Prepared Remarks for Guggenheim Symposium
Sonja Starr

Thanks to organizers; it's an honor to be here.

Anne's been a strong proponent of the expanded use of risk assessment in the criminal justice system, and I've been a vocal critic, though we've had a chance to talk about these issues in recent months and as it turns out we've discovered a lot of common ground. We both believe in using data to improve the practice of criminal justice, and we both care about equality concerns that arise when people are treated differently based on their characteristics. Anne's organization has been working to develop risk assessment instruments that try to address some of these equality concerns, which is a positive development, and she'll tell you a bit more about that shortly.

But my focus, meanwhile, is on calling the attention of the legal community and the media and the public to the very, very serious problems that exist with almost all of the risk assessment instruments that are already in widespread use in criminal justice systems around the country. We are already subjecting millions of criminal defendants to procedures that determine their treatment based on actuarial instruments that explicitly treat socioeconomic and demographic factors as risk factors, and that means that poor people and people with the "wrong" demographics are being systematically and purposely treated more harshly by the criminal justice system. This is a serious injustice that has not received much attention, in large part because the instruments are not transparent and people who are not social scientists tend not to understand how they work. Those of you who are journalists can play an important role in bringing this problem to light.

These actuarial instruments have been around for decades in the context of parole board decision-making especially, and they are also now used in a variety of other criminal justice contexts. My own research has focused mostly on the use of risk assessment in sentencing, which is the fastest-growing trend in this area. In at least 20 states, many or all judges are being given risk scores for defendants before they sentence them, often as part of a presentence investigation report. Many other states are considering legislation to do the same, and prominent organizations like the National Center for State Courts and the American Law Institute, which drafts the Model Penal Code, have called for the expansion of this practice, which they often refer to as "evidence-based sentencing."

I find "evidence-based sentencing" to be something of a misnomer, bordering on doublespeak, because the risk scores don't actually have anything at all to do with the evidence in the defendant's own criminal case, which is normally the main thing that determines the defendant's sentence. Instead the "evidence" in question comes from studies of past offenders with similar preexisting characteristics—it's extrapolating the defendant's future crime risk based on a profile. So really, a better term for this is "profiling." And judges are told to use these profiling-based risk

predictions to determine the defendant's sentence, just like parole boards use them to decide whether to release a prisoner early.

There are a number of reasons to be concerned about this practice, but my primary concern is that many of the characteristics that are included in these profiles are inappropriate--and in some cases unconstitutional--bases for punishment. Put simply, people should not be punished extra, or for that matter punished less, based on who they are or how much money they have.

The instruments being used in sentencing and parole vary, but all contain several variables related to criminal history. Most also contain gender, age, employment status, education level, and marital status. The most popular instruments, like the LSI-R, include a whole battery of questions that relate to the defendant's financial status and history, family background, and neighborhood. For example, from the LSI-R:
--Financial problems, such as past or present trouble paying bills, rated from 0 to 3
--Reliance on social assistance, including welfare, unemployment, disability pensions
--Dissatisfaction with marital or equivalent situation: they rate the happiness of a person's relationship from 0 to 3
--Rewarding nature of a person's relationship with his parents—so an absent parent or one with whom the defendant has a bad relationship counts against him
--Similar ratings for relationships with other family members
--Whether parents or other family members have a criminal record
--Quality of accommodations
--Stability of accommodations—how often the person has moved
---High crime neighborhood
--Participation in organized leisure activities like membership in clubs (lack of this is a risk factor)
--Criminal records of acquaintances.

Another popular instrument, COMPAS, which for example just got adopted statewide in Michigan, includes similar factors, plus others, like chance of finding work above minimum wage, high school grades, whether the defendant's parents have been incarcerated, whether the defendant's parents used drugs, whether the defendant or any of his family members have ever been a crime *victim*.

Essentially, every indicator of socioeconomic disadvantage that you can think of has been included, and all of them add to the risk score. I want to make clear that this happens automatically, mechanically—*every* defendant who is on social assistance will have the same number of points added to his risk score because of it. It's built into the formula. We're used to thinking about disparities in sentencing as being something subtle and unconscious, insidious, something we have to detect through complicated empirical analyses—we look for evidence of whether judges are subtly taking inappropriate factors into account. But this is something different. This is the state *codifying* discrimination on the basis of these factors—it is explicitly built

into the instrument. Any time the judge gives any weight to the risk score, she is giving weight to socioeconomic and demographic factors. The point of this system is that the state *wants* poor people, people with all these risk factors, to be punished extra, and it's directing judges to do so.

