

## Probable Causation, Episode 93: Justin Holz

**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. My guest this week is Justin Holz. Justin is currently a postdoctoral scholar in the University of Chicago's Department of Economics and he'll join the University of Michigan's Ford School as an assistant professor this fall. Justin, welcome to the show and congrats on the new job.

**Justin** [00:00:34] Thank you so much. It's a pleasure to be here.

**Jennifer** [00:00:36] So today, we're going to talk about your research on how injuries to police officers affect other officers 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?

**Justin** [00:00:49] Sure. So so I'm an applied microeconomist who generally studies the role of social influences on decision making. And so what that means is I study questions like how can we use social image concerns to get big businesses to pay their taxes? How to different types of social preferences lead people to punish price gougers? Or how does racial discrimination impact the way hiring managers respond to minimum wage hikes or whether police officers decide to stop motorists? My interest in this particular question really came from a desire to understand the policing production function. That is, what are the inputs that determine policing outcomes, like arrests or physical force. And how do officers convert those inputs into the outcomes that we care about. I mean, there's there's about 15 million arrests in the United States every year. Many of these arrests involve the use of force and lead to about a thousand people being killed by police every year. And despite growing concern about this police violence, there just isn't a lot of rigorous scientific evidence of what types of policies can improve the situation.

**Justin** [00:01:56] And this is mainly because there's just not a lot of empirical evidence about what are the determinants of police force use. And so in this setting, with my coauthors, Bocar Ba and Roman Rivera and I became interested in understanding whether officer injuries on duty lead to peers becoming more likely to use force.

**Jennifer** [00:02:14] Your paper is titled "Peer Effects in Police Use of Force" as you said, it's coauthored with Roman Rivera and Bocar Ba and it's forthcoming in the American Economic Journal Economic Policy. Congrats on that.

**Justin** [00:02:25] Thank you.

**Jennifer** [00:02:26] So let's start with a definition. What type of force are you focused on in this paper and how common is it in police civilian encounters?

**Justin** [00:02:33] Yeah. So so the data for our study comes from Chicago, and the Chicago Police Department defines force as physical contact by an officer as used to compel a subject's compliance. So in this paper, given this definition, we focus on types of force that generate what is called a tactical response report. That is, when either an officer uses more than a minor level of force, when a suspect alleges an injury, when a suspect resists arrest, or in situations where the suspect uses physical violence against the officer. So this is going to include types of force like control tactics, armed guards, escort holds or responses without weapons, open hand strikes, punches, kicks, and then also responses with weapons such as pepper spray, tasers and firearms. Its going to exclude things like

holds or handcuffing and then for reasons related to the identification strategy, it's also going to exclude force involving canines.

**Justin** [00:03:31] Generally, it's hard to know how common force uses in police civilian encounters because we only observe encounters that escalated to the point where a tactical response report needed to be generated. So we don't see that denominator to know how common force use is, but what we can say is that about 2% of the officers in our sample use force at least once in any given week. These are primarily going to be controlled techniques and other types of force without a weapon, but there are also instances of higher levels of force use like tasers and firearms. And so these are going to be the types of force that we focus on in this paper.

**Jennifer** [00:04:09] Okay, great. And I guess probably does not need saying for those who who study this topic, but we might as well say out loud also that use of force in general and in terms of like how frequently is it used in the United States as a whole we just we really have no idea. We have these like data from one city here, one city there, right.

**Justin** [00:04:28] Yeah, exactly.

**Jennifer** [00:04:29] We're working on getting better data in this area, but it is a tough thing to study. And so so on that note, what had we previously known about what affects the amount of force that police use?

**Justin** [00:04:41] Yeah. So, so there are two important things that are that one should understand in this setting. The first is department policy on force use, and then the second is empirical determinants from the literature. So in our setting, the department's policy is that police officers should attempt to gain voluntary compliance whenever possible, and then they are allowed to use force in situations that are, quote, objectively reasonable, necessary and proportional to the subject's actions. However, there's no formal definition of objectively reasonable, and officers have a lot of discretion and have to make split second judgments in these high stakes settings. And so the Chicago Police Department, where the study takes place, instructs officers to consider whether there is an imminent threat to themselves or third parties, how much the threat that is, and then also whether the suspect has immediate access to weapons.

