



The Counterintuitive Consequences of Sex Offender Risk Assessments at Sentencing

Megan T Stevenson, Jennifer L Doleac

University of Toronto Law Journal, Volume 73, Supplement 1, 2023, pp. 59-72 (Article)

Published by University of Toronto Press



➔ For additional information about this article

<https://muse.jhu.edu/article/905266>

🔗 For content related to this article

https://muse.jhu.edu/related_content?type=article&id=905266

Virginia adopted a risk assessment to help determine sentencing for sex offenders. It was incorporated as a one-way ratchet toward higher sentences: expanding the upper end of the sentence guidelines by up to 300 per cent. This led to a sharp increase in sentences for those convicted of sexual assault. More surprisingly, it also led to a decrease in sentences for those convicted of rape. This raises two questions: (a) why did sentencing patterns change differently across these groups, and (b) why would risk assessment lead to a reduction in sentence length? The first question is relatively easy to answer. While both groups saw an expansion in the upper end of the sentencing guidelines, only sexual assault had the floor lifted on the lower end, making leniency more costly. The second question is less straightforward. One potential explanation is that the risk assessment served as a political or moral shield that implicitly justified leniency for those in the lowest risk category. Even though the risk assessment did not change sentencing recommendations for low-risk individuals, it provided a ‘second opinion’ that could mitigate blame or guilt should the low-risk offender go on to reoffend. This decreased the risks of leniency and counterbalanced any increase in severity for high-risk individuals.

Keywords: algorithms in criminal justice, criminal justice reform, evidence-based criminal justice, risk assessments, sex offences

I Introduction

Preventive confinement is the practice of incarcerating someone not as a punishment for past crime but for the purposes of preventing crime that they might commit in the future. This practice is rightly controversial: how can one truly know what someone will do in the future? Should the government ever have such extensive power to take away someone’s liberty on the basis of mere speculation? Yet preventive confinement also has relatively broad support, at least in

* Associate Professor, University of Virginia School of Law, Charlottesville, United States

** Associate Professor, Department of Economics, Texas A&M University, College Station, United States

† The analyses presented here were originally included in an early draft of ‘Algorithmic Risk Assessment in the Hands of Humans,’ which now focuses solely on the impact of the non-violent risk assessment. Thanks to many individuals for their helpful comments along the way; particular thanks to those at the Virginia Criminal Sentencing Commission for providing the data and patiently answering questions. Thanks also to Ben Schoenfeld for providing the web-scraped court records used in this study, to the University of Virginia law librarians (particularly Ben Doherty and John Roper) for their helpful assistance, as well as to Billi Jo Morningstar for copyediting. Replication code and data available upon request. The authors have no conflicts of interest to report.

one key area: the prevention of sexual violence.¹ Those with a history of sexual offences are widely believed to be prone to high rates of recidivism, despite limited evidence to support this claim.² The perception of sexual crime as stemming from a deviant mental disorder has helped spawn widespread ‘sexual predator’ laws designed to incapacitate these individuals through long criminal sentences or civil commitment.³ Statistical risk assessments have been a central part of evaluating which individuals pose a high enough risk of sexual crime to warrant preventive confinement.⁴ These risk assessments are championed because of their relative accuracy,⁵ but there is little information about their impacts in practice.

In this symposium article, we evaluate how sex offender risk assessment has affected sentencing decisions in Virginia. In the early 2000s, Virginia adopted a risk assessment to help inform sentencing for sex offenders. The tool worked as a one-way ratchet, shifting the sentence guidelines range upward for those with higher risk scores. Accordingly, there was a sharp increase in sentences for those convicted of sexual assault. However, those individuals convicted of rape saw a decrease in sentences after risk assessment was adopted. The results pose two questions: (a) why did risk assessment affect sentencing differently for these two groups, and (b) why would risk assessment lead to a decrease in the sentence length? The first question is likely answered by differences in how the risk assessment was incorporated into the sentence guidelines. For those individuals convicted of rape (which includes forcible sodomy and object penetration), the tool worked only to increase the upper range of sentences that are considered to be compliant with the guidelines. Since judges needed to provide written explanation for sentencing outside of the guidelines range, this expanded the set of sentences that judges could give without having to provide a written explanation. For those convicted of sexual assault (which includes aggravated sexual battery and carnal knowledge of, or indecent liberties with, children), the tool also increased the upper range of sentences considered guidelines compliant. But it additionally impacted sentencing at the low range; in particular, probation or short jail sentences were no longer guidelines-compliant for those with a higher risk classification. A judge that wanted to be lenient on this group would need to provide a written explanation for their leniency.

