Probable Causation, Episode 45: Conrad Miller

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 [00:00:18] My guest this week is Conrad Miller. Conrad is an Assistant Professor of Economic Analysis and Policy at UC Berkeley's Haas School of Business. Conrad, welcome to the show.

Conrad [00:00:28] Thanks for having me.

Jennifer [00:00:29] Today, we're going to talk about your research on how the racial composition of a place affects punishment severity. But before we get into that, could you tell us about your research expertise and how you became interested in this topic?

Conrad [00:00:42] Sure. I'm a labor economist, so half of my research focuses on how workers sort across firms and the role that firms play in generating income inequality, particularly focused on racial and gender inequality. The other half of my research looks at discrimination and racial disparities in other contexts. So that includes the criminal justice realm. I first started thinking about racial disparities in criminal justice outcomes as a research topic in graduate school. So actually, my coauthor and grad school classmate, Ben Feigenberg, and I were looking to work on something and we were originally interested in how the increase in incarceration rates in the United States in the 80s and 90s affected local labor markets, including racial inequality in those markets. To study that, we wanted to use variation across counties in how punitive their local courts were in order to construct an instrument for studying the causal effects of incarceration policy. But then in the course of trying to do that, we realized a couple of things. First, actually, just getting the data to execute that was quite challenging, much more challenging than we anticipated. And second, just documenting how local courts varied and how punitive they were and what explains that variation was actually pretty interesting in itself.

Jennifer [00:02:00] So your paper is titled, "Racial Divisions and Criminal Justice: Evidence from Southern State Courts." As you mentioned, it's coauthored with Ben Feigenberg, and it's forthcoming in AEJ Policy. And so some background here, there is, of course, a long running conversation about racial disparities in criminal justice outcomes, particularly incarceration in the US. And a lot of that conversation focuses on how much racial bias matters, that is, to what extent are otherwise similar Black and white people treated differently by law enforcement and the courts in a given place. But in this paper, you're interested in a different way that racial bias might be driving the racial disparities that we see in the data. So what is the hypothesis that you and Ben had in mind?

Conrad [00:02:41] Yeah, let me first take a step back and say a little bit about motivation here. So there's a there's a broad social science literature, mostly outside of economics, that argues that crime and criminal justice in the US is and has been a fundamentally racialized political topic. So in particular, many have argued that race and racial politics play a central role in explaining why the US has such high incarceration rates relative to the rest of the world. So, as you say, there's been this- a lot of focus on documenting and understanding racial disparities in various stages of the criminal justice process so that that would include arrests, charging decisions, sentencing. And these disparities are quite large and striking. Yet if we if we just look at incarceration rates for white Americans, actually, that rate alone would rank exceptionally high among peer nations. So if race

plays a central role in explaining why US incarceration rates are so high, we actually need a mechanism that can explain why the incarceration rate for white Americans is also exceptionally high.

Conrad [00:03:49] Now our basic hypothesis that we're interested in is the places that are racially diverse, which would include the United States, those places are more punitive for everyone. And the mechanism we have in mind is that voters both shape criminal justice policy and they exhibit what's called ingroup bias in their preferences over policy. What we mean by ingroup bias in this context is that voters on average support a more lenient system when they perceive that those likely to face punishment are similar to them on the basis of race. But those same voters may actually support a more punitive system when they believe those likely to face punishment belong to some outgroup—so a different racial group, for example. Now returning to the question of racial disparities, a side effect of this mechanism we have in mind is that one factor that may be contributing to racial disparities in criminal justice outcomes in the US is actually that Black people are concentrated in racially diverse places where criminal justice systems are likely to be more punitive for everyone.

Jennifer [00:04:54] And so before this study, what did we know about how racial composition and ingroup bias affect local preferences and policies?

Conrad [00:05:02] We knew a fair amount, though I'd say a lot of this prior research looks at support for redistribution and public goods provision rather than criminal justice policy per se. So, for example, there are a lot of papers or several papers showing that in surveys, people express more support for redistribution and redistributive policies when they perceive that the beneficiaries of that redistribution belong to their own racial group. There are also survey experiments showing that just priming respondents to think about diversity, either immigration or racial diversity, just priming them to make them consider those issues, makes respondents less supportive of redistributive policies, which suggests that diversity itself can actually shift people's policy preferences in this space. There's also evidence that spending on public goods, including education, is lower in areas that are more racially heterogeneous, more diverse.

