know a bunch about them. And even in the 911 data, we don't know the actual race of the civilian who, like, made the call. Right. So the best that we could do is that we can link to census block group level data that gives us information on civilian race. And we think this is a pretty good - a pretty good proxy because the cities that we're working in are pretty segregated. So we have a lot of variation in civilian race. So like you can look at the 911 calls and see that police officers of both races are going to areas with lots of black folks and also areas with lots of white folks.

CarlyWill [00:31:26] And in the second city, right, where we have mainly white and Hispanic police officers and civilians, we have pretty much the same information on just under 1 million calls. But there we have a bit more - 3,000 use of force incidents. So, yeah, that's the data that we had. And just - I guess if you're someone that's interested in working with this type of data, I like to get my little plug in there for public information offices. That's how I - that's how I started the process of getting this data. I did try to get data from a couple of different cities, and sometimes it can be a challenge to get the data from them. And obviously, in our scenario, most of the data we were interested in was covered by public records law. So that was something that I could lean on. However, at the end of the day, we did need the police departments to agree to release the name of - not necessarily the name, but the race of the police officers to us. So that was definitely the kind of the biggest hurdle in the project.

Jennifer [00:32:39] And that's what required the data use disagreements?

CarlyWill [00:32:41] Yes.

Jennifer [00:32:42] Okay, so what exactly are the outcome variables that you're going to focus on here?

CarlyWill [00:32:46] So we're going to focus on two different outcomes. The first one is whether a specific 911 call ends in a use of force. And then the second one is whether a 911 call ends in a use of force with a gun. In the second city, unfortunately, we're not going to observe gun use of force separately. This use of force information is going to come from a police use of force incident report. And most of the time, these reports are reviewed by police supervisors. Sometimes they're reviewed by civilian review group, although not always. And we decided to focus on use of force because we think it's arguably the most extreme and one of the most important measures of policing severity.

Jennifer [00:33:33] So you're going to be comparing the behavior of officers of different races. But of course, race isn't randomly assigned. So this is - this analysis is going to tell you about how behavior differs across the types of officers who are white and the types of officers who are black or Hispanic. So did these groups of officers differ on any other observable dimensions other than race?

CarlyWill [00:33:55] Yeah. So in our first city, black officers are more likely to be female and have a little bit more experience, although that difference isn't super big. However, in our analysis, right, when we're focusing on kind of average differences between black and white officers and in our police and civilian interaction results, we're going to control for everything that we observe, right, about the police officer. Of course., Right, those are the things that we observe. And it's certainly true that black officers could be different in other ways as well. I will say when we get into the results, we can talk about this a little bit more, but it doesn't seem to be that the results are being driven by female or male officers. Whenever we just look at female officers and male officers separately, our results look pretty much the same. But of course, I do think it would be interesting to do more work if you had a data set that had a lot more information about the different backgrounds of police officers. Although I do think from a policy standpoint, I think the question that people are interested in is not necessarily like how might officers behave differently if you were to randomly, like, reassign race, but nothing else. I do think in the policy setting, people are kind of caring about that package deal of whether white officers behave differently than black ones.

Jennifer [00:35:32] All right. So let's let's get into those results. Let's talk about the first city first, where the population is mostly white and black. So was there any difference in how often white and black officers use force on average?

CarlyWill [00:35:44] Yes. So on average, in our first city, white officers are going to be about 55% more likely to use force compared to black officers. So if you think about this in terms of 911 calls, if black officers use force for 10 and out of - or for 8 out of 10,000 calls, then white officers are going to be using force for 12 out of 10,000 calls. And we also see a similar pattern for use of force with a gun. So for use of force with a gun, white officers are about 65% more likely to use force with a gun compared to black officers. And this result is going to hold up to controlling for all this nice stuff that we observe about the call, as well as the experience level of the officer and the gender of the officer. And like I mentioned, if you were to say, just look at these results for females, the difference between white and black officers looks the same or similar to if you just focused on male officers. So it definitely seems to be something about race that's mattering.

Jennifer [00:36:58] Okay, and then did white and black officers respond to the racial composition of the neighborhood differently in that city?

CarlyWill [00:37:05] Yes. So if you were to look at the odds that a black officer were to use force in white versus black neighborhoods, you're going to see that black officers only increase their force modestly whenever they're dispatched in neighborhoods with a higher proportion of black civilians. However, white officers are going to use significantly more force as they're dispatched to more black neighborhoods. So to put things in more specific terms, whenever a civilian encounters an opposite race police officer there going to be 30 to 60% more likely to experience any type of use of force and 5 times as likely to experience a gun use of force. Okay, so another way to think about this result is to compare police officer behavior - and really white versus very, very black neighborhoods, okay, so white police officers are going to use 85% more force in black neighborhoods

compared to black officers. Okay, so, right, this statistic might not reflect anything about civilian race if white officers are just more aggressive across the board. But whenever we look at that, we can see that white officers only use 30% more force in white neighborhoods compared to black officers.