**Justin** [00:05:31] So that's the department's policy. The literature on this, which comes from settings in Chicago and out, is still sort of evolving on this topic. So there's a couple of papers that I can mention. There's a nice paper by Adger, Ross and Sloan that shows that an officer's propensity to use force depends on the prior propensity to use force of the field training officer that they're assigned to. There's another another paper by Hoekstra and Sloan that shows that white officers use more force than black officers, especially in black neighborhoods. And Roland Fryer has a paper that that shows that minorities are more likely to experience some form of force in interactions with police. And there's also been some randomized controlled trials by Barak Ariel and other people that show that body worn cameras can reduce force use and the presence of tasers can lead to increased force use. So we know some things about the determinants of force use. However, this is still an evolving area and there's a lot more work that needs to be done.

**Jennifer** [00:06:28] And what had we known about peer effects in policing more broadly?

**Justin** [00:06:33] Okay, so there's a literature that shows that peers are important in policing decisions. So there's a body of work by Michael Sierra-Arevalo that has shown

that line of duty injuries play a big role in how officers view their job. There's some work by a Harvard sociologist, Joscha Legewie that shows that fatal shootings of police officers lead other officers to use more force, especially against black suspects. Then there's that Adger, Ross and Sloan paper that I mentioned earlier that suggests officers learn about when to use force from their field training officers.

**Justin** [00:07:06] And then more broadly, if we view peer injuries as an exposure to violence there's several laboratory and artificial field experiments that suggest that exposure to violence can change a person's preferences, such as their willingness to engage in risky behavior or impulsivity. There's also other evidence that exposure to violence can cause people to become more violent themselves. And if we think about peer injuries as emotional events, then previous work has shown that emotional shocks in these cases, triggered by unexpected football game losses, affect violence and law enforcement in a narrow time window around the event. So, for example, there's a famous paper by Card and Dahl that shows that unexpected football game losses leads to an increase in immediate intimate partner violence. There's another paper by Munyo and Rossi that shows that violent crimes increase after unexpected soccer game losses. And then more recently, there's a paper by Eren and Mocan that shows that unexpected football game losses lead to judges assigning longer sentence lengths. So all of these results sort of suggest that we should be seeing officers respond to fewer injuries on duty. However, we're sort of the first people who look at the effect of lower level injuries on police officer decision making.

**Jennifer** [00:08:24] Yes. I mean, this is obviously a hot policy topic -- police use of force, especially the unnecessary police use of force and police killings. And there are too many police killings that wind up making the news. And so why is this so difficult to study? Why don't we know more than we do at this point about both? What determines whether police use force in any sort of incident and what to do about it? What's the hold up here? Is this mostly a data challenge or an identification challenge or both of those things?

**Justin** [00:08:53] Yeah, I would say both of those things. So there are several challenges that make this question very difficult to study. So first, you need to be able to observe police violence and injuries at the officer day or officer week level and up until recently, this was very difficult to do. So the data that we have in this study was a tremendous data gathering effort by the Invisible Institute and my coauthors, Bocar Ba and Roman Rivera. And they've actually made a lot of this data publicly available at the Citizens Police Data project. Second, in this setting, you need to be able to have some measure of an officer's network. So the police force in Chicago has about over 10,000 officers, and not all of these officers are friends with each other, and not all of them are likely to hear about the non-life-threatening injuries that other officers experience. So you need some way of knowing who is likely to hear about the injuries of which officers. And so that, you know, who is essentially going to respond to those those events.

**Justin** [00:09:50] And then even with this data, there's going to be this, you know, identification challenge. You need to worry essentially, that officer injuries and force use are co-determined. So that is any shocks to officer civilian aggression might lead to both officers being injured and officers using force. So if there is something happening in an area that increases the noncompliance rate of suspects, for example, then that can lead to more officer injuries and more suspect injuries. The reasons that are unrelated to peer injuries. And so to solve these issues ideally, you want a network of officers that were formed for reasons other than officer aggression. So that is if some officers were friends with each other because they like to get into fights, then you might find that officer injuries

and force use increase simultaneously for reasons that are unrelated to officers responding to their peer injuries.