The second question provides more of a puzzle. One potential explanation is that the risk assessment served as a political shield that allowed judges to be more lenient than they had been previously, at least for those in the lowest risk

1 See David DeMatteo et al, ‘A National Survey of United States Sexually Violent Person Legislation: Policy, Procedures, and Practice’ (2015) 14 *International Journal of Forensic Mental Health* 245.

2 R Karl Hanson & Dana Anderson, ‘Static-99R: An Empirical-Actuarial Risk Tool for Adult Males with a History of Sexual Offending’ in Kevin S Douglas & Randy K Otto, eds, *Handbook of Violence Risk Assessment*, 2d ed (New York: Routledge, 2021) 106.

3 Robert A Prentky et al, *Sexual Predators: Society, Risk, and the Law* (Oxford: Routledge, 2015).

4 See Christopher Slobogin, ‘Prevention of Sexual Violence by Those Who Have Been Sexually Violent’ (2011) 34 *Int J L & Psychiatry* 210.

5 William M Grove & Paul E Meehl, ‘Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical-Statistical Controversy’ (1996) 2 *Psychology, Public Policy and Law* 293.

category. The risk assessment served as a sort of second opinion that they could point to in order to justify relatively lenient sentences for those ranked as low risk. This reduced the likelihood that they would be blamed if a released individual went on to reoffend. After all, the risk assessment had rated that person as low risk, helping to justify the decision. This shift toward leniency was feasible for those convicted of rape since average sentences were well above the bottom end of the sentence guidelines. In contrast, risk assessment lifted the sentencing floor for sexual assault, raising the costs of leniency and counterbalancing the political shield effect.

II *The impact of sex offender risk assessment in Virginia*

A BACKGROUND

In 1999, the Virginia General Assembly adopted legislation requesting the Virginia Criminal Sentencing Commission (VCSC) to develop a risk assessment to be used in the sentence guidelines for sex offences.⁶ Claiming ‘that certain sex offenders are at high risk for re-offense’ and that ‘such sex offenders typically prey on vulnerable populations, such as children,’ the resolution stated that ‘it is important to identify and incapacitate, to the extent possible, these predatory sex offenders.’⁷ The VCSC proceeded to develop a sex offender risk assessment. They tracked 579 felony sex offenders released between 1990 and 1993 and followed them for a period of five to ten years.⁸ They then sought to identify statistical correlations between sexual or violent outcomes and predictive factors relating to the offence and the individual’s criminal record and personal characteristics.⁹ The sex offender risk assessment was adopted state-wide in the beginning of the fiscal year of 2002. This same risk assessment is still in use today.

One key challenge for sex offence risk assessment is that measured recidivism for a sex offence is quite low. We calculate that less than 2 per cent of convicted felony sex offenders in Virginia are convicted of a new sex crime within eight years after release.¹⁰ At that rate, only twelve of the 579 individuals in the training sample would have received a new conviction for a felony sex offence. Given the challenges associated with predicting a low-frequency event, especially with a small sample, the VCSC chose to train the prediction model (a) on arrest instead of conviction and (b) to predict any against-person offence, including crimes such as simple assault, aggravated assault, and robbery as well as sexual violence.¹¹ The prediction model ultimately included the following factors: age, education, employment, relationship with the victim, whether there was penetration or attempted penetration, location of the offence, history of arrest for against-person crimes, prior incarceration, and prior mental health/substance

6 Background information in this part largely comes from Virginia Criminal Sentencing Commission (VCSC), *Assessing Risk among Sex Offenders in Virginia* (Richmond, VA: VCSC, 2001), online: <www.vcsc.virginia.gov/sex_off_report.pdf> [VCSC, *Assessing Risk*].

7 United States, SJ Res 333, 1999 Sess, VA (1999).

8 VCSC, *Assessing Risk*, supra note 6 at 7.

9 Ibid.

10 We limit our sample to those convicted before risk assessment was adopted.

11 VCSC, *Assessing Risk*, supra note 6 at 7–8.

abuse treatment.¹² This model was then converted into an integer-weighted risk assessment, where the number of points for each predictor correspond with the strength of its correlation with the outcome.¹³ As is common in risk scores, a young age is one of the most heavily weighted factors.¹⁴ Being under thirty-five adds twelve points to the score. Thus, a twenty-two year old with no criminal record, no aggravating factors based on the offence, and no prior mental/substance abuse treatment and who is not regularly employed would receive twenty-one points, which is just shy of the twenty-eight points needed to be labelled high risk.