Somewhat surprisingly, the nature and severity of the crime on which the defendant is being sentenced are not included in any of the instruments. Perhaps it's for this reason that the LSI-R training manual specifically says that it "was never designed to assist in establishing a just penalty," although that is precisely what it is now widely being used for.

Race is generally not included in the assessments, although certainly many of these variables are extremely strongly correlated with race. When you sentence people to extra time for being poor, you are bound to increase racial disparities as well.

The trend toward evidence-based sentencing has been greeted in large part with celebration. Scholars as well as judges, sentencing commissioners, and organizations focused on sentencing reform have embraced it as a new era of scientific, rational, "smarter" sentencing. Perhaps surprisingly, some of the strongest advocates have been progressive critics of mass incarceration, who hope that using risk scores will allow incarceration to be avoided in some cases by helping judges to identify low-risk offenders.

I disagree. It is bad policy and almost surely unconstitutional for the state to direct judges to deem classes of people categorically more dangerous, and sentence them for longer, on the basis of their poverty and their demographic characteristics. I agree that we have a mass incarceration crisis in this country, and we need to think creatively and in data-driven ways about policy solutions, but this particular use of data cannot be the right path. One of the reasons the social impacts of mass incarceration are so worrisome is that they are demographically, socioeconomically, and geographically concentrated. For instance, one in every nine black men under 35 is in prison right now, and one in three young black men will be at some point in his life. And if you narrow your focus to the poorest communities, or to particular crime-ridden neighborhoods, or to young men who are unemployed or lack high school diplomas, you get far higher numbers. There's a large literature documenting the hugely distortive effects on communities when you remove, say, half the young men in them. The risk prediction instruments could exacerbate all of these problems.

And that's one reason that people who maybe don't ordinarily worry so much about discrimination against men, or against the young or the unmarried, for instance, really should worry here. Those are all dimensions along which the impact of the criminal justice system is concentrated and concentration is something we should worry about.

I think that advocates of these instruments are in fact endorsing forms of explicit discrimination that they would never endorse were it not for the fact that they are somehow sanitized by the scientific framing that accompanies them—the fact that it's referred to as "evidence-based" and supported by regressions. But to me, behind this anodyne scientific language is an expressive message that is toxic. Stereotyping groups as criminally dangerous is a practice with a nasty cultural history in this country, and this practice involves the state officially labeling certain groups of people dangerous, on the basis of their identity and poverty, rather than their criminal conduct.

Basing sentences on gender as well as socioeconomic variables is also almost certainly unconstitutional, and my own research has been pitched at lawyers and judges to make this case.

First, gender. It's well established law that gender classifications require an "exceedingly persuasive justification"—this is a really tough test to pass. It's hardly ever legal for the state to treat people differently based on gender. Weirdly, though, even though everybody in the literature seems to take for granted that including race in the instruments would be unconstitutional, the use of gender doesn't seem to bother anyone. If scholars or advocates even mention it, they just say that because men really do on average pose higher recidivism risks, including gender in the instruments advances the state's important public safety interests and thus it passes the "exceedingly persuasive justification" test.

The problem with this response is twofold. First, this assumes the instruments actually advance those public safety interests effectively, which I think has not been persuasively established--I'll address that in a couple of minutes. Second, the argument runs afoul of one of the most central principles of the Supreme Court's gender discrimination jurisprudence: the prohibition on statistical discrimination. In general, the state cannot defend gender discrimination on the basis of generalizations about what men and women tend to do, even if those classifications are not just empty stereotypes but in fact are empirically well supported. In *Craig v. Boren*, for instance, the Court struck down a drinking-age law that discriminated against men even in the face of studies showing that young men posed more than ten times the drunk driving risk of young women.

There are lots of other examples in the case law, and this principle is something that really destroys any attempt to defend gender-based risk assessment, because the whole approach is grounded in reliance on statistical generalizations.

And this same principle is also the reason it's unconstitutional to discriminate in sentencing or parole based on financial factors such as unemployment, education, and income.