**Justin** [00:10:44] And finally, you want officers who are geographically separated in some way so that you can make an argument that the injuries of some officers are unrelated to the job risk that their peers face. And there's a lot of other issues as well as well, but but these are sort of the most pressing issues to deal with.

**Jennifer** [00:11:02] Okay. So we will talk about the data in a moment. But you're basically going to address the identification piece with some natural experiments. And so I as I was reading this, I think of the natural experiments you have as having two components so the first is who an officer's peers are. There are lots of peers or one might be interested in, but you're going to consider the other officers they went to the police academy with. So presumably this is a group they bonded closely with during that training period. So tell us about the process of police recruitment and training in Chicago. How are applicants chosen for police academy cohorts and how are they assigned to duties after that training is complete?

**Justin** [00:11:41] Yeah. So so as you said, the network that we choose to use in the setting is officers who attended the police academy together, but then later went on to work in different areas of the city. And the reason that we choose this group is that there is an institutional structure of Chicago police recruiting that allows us to solve some of those identification issues that I mentioned earlier. So in Chicago, the recruitment process generally follows five steps. First, there's a recruitment call where the police department advertises that they want more officers and then those that are interested in becoming a police officer come in and take an entrance exam. Everybody who passes the exam is given a lottery number, and this lottery number determines if and when they're hired.

**Justin** [00:12:25] And so once their lottery number is called up, they take a battery of physical and mental tests and then they attend the police academy. And so what we have is a bunch of police academy cohorts that are quasi randomly assigned. And then once they're in the police academy, as you said, there's opportunities for these officers to bond with each other. They receive over 900 hours of basic training over six months. This training includes instructions on firearm usage, there's physical training, classroom training. They develop general policing skills. It's meant to enhance leadership abilities and develop a solid ethical foundation. And so these are all the things that they learn in the police academy with those who are randomly assigned to attend the academy with them. And then after they complete the police academy, then they move on to a probationary field training period that lasts roughly a year. And during this training period, they're assigned to different geographic districts across Chicago. And these duty assignments can change day to day.

**Justin** [00:13:29] And then after the probationary period is over, the officers move to a more permanent police unit based on the needs of the police department rather than the preferences of the officer. And so many of these units are geographically designated, not all of them, so the canine unit, the SWAT team, for example, are not and so they're going to be excluded from our setting. But officers who are moved to one of these geographically designated units are going to be included. And so in our analysis, we're going to consider an officer's peer as a member of these randomly assigned academy cohorts who now work in a different geographic unit.

**Jennifer** [00:14:06] Awesome. Yeah. So the random assignment to cohorts is especially neat here because as you said earlier, in general, it's not random who your friends with and so so here you're basically these are randomly assigned friends, randomly assigned peers. Okay. And then the second piece of your natural experiment is that something suddenly happens to one of these peers. So we know policing is a dangerous job and sometimes officers are violently injured on the job. So you use these unfortunate events to measure how peers injuries affect other officers behavior. That is the behavior of the officers they went to the academy with. So walk us through how you do this.

**Justin** [00:14:42] Sure. So so in our data, we have access to administrative records on when officers were injured and the officers stated opinion about the amount of resistance that a suspect used. So we classify injuries as those where the officer reports that they were injured during an encounter and they allege that the suspect attacked them. So this is going to include situations where the officer was physically assaulted by the suspect, but not when the officers themselves were harmed by overexertion or a fall or something like that. Then essentially what we do is we compare an officer's propensity to use force in the weeks before a peer is injured to right after the peer is injured relative to the same time path of force for those who don't have a peer injured during that time. So intuitively, you can think about this as a difference in differences design where one difference is before and after the peer is injured. And then we compare that difference to similar officers who did not have a peer injured during that time. And so our approach is going to allow us to control for time invariant individual differences in the propensity to use force, like someone's latent propensity to use force along with district-week-level-differences in the costs and benefits of using force.