Conrad [00:06:03] Now, particularly, there's some evidence on preferences over criminal justice policy. So, for example, when white survey participants are randomly shown information about racial disparities in punishment, there's evidence that they become more supportive of punitive policies like three strikes laws. But what I think is less clear in this literature is what implications these survey patterns have for policy outcomes. So even if people's preferences over criminal justice policy exhibit ingroup bias, it's not really clear how that bias will translate into actual outcomes. So, for example, into outcomes like incarceration rates.

Jennifer [00:06:39] Right. And so you're going to look at those actual outcomes. But this is, of course, a tough question to study for a few reasons. So what are the hurdles that researchers like yourselves have to overcome to measure the causal effects of racial composition on outcomes like punishment severity?

Conrad [00:06:55] Yeah, there are several challenges here. Our our- really our fundamental question here is kind of a big question. We want to understand does racial heterogeneity in the US contribute to US exceptionalism in its punitiveness? And I would say there are three key challenges we face to credibly answering this question. So first, a research question is fundamentally motivated by a cross-country comparison. We want to

understand why the US relative to other countries has such a punitive criminal justice system. But for a lot of reasons, a cross-country comparison could be very hard to do here. So, for example, we would need to collect data from various countries, somehow harmonize that data, which would include harmonizing across distinct criminal codes. And really that that wasn't feasible for us to do for a large enough set of countries to actually do that kind of analysis. So we take a different approach. Instead, what we do is we exploit variation within the US, across jurisdictions, in criminal justice systems and criminal justice policies. This is feasible as a strategy because judges and prosecutors are going to have significant discretion in determining punishment in the US. And moreover, in the US they're typically locally elected. So there's good reason to think that local punishment practices are actually going to be tied to local electorate preferences. So the question we're going to ask instead is within the US, are more racially heterogeneous jurisdictions more punitive? And you should think of the typical jurisdiction here as a county.

Conrad [00:08:33] A second key challenge for us is is going to be that it's potentially difficult to isolate differences in punishment practices across jurisdictions from differences in crime. So if I if I measure, say, incarceration rates for different counties, we need to figure out whether those differences in incarceration rates are actually driven by how similar offenses are treated versus differences in the composition of offenses between those places. So in this paper, we're trying to estimate what we call "punishment severity measures," and by that we mean variation in how different jurisdictions treat equivalent cases. So in other words, we want to measure what's the causal effect of jurisdiction on arrest charge outcomes. To do this, we're going to exploit relatively rich administrative court data where we're going to observe a lot about cases and their defendants. And these rich controls are going to allow us to credibly compare like with like cases across jurisdictions. Another thing we're going to take advantage of— and that we'll talk more about later—is that in these data, we can actually identify the same defendant arrested in multiple jurisdictions. So one kind of comparison we can do is compare the same person arrested in multiple places and see how their their charge outcomes differ across places.

Conrad [00:09:52] The third challenge that we face is that ultimately what we want to know is whether racial composition in a local area plays a causal role in explaining local punishment severity. That's going to be hard to figure out because even if we can compare punishment severity, say, between racially diverse and homogenous places, we expect those places to differ in other critical ways outside of racial composition. So, for example, they may have different crime rates. So if we find that punishment severity differs between those places, it's going to be hard to pinpoint exactly what the cause is.

Jennifer [00:10:27] All right. So let's talk about how you and Ben approach this. What is the ideal experiment that a researcher might like to run to answer this question? And how does your approach approximate that experiment? How do you think about the intuition of what you're doing here?