The risk assessment then serves to modify the sentence guidelines range. Three cut-offs were delineated in the sex offender risk assessment. Those who scored above each of the respective cut-offs were classified in one of three higher risk categories.¹⁵ Those in the highest risk category had the upper bound of the sentence guidelines range increased by 300 per cent; those in the lower categories had it increased by 100 per cent and 50 per cent respectively.¹⁶ For rape, there was no change to the bottom end of the guidelines range. For sexual assault, however, the risk assessment raised the floor of the sentence guidelines for high-risk individuals.¹⁷ While a low-risk individual convicted of sexual assault could be recommended for jail or probation, all high-risk individuals were automatically recommended for prison (which entails at least a twelve-month sentence). Anyone convicted of rape in Virginia is recommended for a prison sentence, both before and after risk assessment adoption.¹⁸

Both the sentence guidelines and the risk scores are calculated using a set of worksheets that are filled out by a probation officer or prosecutor after conviction. These worksheets are then provided to the judge at the sentencing stage. The cover page of these worksheets displays the sentence range recommended by the guidelines.¹⁹ A sentence is considered 'guidelines compliant' if it falls within this range. Note that the judges are not given any statistical information about the probability of recidivism. The risk assessment information comes solely in the form of an ordinal ranking between zero and sixty-seven as well as classification into one of four bins: the lowest-risk bin as well as the three higher-risk bins corresponding to the 50 per cent, 100 per cent, and 300 per cent increases in the upper end of the sentence guidelines range. Judges are allowed to sentence outside of the guidelines range as long as they stay within the statutory limits. But if they sentence outside the range, they must provide written justification. Average

12 The sex offender risk assessments can be found on the VCSC's website as Section A (Part I) for 'Other Sexual Assault' and Section A for 'Rape.' VCSC, *Other Sexual Assault Worksheet* (Richmond, VA: VCSC, 2021), online: <www.vcsc.virginia.gov/worksheets2021/Other_Sex_Asl.pdf>; VCSC, *Rape Worksheet 3* (Richmond, VA: VCSC, 2021), online: <www.vcsc.virginia.gov/worksheets2021/Rape.pdf> [VCSC, *Rape Worksheet 3*].

13 VCSC, *Assessing Risk*, supra note 6 at 8.

14 Megan T Stevenson & Christopher Slobogin, 'Algorithmic Risk Assessment and the Double Edge Sword of Youth' (2018) 96 Washington University Law Review 681 at 700–2.

15 VCSC, *Assessing Risk*, supra note 6 at 9–10.

16 VCSC, *Rape Worksheet 3*, supra note 12 at 1.

17 VCSC, *Assessing Risk*, supra note 6 at 10.

18 Ibid.

19 See e.g. VCSC, *Rape Worksheet 3*, supra note 12 at 1.

compliance rates are calculated by the VCSC in their yearly report. In practice, judges sentence within the guidelines range 80 per cent of the time.²⁰

After sentencing, the final sentence is written on a separate worksheet. All sheets are mailed to the VCSC, which maintains a database on all felony sentences in Virginia. In addition to housing records of felony sentencing, the VCSC maintains detailed records of every change to the sentence guidelines and every change to sentencing policy going back at least to 1995. One relevant sentencing change is that a separate risk assessment was adopted in the beginning of 2003 – the year after the sex offender risk assessment was adopted – for people convicted of drug, larceny, or fraud offences. We evaluate the impact of the non-violent risk assessment in a separate article entitled ‘Algorithmic Risk Assessment in the Hands of Humans.’²¹ In order to simplify the analysis, we drop anyone charged with drug, larceny, or fraud offences. There were no other major sentencing policy changes pertaining to the time period of our main analysis (from 2000 to 2003).