**Justin** [00:16:00] So that would be shocks to noncompliance within a within a district. And so now the strategy is meant to account for the threats to identification that I mentioned earlier, but there are some downsides to it that most likely lead us to underestimate the effect of peer injuries. So first cohorts have about 40 people in them on average, and not all of these people are friends with each other. And so we're likely, including some people in our definition of peers who have no idea whether their friends got injured because they didn't form relationships while in the police academy. Second, we don't use data during that probationary period that I mentioned because we don't know where the police officers are operating on any given week. So we're dropping the period of time where these friendships may be the closest. And third, we exclude peers who are in the same geographic district because we're worried about these correlated shocks to noncompliance, but these officers are likely better friends than those who attended the academy together and are now in different districts. So we're missing that part of the story. So the approach is imperfect. It likely underestimates the effect of peer injuries, but that's the main intuition behind what we do and why we do it.

**Jennifer** [00:17:10] Okay, great. Yeah. So the punch line here, the data results you're going to talk about are likely conservative. And in the real world, they're really bigger than that.

**Justin** [00:17:19] That's right.

**Jennifer** [00:17:20] Okay. Tell us more about this cool data. What data do you have in this data set that you all have gathered over the years?

**Justin** [00:17:27] Yeah. So so the data that we use in this study comes from merging four administrative sources. So we have administrative records on about 3500 police officers,

start date and unit assignments. This data is going to allow us to understand when police officers started at the academy and the district or unit assignment where they ultimately end up. We're going to merge this to the tactical response reports that detail every time more than minor force was used, the type of force that was used, the injuries that occurred, if any, along with the stated resistance of the suspects from the perspective of the officer. So we don't have any information about the types of injuries officers sustained, but work by by Tiesman and coauthors reports that the the most common injuries to police officers are things like contusions, abrasions, lacerations, fractions or dislocations and these injuries are commonly to the hands, legs, neck, head or shoulders.

**Justin** [00:18:27] And so we have all this data from 2004 to 2016. It doesn't include the nature of the injury, but we're only going to be classifying injuries that occur when the police officer claims that the suspect attacked them. And then we're going to supplement the tactical response reports with additional data that we have on arrests made by the officer, along with formal complaints issued against the officer about a behavior that occurred in the week after a peer injury. And so these last two data sets on arrests and formal complaints, they're going to be used more to help us think about why we find the peer effects that we find. And together, the data sets are going to are going to help us understand both the direct effect and then also disentangle the motivations.

**Jennifer** [00:19:15] Okay. So in your dataset, how often are police officers injured on the job?

**Justin** [00:19:20] Yeah. So first, I should say that the Bureau of Labor Statistics reported in 2014 that police officers across the United States suffered about 2800 injuries and so this is four times higher than than the average of all other occupations. And so it's quite common for a police officer to get injured on duty. The Tiesman study that I referenced earlier reported that the most common cause of these injuries was physical violence. So it seems relatively likely that police officers face physical violence on the job. In our sample, the probability that any particular officer is injured in a given week is 0.3%. So that's relatively low, but this translates into about a 1 in 10 chance that appears injured in any given week and a 9 in 10 chance that that any officer is injured in a given week. So the probability of experiencing an event, a pure injury, is quite high in any given week. And there's also some evidence from the survey of law enforcement officers that suggests that these injuries weigh heavily on the minds of police officers.

**Justin** [00:20:27] So according to those who responded to the survey I mentioned, 42% say that they're nearly always worried about their physical safety at work. Another 42% say that they're sometimes worried about their physical safety at work, while only 16% said that they're hardly ever worried about their safety at work. So this means that officers experience relatively high injury risks while on duty, and it seems to be affecting how they view their job and their interactions with suspects.

**Jennifer** [00:20:55] So then in your analysis, when you're looking at the effects of these injuries on their peers behavior or other officers behavior, what outcomes are you most interested in?