Conrad [00:10:42] So I would divide our paper into two parts, and I think each part has its own research design and really ideal experiment we're trying to approximate. So in the first part, we want to measure these punishments severity. And that, again, is that the causal effect of jurisdiction on arrest outcomes. So in the ideal experiment, you know, one way to think about it would be we would randomize the local system where an alleged criminal offense is prosecuted. So some crime is committed or some crime is alleged, and let's say we could randomize what court system that case moves through. That would be an ideal way we could actually measure how jurisdictions vary in their punishment severity. We're going to try to approximate this ideal in two ways. So one approach we'll take is to use our

pretty rich data on case and defendant characteristics to just compare similar cases. So, for example, you know, we'll be able to compare cannabis drug possession charges for similar amounts of cannabis. In different jurisdictions, we'll have detailed data on the specific arrest charge and say we could make that comparison for two cases where in both cases the defendant has no prior criminal history. And we can ask, are those two sets of cases leading to systematically different outcomes in different jurisdictions? We'll take that comparison as a way of comparing punishment severity across places.

Conrad [00:12:10] The second thing we're going to do as a way of really validating that other that approach where we're just really controlling for detailed characteristics is using the fact that we actually observe defendants that are arrested in multiple jurisdictions. And so we can use that to really see how the same person is treated in different places. Now, in the second part of the paper, we want to understand how racial composition of a jurisdiction contributes to variation in punitiveness across local courts. Now, the ideal experiment here is a little harder to imagine. In some sense, what we'd want is random variation in the racial composition of a local area. Now, clearly, we're not going have a great way of approximating this. Instead, what we're going to do is going to take two approaches, really.

Conrad [00:12:58] So first, we're just going to regress our measurement of punishment severity on local racial composition while controlling directly for other characteristics that we might think of as key determinants of punishment severity. So things like population density, crime rates, and income. But the other feature of our setting we're going to rely on is that our our theory for what's going on here is going to have particular prediction for the relationship between racial composition and punishment severity. In particular, we're going to predict a non-linear relationship between, say, punishment severity and the Black share of the population. So we're going to predict something that's increasing for some range—so punishment severity is increasing in the Black share of the population at low rates in terms of Black share. But then at some point we expect that relationship to be declining for reasons we'll get to. And so we're going to rely on this nonlinear prediction as another way of providing credible evidence for our proposed mechanism.

Jennifer [00:14:01] All right. So let's talk about this very cool data that you have in this paper. I definitely sent some emails to some grad students as I was reading this closely, saying, "They got this data, you should go look for it." So you use data from four southern states. So which states were those and what information were they able to provide?

Conrad [00:14:19] So we use data from Alabama, North Carolina, Texas, and Virginia. And I should say that we focus on the South in large part because those are states that have substantial variation in racial composition across counties. So particular, they're going to be counties in our states with large Black majorities. And you're really- you're not going to find counties like that outside of the south. So the data we have from Alabama, North Carolina, and Virginia come from their administrative Office of Courts, while the Texas data come from the Department of Public Safety. And what's important in the kind of the common elements of the data across our states here is that these data track criminal cases in those four states from arrest through sentencing. So for each case, we know the associated arrest charges. We know what charges the defendant is actually convicted of, if any. And then we know the associated sentences for those convictions. Another nice feature here is that we have defendant identifiers that are going to allow us to match multiple cases to the same defendant. And another feature of the data that turns out to be actually pretty important is that they include charges that are later dropped or dismissed. So it turns out that jurisdictions vary a decent amount in how often charges are dropped

and dismissed. And that actually is an important margin for a variation in punishment severity across places.

Jennifer [00:15:43] So as you mentioned, your first goal in this paper is to document the variation in punishment severity across places within each state. So how do you use all this data you've collected to measure punishment severity?

Conrad [00:15:56] Yeah, so the metric we focus on in the paper is whether a criminal charge leads to a jail or prison sentence. So we call the share of charges that lead to an incarceration sentence, the confinement rate. So essentially, we're going to measure confinement rates across jurisdictions. And our punishment severity measure is just a locale's confinement rate, holding other charge characteristics fixed, as well as defendant characteristics. So our baseline strategy for for estimating punishment severity is we're going to essentially take a selection on observables approach. So for each charge or case, we can regress an indicator for whether that charge led to a confinement sentence. On a rich set of charge and defendant characteristics, say, including the specific offense the defendant was arrested for, a defendant's criminal history, etc. But we'll also include fixed effects for each jurisdiction. So what we're ultimately doing there is we're controlling for these other characteristics in terms of arrest charge and defended criminal history and then we're asking, holding those characteristics fixed, what would a given jurisdiction's confinement rate look like, given the composition of charges that are flowing through the state? And those jurisdiction fixed effect measures, that's going to form the basis for our punishment severity measure. And we also look in the paper at alternative measures of punishment like conviction rates or sentence length, and we find similar results. So our results are not particularly sensitive to what outcome we focus on when we're measuring punishment severity.