B DATA

The primary data source used in this article was acquired via a public records request from the VCSC. This data covers all individuals convicted of a felony in the state of Virginia from 1995 to 2017. It contains information on charges, sentences, sentence guidelines, and the risk score. Race and gender were obtained by matching the VCSC’s data with bulk court records taken from a public website.²² For most counties, court records from as far back as the year 2000 can be found online. Alexandria and Fairfax counties are an exception; we were unable to match the VCSC’s data back to the original court records for cases in these counties. The match was conducted using the fastLink package in R, which conducts probabilistic matching across multiple variables.²³ The variables used for the match included the first name, last name, offence date, birth month, and county of arrest. With a probabilistic match threshold of 0.92, 96 per cent of the post-2000 cases in the VCSC’s data (excluding Alexandria and Fairfax) found a match in the court records.

Descriptive statistics for defendants sentenced in the fiscal years of 2000 and 2001 (the last two years before the sex offender risk assessment was adopted) can be found in Table 1. The first column shows data for people convicted of rape, the second column shows descriptive statistics for those convicted of sexual assault, and the third column shows data for people convicted of non-sex offences. The top panel is the VCSC’s data, the bottom panel comes from court records. The most common non-sex offence categories include robbery, burglary, assault, and traffic. (As discussed, we dropped anyone convicted of drug, larceny, or fraud offences.) Note that, before risk assessment was adopted, those convicted of rape had much more serious sentences than those convicted of sexual assault: 97 per cent

20 VCSC, *Annual Report* (Richmond, VA: VCSC, 2022) at 12.

21 Megan T Stevenson & Jennifer L Doleac, ‘Algorithmic Risk Assessment in the Hands of Humans’ (29 September 2022), online: SSRN <ssrn.com/abstract=3489440>.

22 This data was web scraped and made publicly available by Ben Schoenfeld. It can be downloaded at *Virginia Court Data*, online: <virginiacourtdata.org> (last updated 7 April 2021).

23 Ted Enamorado et al, ‘Using a Probabilistic Model to Assist Merging of Large-Scale Administrative Records’ (2019) 113 *American Political Science Review* 353.

TABLE 1. Summary statistics

	Rape	Sexual assault	Non-sex offense
Sentencing Commission data			
Age	32.95	34.41	31.34
Age<23	0.24	0.23	0.30
Recent felony conviction	0.02	0.02	0.07
Rape	1.00	0.00	0.00
Sex assault	0.00	1.00	0.00
Burglary	0.00	0.00	0.22
Robbery	0.00	0.00	0.13
Assault	0.00	0.00	0.20
Traffic	0.00	0.00	0.19
Guidelines-recommended sentence	213.18	28.09	44.19
Incarcerated	0.97	0.75	0.82
Months incarcerated	173.69	23.65	40.66
Observations	434	836	10621
Circuit Court data			
Black	0.49	0.32	0.52
Female	0.02	0.03	0.09
Observations	272	496	6841

Note: This table provides summary statistics for those convicted of rape, sexual assault, or a non-sex offence. The data contains observations from the two years prior to risk assessment adoption: fiscal year 2000–1. Each entry can be interpreted as the fraction of individuals in that column with that particular characteristic, with the exception of age, which is measured in years, the guidelines-recommended sentence, which is measured in months, and months incarcerated, which is measured in months. The top panel comes from VCSC’s data and the bottom panel comes from Circuit Court data.

of those convicted of rape were incarcerated versus only 75 per cent for sexual assault. The average sentence was 173 months for those convicted of rape, but only twenty-three months for sexual assault.

C DOES THE RISK ASSESSMENT AFFECT SENTENCING?

We begin with a simple inquiry: do judges even use the risk assessment? In particular, we evaluated whether the sharp discontinuities in the sentencing recommendations associated with the risk score translate into sharp discontinuities in sentencing. The risk assessment is incorporated into the sentence guidelines in a way that creates very different sentencing recommendations for people with adjacent risk scores. Those individuals with a risk score of twenty-seven should be very similar to those with a risk score of twenty-eight. However, since twenty-eight is the threshold at which one enters a high-risk category, the sentencing recommendations for them are very different. Those with a risk score of twenty-eight are automatically recommended for prison and have the upper bound of the guidelines range expanded by 50 per cent. Those with a risk score of twenty-seven are not rated high risk. There is no expanded upper bound to their sentence, and, in sexual assault, they may receive a lower recommended sentence




FIGURE 1. Sentencing discontinuities in the risk score

Note: This figure plots sentences against the sex offender risk score. The graph on the left shows the probability of incarceration and the graph on the right shows the sentence length with an arcsinh transform. The data is divided into bins of three risk scores; each dot shows the average incarceration rate or sentence for that bin, and the vertical whisker bar shows the 95 per cent confidence interval for the mean. Those with a score of twenty-eight or above are rated high risk.

than prison. If we see sharp differences in sentencing practices for those with the risk score of twenty-seven versus twenty-eight, one can infer that this is because the sentencing recommendations were influential.