**Justin** [00:21:05] The primary outcomes that we're interested in this study are whether an officer uses force in a given week and whether they injure a suspect in a given week. So both of these outcomes are coming from the tactical response reports that I mentioned earlier and we also look at these other outcomes to help us determine what's driving the effects on on forces and injuries. And so these are going to be things like whether the

officers injured themselves, whether they make any arrests and whether any individuals issue formal complaints against the officers. But primarily, we're interested in the officer's decision to use force and the officers when they ever they injure suspects.

**Jennifer** [00:21:45] And how common are those outcomes in a given week?

**Justin** [00:21:48] In our sample, about 2% of officers use force in any given week. Again, these are primarily controlled techniques and other types of force used without a weapon, but we do observe instances of force use with a baton, taser or a firearm they're just they're just more rare. An officer allegedly injures a suspect in about 0.6% of weeks. And again, we don't have information about the types of injuries the suspects sustained, but that's the risk that they face. And then the probability that an officer in our sample receives a complaint is about 1.6%, with most of these complaints being either for constitutional violations like excessive force, false arrests or improper search or other types of violations, like a failure to provide service. Arrests are a lot more common with an officer having about a 50% chance of making an arrest in a given week.

**Jennifer** [00:22:39] Okay. Let's dive into the results. What is the effect of a peer's injury on other officers use of force?

**Justin** [00:22:46] So in the weeks before a peer is injured, we find that there is no difference in outcomes between officers who have a peer injured and those who don't. And so this pattern tells us that officers who do not experience a peer injury are a good comparison group for officers who do experience a peer injury.

**Justin** [00:23:05] And then in the week after a peer is injured, we find that officers are about 7% more likely to use force against suspects. And this translates to about a 10% increase in the probability that a suspect is injured in that week. And so similar to other papers studying emotional responses to events, we find that the effects are concentrated in a narrow time window around the event and then dissipate soon afterwards. And so we find that this effect occurs the week after a peer is injured and then decreases again. We also find that the decrease in subsequent weeks comes from the treated group. So those who experience a peer injury returning to their lower baseline propensity level of using force and not the comparison group learning about the injuries more slowly and increasing their use of force to meet the treated group.

**Justin** [00:23:56] So to summarize, we find that police force use and suspect injuries increased by 7 and 10%, 10% respectively, and this happens immediately after the event and then officers quickly return to their baseline rates after that week.

**Jennifer** [00:24:10] Are there particular types of force that are driving that overall effect?

**Justin** [00:24:15] Yeah. So the largest increases in the number of force uses is coming from low level uses of force. So this is going to be control tactics and use of force without a weapon so strikes, punches, that type of thing. There's also a substantial increase in officers brandishing or using a firearm in the week after a former peer is injured, amounting to about a 46% increase relative to baseline. However, this represents a very small percentage point increase due to the rarity of firearm usage more generally. And so, you know, if we're assuming that officers are using force in alignment with the the police department's use of force model, the increase in force that we see that's mainly driven by these relatively lower levels of force use is consistent with officers encountering low

resistance suspects and using force in situations where they otherwise would have felt it was safe had their peer not been injured.

**Jennifer** [00:25:12] Okay. And what is the effect of appears injury on complaints against other officers?

**Justin** [00:25:19] Yeah. So in the week following a peer injury, we find that officers are about 7% more likely to engage in behavior that leads to a complaint of any type. And this 7% masked some heterogeneity across the different types of complaints that suspects make about the police officers behavior. And so, for example, we find no change in complaints for excessive force or complaints or verbal comments that the police the police make. And we also don't find any evidence that there is an increase in complaints about behavior unbecoming of an officer. However, we do find that there's increases in complaints about false arrest or improper search. Both of these increase by about 11% in the week following a peer injury. And we also find that there's an increase in complaints, citing a failure to provide service. So complaints that an officer failed to provide a service increase by about 16% in the week following appear injury. And so officers engage in behavior after peer is injured that lead to civilians being more likely to issue a formal complaint.

**Jennifer** [00:26:31] Do any of these effects vary in any interesting ways across different officers or different events or different suspects?