Jennifer [00:17:34] And so the basic intuition here is, you know, you kind of imagine having the same- like an identical person commits the same offense in one jurisdiction or another. And the first one, it just happens to be harsher. And the likelihood that any offense leads to incarceration is higher. And so the idea then is that person is more likely to be incarcerated. And so what you're really assuming here, I think, is that, you know, you're able to see everything that matters about these different cases. Right. And that's why that's- the really rich data is so important. And we'll talk more about all the checks that you do to convince yourselves that that's that that's holding. But that's basically the intuition, is that right? It's that conditional on all the stuff you can observe in the data—the type of charge, the other characteristic of the person's criminal history, if an arrest was made in one jurisdiction or another—those cases are are the same, is that, right?

Conrad [00:18:26] Yeah, that's exactly right. And in other contexts, we might be skeptical as to whether we can really observe enough about different cases or different contexts to say, you know, we've really held everything else relevant fixed. This is actually a context where I would argue that you could credibly make that claim. So in particular, across states we'll have hundreds of specific offense charges in terms of codes for specific offenses. So we can actually get pretty granular in terms of what someone has been arrested for. And once you have those kind of granular measures, it becomes a lot more, I think, credible to argue that you're actually comparing pretty similar cases across places. But we'll have other we'll have other approaches to, I think, even further support and corroborate that basic approach.

Jennifer [00:19:18] Yeah, as listeners might imagine, this is like the meat of the paper is checking all of these assumptions. All right, so so what do you find here in terms of how

much local jurisdictions matter? How much of the variation in confinement rates is explained by the location of the case?

Conrad [00:19:34] So we find that local jurisdictions matter a lot. So one way to summarize how much they matter is that within a state, a defendant charged in the top quartile of jurisdictions by punishment severity is two to four times more likely to be incarcerated for a given charge than the same defendant charged in a jurisdiction in the bottom quartile. When we look at variation in confinement rates across jurisdictions, we find that 80 to 90 percent of that variation is explained by punishment severity. In other words, it's explained by the causal effect of jurisdiction. So another way of putting that is nearly all the differences we see in confinement rates across jurisdictions can be explained by how jurisdictions treat equivalent cases rather than, say, differences in the composition of cases across jurisdictions.

Jennifer [00:20:23] And then did those measures vary across Black and white defendants? Or does place matter more than race in a particular case that's in front of a court?

Conrad [00:20:32] Yeah, this is an interesting question. So one thing we found that I think we didn't necessarily anticipate is that these punishment severity measures we construct are actually very highly correlated for Black and white defendants. So there's certainly racial differences in confinement rates within jurisdictions. That's a that's a feature that's certainly been documented in other contexts. In our case, across states, we find gaps in confinement rates conditional on other charge and defendant characteristics on the order of 20 percent in terms of confinement rates. But places that are more punitive for Black defendants are also more punitive for white defendants. So place matters a lot above and beyond racial differences in outcomes within within places.

Jennifer [00:21:19] So as we were just discussing, a big challenge in all of this is that you're using cross-sectional data to create these punishment severity measures, so you're comparing places with different confinement rates rather than maybe changes in confinement rates over time due to some policy change or other natural experiment. So this means that we might worry about whether differences in confinement rates are driven by something else. Maybe a larger share of arrests lead to incarceration because the police officers had a higher bar for making an arrest in the in the first place. Or maybe the composition of offenders is different in a way that makes selection bias a problem. So you and Ben run a bunch of checks to convince yourselves that your measure of punishment severity is really a good proxy for the harshness of the criminal justice system in place. So walk us through some of those checks and why they're useful.