Until recently lawyers and legal scholars really had overlooked this problem. The reason for that is that generally, the courts are very tolerant of discrimination on the basis of socioeconomic status—they tend to defer to legislative judgments on that.

And so lawyers tend to think: Bringing a constitutional challenge based on socioeconomic discrimination is a loser.

But that's just not true when it comes to socioeconomic discrimination in the criminal justice system. For more than half a century, the Supreme Court has applied especially demanding scrutiny to policies adversely affecting poor defendants. The seminal case is *Griffin v. Illinois*, which described the provision of equal justice for poor and rich as the "central aim of our entire judicial system."

In *Bearden v. Georgia*, in 1983, the Supreme Court unanimously held that a defendant's probation could not be revoked because after losing his job he had become financially unable to pay a restitution order, since that would impermissibly make his sentence turn on his socioeconomic status.

Crucially, the Court in *Bearden* squarely rejected the state's attempt to argue, based on empirical studies of recidivism risk, that the defendant's unemployment and financial status rendered him an elevated public safety risk. The Court's response to this was much like its response to statistical discrimination in the gender context. It wrote:

"This is no more than a naked assertion that a probationer's poverty by itself indicates he may commit crimes in the future. …[T]he State cannot justify incarcerating [him] solely by lumping him together with other poor persons and thereby classifying him as dangerous. This would be little more than punishing a person for his poverty."

And that's exactly the problem with so-called "evidence-based sentencing." These actuarial instruments lump defendants together with other people who share their socioeconomic characteristics, and on the basis of those other people's past conduct, they classify defendants as dangerous. They punish a person for his poverty. And the Supreme Court has already unanimously held that unconstitutional—it just seems like everyone's forgotten.

OK, so what if we tried to predict risk statistically, but *didn't* use these demographic and socioeconomic characteristics? Suppose instead, we took into account the nature of the defendant's crime, which current risk instruments mainly ignore, as well as past history? That would be far less morally and legally problematic, because it would be based on the defendant's criminal conduct. And I think there's good reason to believe it would be about as accurate. Nothing predicts future behavior like past behavior, and I know Anne and the Arnold Foundation have found that a behavior-based risk assessment instrument at least in the bail context gets quite accurate results, which is a big step forward.

The thing is, factors like demographics and socioeconomics *are* correlated with crime, but once you already have behavioral factors, current and past criminal conduct, in your model, adding those problematic variables might not add much

marginal predictive value. Sure, adding more factors might get you another percentage point or two or three of accuracy, but I don't think we should pursue every last marginal improvement in predictive accuracy at all costs, at the cost of our most fundamental principles of equality and justice.

Now, beyond these equality concerns, I do have a few other concerns about these risk assessment instruments.

One problem is that if the purpose of risk assessment is to protect the public from the defendant's future crimes, these actuarial analyses are not actually asking the question they would need to ask in order to advance that purpose. They predict recidivism risk in the abstract—just "how risky is this person." They make no attempt to predict how the judge's sentencing choice would affect that risk—i.e., the responsiveness or "elasticity" of recidivism risk to differing lengths of incarceration.

And the people who have the highest recidivism risk are not necessarily the people whose recidivism risk is going to be the most reduced by incarceration—in fact, people who are more crime-prone to begin with may also be more likely to be hardened rather than helped by prison. We really don't have the science in place to know what subsets of people will have their behavior changed for the better by prison. Investigating this question requires studies that use rigorous causal inference methods—it's a very challenging empirical question. And so far, the best research on the way incarceration affects recidivism risk has been more general—does incarceration *generally* reduce crime risk--rather than focused on which characteristics are most associated with a greater responsiveness to incarceration.

Then there are some procedural concerns about risk assessment. One major concern is lack of transparency—people in many states are being sentenced on the basis of corporate, proprietary products that they don't have access to. Neither the defendant nor the judge knows the weight that has been given to each specific variable in producing the risk score. That's outrageous in my view.

In addition, defendants are essentially being forced to participate in an assessment interview, which includes detailed questions about their past and about their mental states. If they don't participate, they will be scored as uncooperative and may be punished for it. That seems like compelled self-incrimination, a violation of the Fifth Amendment.

I've got various other methodological objections that I've outlined in my paper, but I'll stop here. Thanks again, and I look forward to the rest of our discussion.