**Justin** [00:26:38] Yeah. So there are a couple of interesting ways that these effects vary. So first, in order to investigate the robustness of our results and make sure that we're not just uncovering a false positive, we considered two different types of specifications. So the first investigates whether our effects are underestimating the true effects of officer injuries by changing the peer group to one that has plausible less social distance.

**Justin** [00:27:00] So we want to find a group of officers using the observables in our dataset that were more likely to be friends in the police academy. And so we do this by only considering injuries to peers who are of the same race of the officer in question. So this is going to be the exact same definition as before, but now we're only considering somebody appear if they're of the same race and we're in the same police academy cohort. This is going to be because other work has shown that there's homophily in friendships. And so it's plausible that people of the same race were more likely to become friends at the police academy. And when we do this, we find that the effects are nearly double what they were when considering all classmates. So when we zoom in on a group that has plausibly lower social distance, we find that officers respond more to the peer injury.

**Justin** [00:27:51] Next, we look at the effect of officer injuries that occurred during events where the officer did not allege that a suspect attacked them. So this could be injuries like falls, slips, trips or overexertion and so these are going to be injuries that are not caused by the suspect in question. And when we do this, we find no evidence that these types of injuries affect the behavior of other officers. So this suggests that officers are really responding to their peer being injured by a suspect and not other things that might be going on contemporaneously. And these two pieces of evidence reinforce the effects that we're finding, sort of reinforces that these effects are causal and not driven by some unobserved confounder.



**Justin** [00:28:35] And then in terms of heterogeneity, one interesting thing that we find is that as officers gain more experience on the job, they learn how to avoid responding emotionally to peer injuries. So what that means is that the more experienced officers are less responsive to peer injuries in terms of using force, injuring suspects or acting in a manner that causes a civilian to issue a formal complaint, then less experienced officers. And this lines up with some work by by Ta, Lande and Suss who use body camera data to document that more experienced police officers have lower emotional reactivity to situations than less experienced police officers. And then, similar to the effect of tenure, we find that repeated exposure to injuries so remember, there's about 1 in 10 chance that somebody experiences an event in a given week repeated exposure to injuries appears attenuates the effect of these peer effects. So that means that officers respond most strongly to the first events and then less strongly to the subsequent events. And so this heterogeneity suggests that perhaps more training or other things that can induce more experience in police officers reduce the consequences of on duty injuries.

**Jennifer** [00:29:51] And finally, you consider a few possible mechanisms that might be driving all of these effects. So what do you find there?

**Justin** [00:29:58] Yeah. So. So there are several possible reasons that officers may respond to their friends getting injured. Our preferred interpretation of the data is that this is an emotional response that's triggered by a friend being injured and that can explain why we see this immediate response that quickly dissipates like Card and Dahl or Munyo and Rossi found. However, we have to be careful to rule out other plausible alternatives. So first, we want to rule out that officers are actually responding to their peers force use rather than their peers injuries. So there's a very large body of work in the peer effects literature that shows that the actions an individual takes affects the actions that their peers take.

**Justin** [00:30:38] And so it's plausible in our setting that officers use more force when they see their friends using more force, and perhaps they're not not punished for using more force. And so what we're going to do here is we're going to investigate this channel by restricting our attention to instances where the police officer, the peer police officer chose to use force against the suspect, but they were not injured during that encounter. And so when we look at the effect of officer force use that's unaccompanied by an injury, we don't see any effect on their peers behavior. And so that suggests that officers aren't just mimicking the force use of their peers. They're instead, they're responding to these these injuries of their peers, specifically injuries that are caused by the suspects in the interaction.

**Justin** [00:31:25] Second, we want to rule out social learning. So that is officers might be learning about their true injury risk through observing the experiences of their peers. Now, under this interpretation, we should see that the induced force use from a peer being injured led to a lower likelihood that officers are injured in the next week because in this case they're rationally updating about the riskiness of the situation based on the information from their peers and responding appropriately. However, we don't find any evidence that officers are made any safer by the force use in the week after a peer is injured. So it doesn't seem like this is officers rationally updating about the true risk of the world, but instead using force in situations that maybe they shouldn't because it's not making them any safer. And so having ruled out these alternative hypotheses, we also want to learn some more information about what type of emotional response officers are experiencing.