Conrad [00:22:07] Yes. So, as you say, really a key concern for us in constructing this punishment severity measure is that we really do want to be comparing equivalent cases across jurisdictions. We don't want to be making really unfair comparisons in the sense that we're comparing different types of cases in different places. So we do we do several things, as you say, in the paper. I really think the most compelling evidence we have that the punishment severity measure we're constructing actually reflect the causal effect of jurisdiction on outcomes. Is this evidence we have from looking at defendants that are arrested multiple times and in different jurisdictions? So you can think of our design here as following in this recent literature that uses what's now known as a mover design, where the goal with this kind of design is to measure the causal effect of location on individual outcomes. And the general approach here is to identify those causal effects by looking at how individual outcomes when people essentially move across locations. So researchers

have used these kinds of methods to study a variety of topics now, including the effects of teachers on test scores, neighborhoods on intergenerational mobility, and placed on medical spending.

Conrad [00:23:31] But the basic idea for us here is, as you as you mentioned earlier, ideally what we want to be doing is we want to say for the same defendant really committing the same type of offense or having the same offense allegation in different places, how is that same kind of case treated differently in different places? Now we can try and approximate that just by, say, controlling for defendant's criminal history and their underlying arrest charge. But we can actually do potentially even better than that in the sense that we can actually look at literally the same person. And see how when they're arrested in different places, how their outcomes differ. So there, what's what's nice about that comparison, is we can essentially net out all of the fixed characteristics about that person that might contribute to their charge outcomes. So, for example, you can net out differences in their perceived guilt or their income at some level and really compare a very kind of similar case in the sense that we're actually looking at the same person.

Conrad [00:24:40] So what we find is that if we take our punishment severity measure using our baseline approach where we're just controlling for observable case and defendant characteristics, we ask how does that baseline punishment severity measure predict changes in charge outcomes for the same defendant that we see arrested in multiple places? So we ask if you're arrested in place A versus place B, how does your charge outcome differ between those two cases? And we compare that to what we would predict your change in outcomes would be, given our punishments severity measure. And what we find from that exercise is that our punishment severity measure provides an unbiased measure of a person's changes in charge outcomes across cases. So in other words, if we predict that you are 25 percent more likely to be incarcerated in jurisdiction A versus B, then when we compare people that are charged in jurisdiction A versus B, we actually find for that same person a difference in their incarceration rates across cases pretty similar to 25 percent. So that gives us confidence that this this punishment severity measure we're constructing actually does reflect the causal effect of jurisdiction for equivalent cases.

Jennifer [00:26:02] All right. So once you've documented the variation in punishment severity, you then move on to the second part of the paper where you test to what extent racial composition explains that variation. And in particular, you're interested in whether a model of ingroup bias can explain the differences in punishment severity. You talked about this a little bit up front, but tell us more about what relationship you expect to see in the data and how you test for it.

Conrad [00:26:28] Yes. So the basic model we have in mind here is this. So voters have preferences over what punishment severity they want to see in their local courts. Do they want to see a lenient local court system or punitive local court system? And we assume that voter preferences exhibit ingroup bias. And as we talked about earlier, this is supported by evidence across various disciplines. So what we're thinking of is when a voter considers what level of punishment severity they want for the local courts, they think about who is likely to be a defendant in their context, in their local context. So when voters perceive that defendants are likely to belong to some outgroup, say, different racial group, they're going to prefer more severe punishment. While on the other hand, if a voter thinks the typical defendant belongs to the ingroup, they're going to prefer a more lenient approach. What that model is going to predict is an inverted U-shaped relationship between the Black share of the population and punishment severity. So in particular here,

we're thinking of a case where there's two racial groups, you know of Black and white, and white voters are going to prefer more punitive policy as the Black share of defendants increases. And you can think of you can think of kind of the opposite relationship for Black voters.