Figure 1 provides visual evidence. The graph on the left plots the probability of incarceration for people with different risk scores, and the graph on the right plots the sentence length with an inverse hyperbolic sine transform. (This transform reduces the influence of extreme outliers; its interpretation will be discussed in the subsequent subpart.) Each dot represents the average incarceration rate across a bin of three risk scores. The vertical whisker bars show the 95 per cent confidence interval for the mean. Scores to the right of the vertical line are in one of the high-risk categories. As can be seen, there is a sharp increase in sentences for those with a score of twenty-eight versus twenty-seven.

One concern with this analysis is that the discontinuity that is documented does not come from the influence of the sentencing recommendations associated with the risk score but, rather, because those on either side of the threshold have starkly different characteristics. Perhaps there is some attempt to game the risk score, pushing certain unfavoured individuals up to a score of twenty-eight or other favoured individuals down to a score of twenty-seven. We can test for this in two ways. First, we can look for ‘bunching’ in the distribution of the risk score. If there is a bunching of people right below or right above the threshold, this could be evidence that the score is being manipulated one way or the other. We can also test to see if aspects of the criminal record, defendant demographics, or charge of conviction differ discontinuously across the threshold. We provide these graphical analyses in Figure A1 in the Appendix. There is no evidence of gaming the risk score in this manner: no bunching around the threshold and no discontinuities in defendant characteristics. Finally, we note that there is no discernible sentencing discontinuity across the upper risk score thresholds at thirty-four and forty-five. It may be that the first threshold between the low-risk

category and one of the high-risk categories is the most influential one. It may also be that our sample is smaller in these higher threshold ranges so the discontinuity is harder to detect.

D WHAT IS THE NET IMPACT ON SENTENCING?

In this subpart, we evaluate how the adoption of risk assessment affects average incarceration rates and sentence lengths compared to the *status quo* (that is, sentencing without risk assessment). Our method is a difference-in-differences analysis that compares outcomes before and after the adoption of risk assessment for defendants who are or are not eligible for risk assessment. This involves a two-way comparison. Effectively, this entails evaluating any change in sentences for those convicted of sex offences before and after risk assessment is adopted. This is then compared to any before/after change in sentences for those convicted of some other offence, which forms a sort of benchmark that can capture time trends in sentencing that might be expected to affect both sex offenders and non-sex offenders similarly. Specifically, we use the regression equation below:

$$Y_i = \alpha + SexOffense_i * Post_i \beta_0 + SexOffense_i * \beta_1 + Post_i * \beta_2 + X_i * \beta_3 + \psi_i$$

Each observation i consists of a single case. *Post* refers to the period after the sex offender risk assessment was adopted, and *SexOffence* indicates that the person was convicted of a sex offence. Covariates (X) include the offence, age, race, gender, whether there was a prior felony conviction within three years, judge fixed effects, and fixed effects for the month, year, and weekday of sentencing.²⁴ For those convicted of rape, our covariates also include the guidelines-recommended sentence without the risk assessment modifications. (This data was not available for sex assault.) Our sample includes the four years encompassing the fiscal years 2000–3. We cluster standard errors at the level of the judge. Given the different way in which risk assessment was incorporated into the guidelines, we show results separately for rape and sexual assault. In other words, we conduct a difference-in-differences analysis where we compared sentencing outcomes before and after risk assessment for those convicted of rape against those convicted of a non-sex offence and similarly for sexual assault.

We first show the impact for those convicted of sexual assault. Our difference-in-differences estimates are shown in Table 2. The odd columns do not include controls for offence, criminal record, and demographics; the even columns do. The first two columns show the impact on the probability of incarceration. The last two columns show the impact on sentence length. Sentences have a long right tail, meaning that there are a number of outliers with extremely long sentences. In order to reduce the influence of these extreme outliers, we use an arcsinh transform on the sentence length variable, which also allows us to interpret the coefficients as an approximate percentage change.²⁵ For those convicted of sexual assault, risk assessment adoption increased the probability of incarceration by 5.5 percentage points and increased the sentence length by approximately 45 per cent. The inclusion of controls modifies the point estimates only slightly, suggesting that trends in offence, criminal record, and demographics do