**Justin** [00:32:22] And so you might imagine that if officers are more afraid, this is leading them to interpret unsafe situations as risky situations. Or you can also imagine that they're angry that their friend got hurt and now they're taking this frustration out on civilians. And so to understand which one of these emotional responses are occurring, we're going to rely on an argument about sorting into interactions. And so if an officer is afraid, they're not only going to use more force when encountering a suspect, but they're also going to try to avoid interactions with suspects so they don't get injured themselves. If an officer is angry, then they're going to be using more force when they encounter a suspect, just like if they're afraid, but they may also seek out encounters where they can interact with suspects. And so while we can't observe the encounter decision, we can use the arrest and the complaint data to give us some more information about the motivations behind police officers.

**Justin** [00:33:18] And so overall, we find very limited or no evidence that arrests are increasing in the week after a pure injury, which suggests that there isn't officers sorting into encounters. However, if you go back to the complaint data, I can remind you that we see about a 16% increase in complaints for failure to provide service, some service that was requested by the officer. So together, these these results suggest that officers might be avoiding interactions with suspects after their friend is injured. And that's sort of more consistent with a story that peer injuries scare officers leading them to interpret safe situations as risky and therefore use force in situations where they otherwise would not have.

**Jennifer** [00:34:00] So what are the policy implications of these results? What should policymakers and practitioners take away from your study on this very hot topic?

**Justin** [00:34:09] Yeah, so broadly speaking, there's two policy implications from these results. The first is that we need to consider officer safety when designing policies that are meant to reduce the use of force. So any policy that decreases officer safety will have the unintended consequence of increasing the use of force through these peer effect externalities. In contrast, policies that increase officer safety will have the added benefit of reducing police use of force again through these peer effect externalities.

**Justin** [00:34:42] Second, we can reduce the use of force more generally by enacting policies that attenuate these peer effects. So this could mean providing more training to officers so that they have more experience. Or it could mean providing counseling and other support services to officers when their friend is injured so they don't respond in the same way that they do right now. However, I should mention that this paper is more about the mechanisms driving police force. And so we should be using randomized controlled trials or observational methods to do more thorough evaluations of the policies suggested by these findings, because there are more suggestions and definitive truth that these types of things will reduce use of force.

**Jennifer** [00:35:23] Yeah, no, that's an excellent, excellent caution, but it is really interesting, I think, and that you seem to be finding that it is this fear component that's driving this behavior rather than something else. And I think the various tests that you're doing are really compelling there. And so I completely agree that, like that leads me to think stuff like counseling could be very interesting to try. And I would love to see police departments having better access to various kinds of mental health care, especially in the wake of various types of traumatic events. As in other settings, it seems like there are a lot of reasons to think that that would be effective.

**Justin** [00:36:03] Yeah, absolutely.

**Jennifer** [00:36:05] Yeah. Have any other papers related to this topic come out since you first started working on the study?

**Justin** [00:36:10] Yeah. So there's been a lot of work in the past few years on this topic. I think I want to highlight a few papers, although there's there's a lot out there now. So there's a paper by Cho, Goncalves and Weisburst where they show that police deaths on duty lead to a decrease in arrests, and they attribute this decrease to fear of job risk. Kind of like the same mechanism that we have in mind here officers are afraid and so they avoid interactions with suspects. My coauthor, Roman Rivera, has a new paper using the same Chicago data, which shows that having more minority coworkers reduces in officers arrest, propensity through reducing low level arrests of minorities. And then, broadly speaking, in more work to try to unpack the policing production function Bocar Ba, Roman Rivera and others have done a lot of good work trying to understand how police officer assignment oversight and minority representation of police officers affects policing outcomes more broadly. And then there's, you know, Andrew Jordan and Taeho Kim have a nice paper on when, you know, police oversight can and cannot be expected to improve officer behavior. And then there's some more new work by Mackenzie Alston and Emily Owens that have started to use laboratory experiments to try to understand how individuals respond to policing. And so this is all to say that there's been a lot of work in this area since it since we started on this paper, but there's still a lot more work to do if we want to try to fully understand the policing production function and what types of policies can improve policing outcomes.