Conrad [00:27:52] But for jurisdictions with sufficiently large Black populations, the pivotal voter is actually going to be Black. So for smaller Black shares, we'd expect this increasing relationship between the Black share defendants and punishment severity because there, the pivotal voter is more likely to be white. But on the other hand, for sufficiently large Black populations, the pivotal voter is more likely to be Black. And in that case, you'd expect punishment severity to be declining in the Black share of defendants. So ultimately, we take this this this model, which gives us this prediction for an inverted U-shaped relationship between the Black share of the population and punishment severity and we take that to the data.

Jennifer [00:28:35] And what do you find when you run that analysis?

Conrad [00:28:38] So consistent with this model we really we find this striking inverted U-shaped relationship between punishment severity and the Black share of the population. So in particular, what does that mean? It means that jurisdictions with very large white majorities are relatively lenient, as well as jurisdictions with very large Black majorities, while it's the more diverse places in between that are substantially more punitive. So one way of quantifying how important this inverted U-shaped relationship is is that we find that within states, defendants are 27 percent to 54 percent more likely to be incarcerated if they are arrested in these peak heterogeneity jurisdictions. So by peak here, I mean at the peak of the inverted U-shape. So they're more likely to be incarcerated in those jurisdictions relative to those arrested in homogenous jurisdictions. And as we talked about earlier, that same pattern actually holds for both Black and white defendants. So you don't want to be a Black or white defendant arrested in these relatively diverse jurisdictions.

Jennifer [00:29:44] OK, so you're finding that punishment severity follows this inverted U-shape with the racial composition of the place. You also test for whether the Black-white gap in punishment severity follows the same pattern. So what was the question you're trying to answer when you do that and what do you find?

Conrad [00:30:00] Yeah, so in our in our basic model, it's quite restrictive in the sense that we've assumed that voters can affect how punitive their local court for all defendants. So even if voters prefer to punish outgroup defendants, but not ingroup defendants, we've assumed in our in our basic model that voters can really only choose the punishment severely for all defendants. But that assumption we're making you can think of as optimistic in the sense that maybe voters are actually able to influence not just the overall level of punishment severity in their jurisdiction, but also the level of racial disparities in punishment severity. So in other words, maybe a racially biased electorate can institute racially biased punishment severity. And I think in principle, this seems somewhat plausible as the potential influence the electorate could have on local policy. But interestingly, what we find is that racial gaps in punishment severity are not systematically related to the racial composition of the local population. So we find that there are racial gaps in outcomes, but the magnitude of those racial gaps do not seem to be related to racial composition in the same way that overall levels of punishment are related to racial composition. So it's not as though more racially diverse jurisdictions also have larger racial gaps in confinement rates.

Jennifer [00:31:24] So again, your main hypothesis is that local racial composition affects punishment severity through the preferences of this local electorate. So another test that you run that I really liked to support this is that you look at whether support for statewide ballot measures related to punishment has the same inverse U-shaped relationship with the Black share of the population. So tell us a little bit about that and what you find there.

Conrad [00:31:49] Sure. So, as you say, the underlying mechanism we have in mind for what drives this relationship between local racial composition and punishment severity is that it's driven by the preferences of the local electorate, that the local electorate can affect policy, say, by voting for their local judges or prosecutors. I think one shortcoming of our main analysis is that we don't actually have direct data on that mechanism. We're not necessarily looking at the voting patterns of the local electorate in these different jurisdictions. So one way we try and get at this in a supplementary analysis is we look at how local voters vote on statewide ballot initiatives related to criminal justice policies. So, for example, one kind of ballot initiative here would be what kinds of defendant rights should there be in a given state? So we can code up how people vote on those ballot initiatives as either relatively punitive or relatively lenient. I should say, we're using data here on ballot initiatives collected by Claire Lim, Jim Snyder and David Stromberg in another paper that they generously provided to us. And so what we could do in our four states is we can actually compare our punishment severity measure to how voters in a given jurisdiction vote on these ballot initiatives. And so what we find is that both- we find two things.