TABLE 2. The impact of risk assessment on sentences for sexual assault

	(1) Inc.	(2) Inc.	(3) Sentence length	(4) Sentence length
Impact of risk assessment	0.0322 (0.0203)	0.0556*** (0.0204)	0.396**** (0.0865)	0.454**** (0.0840)
Observations	26832	26832	26832	26832
R ²	0.00335	0.111	0.00140	0.258
Mean DV	0.739	0.739	2.563	2.563

Note: This table presents difference-in-differences estimates of the impact of risk assessment on sentencing for those convicted of sexual assault. The odd columns omit covariates for offence, criminal record, and demographics. The first two columns show the impact on the probability of incarceration. The last two columns show the impact on the sentence length, with an arcsinh transform. Standard errors are clustered at the judge level. The mean dependent variable for sexual assault offenders prior to risk assessment adoption is shown in the bottom row.

not vary substantially across individuals convicted of sexual assault versus those convicted of non-sex offences.

We then show the impact for those convicted of rape. The difference-in-differences estimates are shown in Table 3, which is formatted similarly to Table 2. We see a striking difference from the previous table. Whereas risk assessment led to a net increase in sentences for sexual assault, it led to a net decrease in sentences for those convicted of rape. After risk assessment was adopted, the probability of incarceration declined by 4.5 percentage points and the sentence length declined by approximately 18 per cent (significant only at the 10 per cent level). In the Appendix, we show that the results are generally robust to variations in the specification: Table A1 shows results with only violent offences included in the control group, and Tables A2 and A3 show results with varying time windows. Instead of the four year sample in our main specification, we show results with a two-year, six-year, and eight-year sample, centred around the date of risk assessment adoption.

We then explore the timing of this effect using a dynamic difference-in-differences approach which entails replacing *Post* with dummies for each year from 1998 to 2005. Figure 2 shows the results for sexual assault. The graph on the left shows the impact on the probability of incarceration; the graph on the right shows the impact on the sentence length with an arcsinh transform. The year prior to risk assessment adoption serves as a benchmark and has been normalized to be equal to zero. Therefore, each dot on the graph can be thought of as the difference in sentences for sexual assault and non-sex offences in the given year as compared to the difference in sentences in the year before risk assessment adoption. The vertical whisker bar shows the 95 per cent confidence interval for that difference. We see that sentencing trends were approximately parallel for sexual assault and non-sex offences between 1998 and 2001. Visually, the 95

24 Since we do not have complete race and gender data, the controls are indicators for whether the defendant is Black, non-Black, or race missing. For gender, the controls are indicators for being female, non-female, or gender missing.

25 An arcsinh transform is similar to a log transform but can be applied to variables that take a zero value.

TABLE 3. The impact of risk assessment on sentences for those convicted of rape

	(1) Inc.	(2) Inc.	(3) Sentence length	(4) Sentence length
Impact of risk assessment	-0.0502*** (0.0157)	-0.0455*** (0.0151)	-0.159 (0.115)	-0.179* (0.0989)
Observations	25918	25918	25918	25918
R ²	0.00565	0.325	0.0489	0.608
Mean DV	0.962	0.962	5.311	5.311

Note: This table presents difference-in-differences estimates of the impact of risk assessment on sentencing for those convicted of rape. The odd columns omit covariates for offence, criminal record, and demographics. The first two columns show the impact on the probability of incarceration. The last two columns show the impact on the sentence length, with an arcsinh transform. Standard errors are clustered at the judge level. The mean dependent variable for rape offenders prior to risk assessment adoption is shown in the bottom row.




FIGURE 2. Sentencing time trends for those convicted of sexual assault

Note: This figure presents dynamic difference-in-differences estimates of the impact of risk assessment on sentencing for those convicted of sexual assault. The year before risk assessment is adopted is delineated in the vertical dashed line. The left graph shows the impact on the probability of incarceration and the right graph shows the impact on the sentence length, with an arcsinh transform.

per cent confidence interval for each of these pre-risk assessment adoption years contains zero, meaning that the sentencing differential between sexual assault and non-sex offences was not statistically different from the sentencing differential in the year prior to risk assessment adoption. However, in the years after risk assessment, sentences were noticeably higher. The increase began in the first year after risk assessment adoption, accelerated the second year, and remained escalated throughout the time period of our analysis.