Jennifer [00:38:35] Right. So they're both scaling up the amount of force they're using, but white officers are scaling up much more.

CarlyWill [00:38:41] That's right.

Jennifer [00:38:42] So then let's go to the second city where the population is mostly white or Hispanic. So is there any difference in how often white and Hispanic officers use force on average in that city?

CarlyWill [00:38:51] So this this finding is going to be different than what we found in the first city. Right. So in the Hispanic and white city, we're not going to document any meaningful difference between white and Hispanic officers in use of force. And obviously, you know, this is different than our first city where we have really big average differences between black and white officers.

Jennifer [00:39:15] And then when we turn to that interaction, so when you're looking at white and Hispanic officers and how they respond to the racial composition of the neighborhood, do you see any differences in that city?

CarlyWill [00:39:27] Yes. So here our results are going to reflect the interaction results in the first city. So, right, our results are going to show that the rate at which white officers use force increases by more as those officers are dispatched to more Hispanic neighborhoods compared to Hispanic officers. So whenever a civilian encounters an opposite race police officer, they're going to be about 75% or more likely to experience any type of use of force.

Jennifer [00:40:01] So you then conduct a whole bunch of additional tests to explore these results and convince yourselves that your estimates represent the causal effects of which officer responds to the 911 call and in particular, the race of the officer. So talk us through some of those and what you find.

CarlyWill [00:40:16] Sure. So I'm just going to talk through one here. But I think it's nice because it is going to be a test that uses a lot of the information that nice information that we observe on our 911 calls. Right. So one way in which we can assess if our research design is working is that we're going to use all the call characteristics that we observe, minus police officer race, to predict police officer use of force. Right. So specifically, we're going to regress police use of force on our beat by time fixed effects, right, that allows us to isolate our random variation. Then we're going to take those residuals, which are just going to capture changes or deviations from average use of force for that specific beat and time on every covariate we observe for each call. Okay, so that's going to include like call priority, latitude, longitude, call time between the time the call was received and the calls dispatched, home beat, per capita income. So lots of important stuff as well as fixed effects for like the call description and call taker. Then we're going to use that resulting regression equation to predict the likelihood of force that would be used for each officer and each call. So intuitively, right, this is going to produce a linear combination of all these different call characteristics where the weight or the importance is being chosen to best predict the odds of force being used.

CarlyWill [00:41:54] Okay, so another kind of silly way to think about this is that I could show you a result that says that the number of hot dog stands in an area or in a neighborhood is not correlated with use of force in the area. Right. However, we probably don't think that's super informative because hot dog stands probably don't help us predict where use of force might occur. But when we use this linear combination of call characteristics, instead this predicted use of force is going to emphasize the characteristics that matter the most or that are most predictive for use of force to begin with. Then we're going to ask whether white and black officers are dispatched to calls of similar underlying danger when assigned to a neighborhood of a given racial composition. So if our identifying assumption of our approach is valid, then we would expect that predicted use of force should be the same for white officers and black officers, and that is what we find. So we're pretty convinced that we are actually identifying the effect of police officer race or that interaction.

Jennifer [00:43:02] Great. So I want to go back to our earlier discussion, at least briefly, of what it means to be a white or black or Hispanic officer in this context. So, as we were saying, race isn't randomly assigned. So race is going to come as a bundle with other characteristics that might differ across race or ethnic groups. So the kinds of things I have in mind here are white and black officers might be motivated by different goals when they decide to become officers or race could be correlated with other factors like life experience that changes how they behave on the job or how they interpret situations that they're dealing with. So especially when you're thinking about policy implications, which we'll talk more about in a few moments, how do you think about this issue in the context of your study, that the race of the officer could be a proxy for those other things that are correlated with race?

CarlyWill [00:43:51] Yeah, so I think that this is definitely an interesting question, something that I would want to know more about. But, you know, I do think that - it's I guess it's challenging for me to think about how from a policy perspective if it's - if it really matters, that we can separately identify the things that are really correlated with race? Right. I think the question you're asking is like, how would some police officer behave differently if we could somehow only just change his race, but not like his education experience or something? I don't know. I think that maybe that's a little bit different than the question that most people are thinking about. Like, for example, like if you're a police chief or something. I do think one of the questions you want to know is like, all right, well, what if I just replace a lot of my white officers with black officers? Right. And I think that is something that we're able to get at in this paper.