Conrad [00:33:23] So first, we find that indeed, places where we measure that local courts are punitive are also places where the local population tends to vote on ballot- these ballot initiatives in a way that indicate preferences for punitive local courts. And we also find this inverted U-shaped relationship between racial composition and local preferences, where instead of using our punishments severity measure, we just look at how people vote on these ballot initiatives. Now, that exercise is is more noisy with these ballot initiatives than what we can do with the punishment severity measures, but you see generally the same pattern. Now one thing that I think is important to keep in mind with this exercise is that these ballot initiatives are not actually about what people want to see in their local courts. This is about state wide policy. And the mechanism we really have in mind is electorate preferences for local courts. So it's not the exercise is not exactly what we would want, but we think it kind of gets part of the way there.

Jennifer [00:34:31] Absolutely. OK, so what are the policy implications of your results? What should policymakers and advocates who are concerned about racial disparities in the criminal justice system take from all this as they try to find ways to reduce those disparities?

Conrad [00:34:44] Unfortunately, I don't think the policy implications of our work are particularly clear, and that's probably true of a lot of work studying the effects of racial heterogeneity in policy. But at a high level, I think our work highlights that racial divisions can affect not just disparities within jurisdictions, but also disparities between jurisdictions. So those interested in racial disparities should be monitoring those between jurisdiction differences and how those differences may be influenced by race. Our work also highlights perhaps unanticipated side effect of racial divisions in society broadly. So in the same way those divisions shape our spending on the social safety net, they can also shape incarceration rates. And as a consequence, policies that somehow improve cohesion may also influence incarceration rates by affecting voter preferences for criminal justice policy. Our work doesn't say anything about what those policies would look like, but it does

suggest the consequence of such policies. I think our work finally highlights a potential downside of having important and powerful criminal justice actors as elected officials as we do in the United States. So it's likely the case that this setup contributes to the fact that we have such dramatic variation in how local courts operate under essentially the same criminal code within a state as we measure in our study.

Jennifer [00:36:11] So moving beyond your paper here, um I know it's been out for a little while, have any other papers related to this topic come out since you and Ben first started working on this study?

Conrad [00:36:23] Yeah, I think there are a few papers that have come out that I wouldn't say are directly in the space we're operating, but it definitely, I think, shaped how I interpret our paper and this sort of general line of research. So, for example, there have been a couple of recent job market papers actually looking at how media coverage of crime, local media in particular, affects police behavior and some indirect measures of voter beliefs. So in particular here, I'm thinking of a recent paper by Nicola Mastrorocco and Arianna Ornaghi. And a job market paper by Jonathan Moreno-Medina. And both of these papers are using these pretty clever research designs, trying to isolate variation in local crime coverage and looking at local crime coverage effects, sort of how the local criminal justice system operates, in particular, focused on police behavior and housing market prices. What I like about this research and how I think it connects our work is that it really, in a more granular way, I think considers one mechanism that we have in the back of our mind, which is that these local, very local preferences and local beliefs that voters have actually influence local outcomes in a meaningful way in terms of how various criminal justice actors operate.

Conrad [00:37:51] Another, I think related work here is the job market paper of Ellora Derenoncourt, who is my colleague here at UC Berkeley and I believe was a guest on your podcast. So she has this amazing job market paper looking at how- looking at the effects of the Great Migration on places, essentially. And she's looking at primarily intergenerational mobility and how black in migration to Great Migration destination cities affects intergenerational mobility in those areas. One mechanism she has in mind is that black in migration could affect how local public goods and kind of local public spending is distributed. So in particular, she she finds that police spending increases in these areas where black in migration increases. And I think you can you can think of that as related to the mechanism we have in mind here, where, by contrast, what we have she actually has, in some sense, a natural experiment that generates quasi experimental variation in the racial composition of different jurisdictions. So she's not looking at a punishment severity per se. Instead, she's looking at spending on state police. But potentially those those things are quite related.