**Jennifer** [00:37:42] Yeah, this is such an exciting space to be following right now. So what's the research frontier? What are the next big questions in this area that you and others are going to be thinking about going forward?

**Justin** [00:37:53] Yes, I think the open questions here are pretty clear. We need to understand the important determinants of police force and misconduct more generally. And we need to understand the mapping of these determinants to the outputs. I think the the frontier in this area is going to be bringing in new data to allow us to answer these questions. So that means creating new administrative datasets like Bocar and Roman are doing with the Citizens Police Data project. But I think we also need to do a lot more field experiments like those by Barak Ariel and coauthors, where we use experimentation to figure out what kinds of interventions work at the police department level and then how to scale up those interventions nationally.

**Justin** [00:38:34] I also think that we need to leverage the data that's already being generated by private companies. And so what I mean by that, for example, is I have a working paper with John List, Alec Brandon, Pradhi Aggarwal and other people that uses Lyft rideshare data to try to understand how police officers use a suspect's race when deciding whether to stop motorists. So in traditional data sets, like the administrative data that we can get from police departments, we only observe encounters that escalated to the point where a citation is used. And this can lead us to underestimating discrimination if officers choose to only arrest the worst offending white drivers and a mix of high and low offending minority drivers. And so we can get around this problem in the Lyft data by using the high frequency location data that comes from a drivers phone whenever they're using the app.

**Justin** [00:39:26] And this data is important, they're going to tell us who is driving at any given moment, particularly those who are choosing to speed, but the police choose not to stop. And so it's also going to allow us to measure precisely how fast every driver is going at any given moment. And so we can also control for how much each driver is breaking the law, if at all. And what we find in this case is that police officers choose to stop minority drivers, about 30% more than white drivers who are traveling the exact same speed in the same place and this result is a lot of important policy implications. But it wouldn't be possible without partnering with a private company like Lyft. And so this study highlights how private data can be used to evaluate social issues like policing. And we should be trying to leverage more of these data sets to understand these types of questions. So to summarize, I think that we need to do more to understand the policing production function, and we need to do this by leveraging or generating new datasets.

**Jennifer** [00:40:28] Love it. Yeah, that last study, it's super interesting. And also makes me wonder how many defense attorneys are going to be trying to subpoena data like that to defend their clients. It just seems like so useful in court to be able to show you didn't stop this other guy and you stopped my client. I will add just one thing, one thing to your last you mentioned, you know, running actual field experiments, which of course, we should definitely be doing. But my sense is that a lot of police departments are already out there trying lots of stuff and, you know, rolling out new training programs or existing training programs in ways that are naturally capacity constrained and take time or something like that. And so some researchers are taking advantage of situations like that, but I get the sense there are a lot more natural experiments like that out there that have yet to be exploited. So if there are grad students listening, it's a great, great opportunity to just go start talking to different agencies and hear about what they're doing and see if there is if we can learn from it. You know, we've got 12,000 police departments in this country. They're all doing different stuff. That should lend itself to learning things pretty quickly. So this area is just fascinating and really exciting to watch. My guest today has been Justin Holtz from the University of Chicago and moving soon to the University of Michigan. Justin, thank you so much for talking with me.

**Justin** [00:41:47] Thank you for having me.

**Jennifer** [00:41:53] You can find links to all the research we discussed today on our website [probablecausation.com](http://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 and other contributors. Probable Causation is produced by Doleac Initiatives, a 501(c)3 nonprofit, so all contributions are tax deductible. If you enjoy the podcast, please consider supporting us via Patreon on or with a one time donation on our website. Please also consider leaving us a rating and review on Apple Podcasts this helps others find the show, which we very much appreciate. Our sound engineer is Jon Keur with production assistance from Nefertari Elshiekh. 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.