Jennifer [00:44:54] You're able to compare the average current white officer with the average current black officer. I guess the reason I'm bringing this up again is that, you know, I think the policy that we could imagine putting in place here is, well, we're just going to try to hire a lot more black officers. Right. And I guess the question is, what is it about the current black officers that makes them respond differently than our current white officers? Is it just the fact that they're black and all the stuff that is typically correlated with that? Or is there something about the current sample who selected in to be an officer in the past that might be different than if we were to dramatically expand our recruitment of black officers, for instance. So I guess I'm trying to get at, like, are there things about the - to what extent could there be things about the black and white officers that might change if we change who we recruit?

CarlyWill [00:45:52] Yeah, no, I think that's definitely an important question that I - yeah, I'm not sure that we can speak to that in our paper. Right. I think kind of the thought experiment in our paper is like almost imagine if you could like rearrange the officers that you have on your force within a city. I think we're getting at that pretty well. But yeah, like I think a really cool field experiment would be, okay, what if we try to recruit a different group of police officers? Like would we see different effects on racial disparities or maybe just the way that they police in general and maybe someone is out there working on that. But I definitely think, yeah, that would be a cool endeavor to take on.

Jennifer [00:46:36] Yeah, it's a really interesting point about, you know, yours - your being able to speak to potentially like how to rearrange officers. So one potential take away could be these, especially in the first city, maybe you should be assigning the black officers to only patrol black neighborhoods and white officers to only patrol white neighborhoods. And that sounds politically problematic for a whole bunch of reasons.

CarlyWill [00:46:57] Yes, yeah.

Jennifer [00:46:59] But like, it could be more efficient or maximize welfare in some way.

CarlyWill [00:47:06] Yeah. I mean, if your only goal is to reduce racial disparities in the way that we're defining it. If you just made everyone have interactions with their same race like that would fix the racial disparities as we define it. Obviously a separate question is, again, like a limitation of this design is that we don't know the truth. So we don't know like which groups are getting it wrong. And perhaps someone is going to think of a clever way to get at that. But right now, with the data that we have, I'm not sure that that's something that we can speak to.

Jennifer [00:47:39] Right. Yeah. So is there any other recent research about race and police behavior that's come out since you first started working on this project?

CarlyWill [00:47:49] Yeah, I mean, especially with the things that - the protests that have been going on, I think a lot of people are thinking about recent police behavior more than they were in the past. But kind of two papers that come to mind is, first, Elizabeth Luh's work on racial bias and the systematic misreporting on trooper reports in Texas. So she can show that police officers miss record on purpose the race of who they encounter. I thought that was really interesting and I hadn't seen another paper like that. And I think her paper is nice because it's - it speaks to this data problem that we have in policing. It's not just about not having some incidents recording, but just like the quality of the data that comes from police reports in general. I also think it shows that race of the civilian is apparently important if police officers are willing to kind of go through the effort to misreport it.

Jennifer [00:48:54] Right. And just to elaborate on that a little bit. So basically what she finds in that paper is that officers are strategically misreporting the race or ethnicity of particularly Hispanic drivers - I think she's looking at the drivers here - when they search them and don't find anything. So they misreport them as being white so that they look less biased. This is very provocative results that then there was a policy change that told them they have to actually ask the driver what their race is and that fixed it, which is sort of amazing. So I agree that's a really cool paper.

CarlyWill [00:49:28] Yeah. And then I know that really recently Bocar Ba and his coauthors, which I believe are Knox, Mummolo, and Rivera - I hope that is correct. They

have a really timely paper that's looking at diversity in policing. So they document that when black, white, Hispanic officers go to really similar areas, that black and Hispanic officers are going to make fewer stops, arrests, and use force less often than white officers. So this is exactly in line with what we find. I think that's - something that's really informative about their paper that I hadn't seen a lot of work on before is that they're able to show that Hispanic officers that speak Spanish make fewer arrests as well. So I think it that brings up an important aspect of policing, which is like just communicating with the civilians that you're involved with. So I'm excited to see more work kind of about these other different police officer characteristics.

Jennifer [00:50:38] Yeah, and I mean, I guess something we haven't talked about with respect to your paper is that that kind of gets at is like what it is that makes white versus black officers respond differently. Is it - is it all a communication issue or is it just flat out discrimination? Is it that white officers are less familiar with the area, and so they're more likely to get scared when they're dealing with a black offender? Or is it just kind of all the different possible things that could - that could lead to these racial disparities? That, again, I think are interesting just because they could potentially point to different policy solutions, like maybe training matters or maybe, yeah, in the Hispanic case hiring more people who speak Spanish - in that case they don't have to be Hispanic, right, if you can get white officers that speak Spanish. So, yeah, that's a - that is - that's really neat.