Conrad [00:39:15] And then I think there's there's a series of papers that have document patterns that I think are consistent with our findings, although they're not exactly framed as such. So there are a few papers documenting racial disparities in courts and in policing behavior. So here I'm thinking of Rehavi and Starr's paper in the Journal of Political Economy; Goncalves and Mello, their recent AER paper; and Raphael and Rozo, who have a recent paper in the Journal of Labor Economics. And what's similar across these three papers is they're interested in documenting racial disparities in different aspects of the criminal justice system. So in prosecutor behavior and sentencing. In how police give speeding citations. And then in the Raphael and Rozo paper, thinking about how police decide whether to book or simply release juvenile defendants that they arrest. And what those three papers find is that—they're all focused on racial disparities in these outcomes,

so they document racial disparities—and one thing they do is they show actually that including jurisdiction fixed effect, so essentially controlling for location, in all three of those papers actually reduces racial disparities and outcomes quite substantially. And in the context of those papers, they're primarily doing this just to try and make cases across racial groups as comparable as possible. But the fact that including jurisdiction fixed effects, these controls, actually affects disparities so dramatically really suggests that the Black population and in some contexts Hispanic populations are concentrated in places that are generally more punitive. So I think they're documenting a pattern that we're really focused on in our paper, understanding why that pattern across jurisdictions exists.

Jennifer [00:41:10] Your paper also reminds me of a paper by Alex Albright, who's a grad student at Harvard, and she looks at the implementation of risk assessment tools in pretrial detention decisions, so on bail decisions, basically finds that there is this policy change in Kentucky that implemented these tools and it seemed to increase racial disparities in outcomes. And she finds in her paper that that's in large part due to different take up rates across jurisdictions. So it's like it's very much in line with this story that you're telling in this paper, that it's about people- Black and white defendants in a particular county aren't necessarily treated differently, but it happens to be the whiter counties that adopt this tool and become more lenient. And so when you look statewide, the racial disparities increase, which is yeah not not obvious at first glance. And if you were just controling for jurisdiction, it would wash out all that effect. So super interesting.

Conrad [00:42:09] Yeah, exactly. I have I have not seen that paper I will look out for it. But that's that's exactly the kind of- yeah, that sounds very consistent with what we're finding.

Jennifer [00:42:18] Yeah. So what's the research frontier here? What are the next big questions in this area that you and others will be thinking about going forward?

Conrad [00:42:26] So I can tell you about another project that I'm working on that I think of as as a follow up in this literature. So this is a project is actually joint with Ellora and Ben Feigenberg, my coauthor, as well as Heather Sarsons, where we're essentially trying to take lessons from Ellora's job market paper and our paper and trying to combine them. So as I said, I think a shortcoming of our paper in particular relative to Ellora's work is we really don't have a clean source of variation across areas in terms of racial composition. So we can make the case that the fact that we see this striking non-linear relationship between racial composition and punishment severity reflects the fact that racial composition itself is contributing to variation in punishment severity.

Conrad [00:43:22] But I think the thing that kind of lurks in the background here is that you might worry that there's just other differences across these places that are also contributing to variation in punishment severity that we're not capturing. So what we're trying to do in this follow up work is using the research design from Ellora's paper that measures variation in migration flows coming through the coming from the Great Migration and combining that with our punishment severity measures in this paper. So the idea would be for a different set of states, in particular states that received inflows of migrants in the Great Migration, construct punishment severity measures for those locations and ask whether places that received larger inflows of Black migrants from the Great Migration due to plausibly exogenous push factors and migration flow patterns, see whether that variation in migration actually seems to affect local punishment severity measures. So kind of using a better or cleaner source of variation in terms of racial composition.

Conrad [00:44:28] I mean, I also think another really interesting area of research here is trying to better understand where do preferences, voter preferences come from for these types of policies? So I'm thinking of here work by Stephanie Stantcheva and her many coauthors, where they've run a series of large scale online surveys and survey experiments, where they ask things like, how do your perceptions of the share of immigrants in your population affect your preferences for redistributive policies, for example. And and above and beyond that actually, shocking people's beliefs about local immigrant shares by providing- by randomly providing information about that and asking how that affects their support for different policies, I think that kind of approach with these kind of large scale surveys and looking at variation across areas, I think that kind of research, looking at criminal justice policy in particular, would be a really exciting thing to do to get a better handle of understanding what actually shapes people's preferences for these types of these types of policies.

Jennifer [00:45:44] My guest today has been Conrad Miller from UC Berkeley. Conrad, thanks so much for talking with me.

Conrad [00:45:50] Thanks, Jennifer.

Jennifer [00:45:51] 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 John 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.