For rape, the pattern is reversed. Immediately after the adoption of risk assessment, there is a decline in both the probability of sentencing and the sentence length, as shown in Figure 3. Sentences remain low throughout subsequent years.

As we have shown, the net sentencing effects are quite different from what one might have expected based on how the risk assessment was incorporated into




FIGURE 3. Sentencing time trends for those convicted of rape

Note. This figure presents dynamic difference-in-differences estimates of the impact of risk assessment on sentencing for those convicted of rape. The year before risk assessment is adopted is delineated in the vertical dashed line. The left graph shows the impact on the probability of incarceration and the right graph shows the impact on the sentence length, with an arcsinh transform.

sentence recommendations. The sex offender risk assessment was incorporated in a unidirectional manner: to increase sentences for the highest risk offenders. Yet sentences for those convicted of rape went in the opposite direction. What is going on here? The previous subpart showed us that, at least around the margins, the sentencing recommendations associated with the risk classification affected sentencing decisions. But apparently, judges responded as much to the absence of the recommendation as they did to the recommendation itself. To whatever extent risk assessment increased sentences for rape offenders who were recommended for a longer sentence – those who scored above the cut-off – it must have decreased sentences even more for those who were not. One might speculate that the risk assessment taught judges that, on average, offenders convicted of rape had lower levels of risk than they had previously thought. However, the risk assessment conveyed no information about the probabilistic likelihood of reoffending.²⁶ Instead, the risk assessment only provided a numerical ranking and coarse classification – for example, ‘[t]his person is a 35 and is rated Level 2.’ Therefore, the only information it provides is relative information across different people convicted of sex offences. And those convicted of rape had slightly higher risk scores on average than those convicted of other sexual crimes. There is no straightforward reason why such information would lower expectations about recidivism.

The evidence is more consistent with a story in which the risk assessment altered incentives as well as information. One such explanation is captured in what we refer to as the ‘political shield’ hypothesis. Under this hypothesis, the risk assessment serves to mitigate some of the responsibility judges face when making difficult

26 Many argue that risk assessment information should provide statistical predictions of risk rather than hard-to-interpret ordinal classifications. One counterpoint is that people may have trouble cognitively processing low probabilities. Weather forecasting, for instance, uses statistical probabilities for high frequency events like rain and categorical predictions for low frequency events like hurricanes. See John Monahan & Henry J Steadman, ‘Violent Storms and Violent People: How Meteorology Can Inform Risk Communication in Mental Health Law’ (1996) 51 *American Psychologist* 931 at 934–5.

decisions about sentencing. By mitigating responsibility, it also mitigates public blame and personal remorse when a released person goes on to reoffend. A mistaken decision can be blamed on a faulty risk score rather than a faulty personal assessment.

This theory was expressed to us by multiple judges during the course of our research. And concern about public backlash when granting lenient sentences to those convicted of sex offences is not just hypothetical. In 2016, Judge Aaron Persky gave only a six-month jail sentence to a Stanford student who had been convicted of sexual assault. This sentence resulted in widespread public outrage. Judge Persky became a household name, and activists mounted a successful recall movement that cost him his job.²⁷ If the Stanford student had received a low-risk score, this would almost certainly be cited to help defend the low sentence he received. It may not have changed the fate of Judge Persky, but it may change outcomes in more marginal situations or in cases where an apparently low-risk offender goes on to commit another crime.

Of course, this is only speculative. We cannot say for certain why a risk assessment schema designed to increase sentences ultimately lowered sentences for rape. Nonetheless, it provides a stark example of how risk assessment can have counter-intuitive consequences.

III *Conclusion*

We evaluate the impacts of adopting sex offender risk assessment at sentencing. This risk assessment was adopted as a one-way ratchet toward higher sentences. It expanded the upper range of the sentence guidelines for both sexual assault and rape and lifted the floor of the guidelines for rape. For sexual assault, results were as expected: sentences increased. For rape, the impact was surprising: sentences decreased on average. The fact that risk assessment impacted these two groups differently can be explained by differences in how risk assessment was incorporated into sentence guidelines across these groups – there was an increase in the costs of leniency for sexual assault but not for rape. But the fact that sentences declined for rape is more of a mystery. We think that it is unlikely to be solely because of the improved information provided by risk assessment. First, the risk assessment showed that those convicted of rape were higher risk, on average, than those convicted of sexual assault. Second, the risk assessment provided no information about actual statistical risk – that is, the probability of new against-person arrest. Rather, it provided only a relative ranking of risk. We propose one potential explanation: the risk assessment served as a sort of second opinion that made judges feel more comfortable granting leniency for those in the lowest risk category. Without this second opinion, a judge bears the weight of the risk entirely on their own shoulders. Should the person go on to reoffend, there is no one but the judge to blame for being so lenient. With risk assessment,

27 For a description of the case and Judge Persky's recall, see 'Recent Election – Sexual Assault Law – Judicial Recall – California Judge Recalled for Sentence in Sexual Assault Case' (2019) 132 *Harvard L Rev* 1369.

the judge can point to a faulty score when a low-risk person goes on to reoffend, thus mitigating internal guilt and external blame.