Jennifer [00:51:29] So what should policymakers be taking away from the literature in this area, including the study that we talked about today? What do you see as the most important implications here?

CarlyWill [00:51:38] Yeah, so I think the thing is that there's a bunch of disagreement about whether we have systemic problems when it comes to race and police use of force, right. I think some people are very convinced that that there is a problem and other people might say something like, okay, these high profile incidences that we're seeing in the news are really rare and terrible, but maybe they're just one offs, right? So I think that moving towards a path of reform is going to be a lot more challenging if we don't have consensus on whether there's a systematic problem or not. Like a lot of the examples or - and the questions that you're asking about, you know, like what are the different characteristics of a police officer that might matter? I think those are really important questions. But if we can't, you know, kind of come to agreement that we do have a systematic problem, I think it's going to be challenging to convince some people to take on a reform minded path. Right. So for me, the most important policy step forward is just that police departments and cities should make data on police behavior much more accessible. And it would be over the top helpful if cities could think about implementing a dispatch procedure, or at least some of their officers are not getting a bunch of discretion. So that would make it a lot easier for researchers or police departments themselves to do the type of analysis that we do in our paper. And I think Mark and I take on this is that if other cities are wanting someone to do analysis on their data, like we would be more than happy to do that type of work.

Jennifer [00:53:27] Yeah, having sunlight - sunlight can be very helpful in these sorts of situations, whether it's a - you know, formal academic researchers or just concerned citizens with some data skills. I agree having more data out there is super helpful all around. So what's the research frontier here? What are the next big questions that you and others are thinking about going forward?

CarlyWill [00:53:48] Yeah, so I think there are two big questions that I'm really interested in right now. And one of them is what we've been talking about a bunch is that at least for me, one of the most striking results is just the average difference between black and white officers. Like obviously we have - there's some fundamental difference between the type of people that are the type of black officers and the type of white officers that we have. So I'd love to be involved in thinking about why black white officers behave so differently. You know, I think there's a little bit of work about how characteristics of officers - so there's a paper by Rim, Ba, and Rivera about ingroup bias that looks at award nominations and then also Holz, Rivera, and Ba have a paper on peer effects and police use of force. But beyond that, I think there's definitely a lot of different characteristics and work to be done to get a better answer to this. So we're probably interested in what happens if we try and recruit different types of officers. Like what if we get officers with different education backgrounds or military backgrounds? Could that help with racial disparities? And then the second thing that I've been thinking a lot about is that, of course, all these interactions involve like at least two people, a police officer and a civilian. And I've spent a lot of the podcast whining about how we don't have that much data on police officer behavior, but like, if - we really, really, really don't know much about, like, how civilians are behaving and how they react to police officers. So there's that one survey I mentioned and then there's maybe a little bit of information in police reports. But beyond that, we really don't know anything. So I think it would be cool to do some type of field experiment where maybe a researcher is going to systematically vary the behavior of a civilian to look at how a police officer responds to that. So you could think of a field experiment where, like, you randomly assigned individuals to be rude or not during a traffic stop, for example, and then think about how police officers respond to that. Obviously, I don't know how much IRB approval that would require.

Jennifer [00:56:09] I was just thinking that sounds like it could be potentially dangerous, but I agree with you, very cool.

CarlyWill [00:56:13] But, you know, I'm sure people out there are thinking about similar types of things, so maybe they have a unique idea for how this could be done, keeping people safe, of course. Yeah, but like, I think that would directly to speak to, you know, the idea that it's just not police officers operating in a void. They do have to deal with us civilians.

Jennifer [00:56:36] My guest today has been CarlyWill Sloan from Claremont Graduate University. CarlyWill, thanks so much for doing this.

CarlyWill [00:56:42] Yeah, thanks for having me.

Jennifer [00:56:50] You can find links to all the research we discussed today on our website, probablecausation.com. You can also subscribe to the show there or wherever you get your podcasts to make sure you don't miss a single episode. Big thanks to Emergent Ventures for supporting the show. And thanks also to our Patreon subscribers. This show is listener supported. So if you enjoy the podcast, then please consider contributing via Patreon. You can find a link on our website. Our sound engineer is Jon Keur with production assistance from Haley Grieshaber. Our music is by Werner and our logo was designed by Carrie Throckmorton. Thanks for listening and I'll talk to you in two weeks.