Appendix




FIGURE A1. Those above or below the threshold have similar characteristics

Note: The top left figure shows the distribution of the sex offender risk score. Each vertical line indicates a threshold between risk levels. The remaining figures plot various defendant characteristics against the sex offender risk score. The data is divided into bins of three risk scores; each dot shows the average incarceration rate for that bin and the vertical whisker bar shows the 95 per cent confidence interval for the mean. Those with a score of twenty-eight or above are in a high-risk category. The characteristics considered include whether or not the person has a prior felony conviction within three years, whether they are black or female, what their age is, and whether they were convicted of rape.

TABLE A1. Robustness tests for difference-in-differences results: only violent offences as a control group

	(1) Inc.	(2) Sentence length	(3) Inc.	(4) Sentence length
Impact of risk assessment	0.0552** (0.0214)	0.442**** (0.0886)	-0.0336** (0.0155)	-0.116 (0.101)
Observations	10903	10903	9989	9989
R ²	0.0800	0.278	0.295	0.606
Mean DV	0.739	2.563	0.962	5.311

Note. This table presents difference-in-differences estimates of the impact of risk assessment on sentencing. The control group includes only those convicted of violent offences. The first two columns show the impact on those convicted of sexual assault, the last two columns show the impact on those convicted of rape. The outcomes are the probability of incarceration and the sentence length, with an arcsinh transform. All specifications include the full set of controls. Standard errors are clustered at the judge level. The mean dependent variable for rape or sexual assault offenders (respectively) prior to risk assessment adoption is shown in the bottom row.

TABLE A2. Robustness tests for sexual assault difference-in-differences results: varying time windows

	(1) Inc.	(2) Sent. length	(3) Inc.	(4) Sent. length	(5) Inc.	(6) Sent. length
Impact of R.A.	0.0368 (0.0267)	0.300** (0.118)	0.0604**** (0.0175)	0.489**** (0.0752)	0.0719**** (0.0149)	0.516**** (0.0658)
Observations	14129	14129	39191	39191	52149	52149
R ²	0.122	0.269	0.112	0.258	0.114	0.258
Mean DV	0.739	2.563	0.739	2.563	0.739	2.563

Note. This table presents difference-in-differences estimates of the impact of risk assessment on sentencing for those convicted of sexual assault. The sample spans fiscal years 2001–2 in the first two columns, 1999–2004 in the middle two columns, and 1998–2005 in the final two columns. The outcomes are the probability of incarceration and the sentence length, with an arcsinh transform. All specifications include the full set of controls. Standard errors are clustered at the judge level. The mean dependent variable for sexual assault offenders prior to risk assessment adoption is shown in the bottom row.

TABLE A3. Robustness tests for rape difference-in-differences results: varying time windows

	(1) Inc.	(2) Sent. length	(3) Inc.	(4) Sent. length	(5) Inc.	(6) Sent. length
Impact of R.A.	-0.0463* (0.0253)	-0.274* (0.159)	-0.0444**** (0.0128)	-0.111 (0.0885)	-0.0383**** (0.0115)	-0.00921 (0.0781)
Observations	13684	13684	37906	37906	50430	50430
R ²	0.126	0.310	0.117	0.299	0.119	0.298
Mean DV	0.962	5.311	0.962	5.311	0.962	5.311

Note. This table presents difference-in-differences estimates of the impact of risk assessment on sentencing for those convicted of rape. The sample spans fiscal years 2001–2 in the first two columns, 1999–2004 in the middle two columns, and 1998–2005 in the final two columns. The outcomes are the probability of incarceration and the sentence length, with an arcsinh transform. All specifications include the full set of controls. Standard errors are clustered at the judge level. The mean dependent variable for rape offenders prior to risk assessment adoption is shown in the bottom row.