

VIRGINIA LAW REVIEW ONLINE

VOLUME 103

DECEMBER 2017

94-102

ESSAY

ACT-SAMPLING BIAS AND THE SHROUDING OF REPEAT OFFENDING

Ian Ayres, Michael Chwe and Jessica Ladd.

A college president needs to know how many sexual assaults on her campus are caused by repeat offenders. If repeat offenders are responsible for most sexual assaults, the institutional response should focus on identifying and removing perpetrators. But how many offenders are repeat offenders? Ideally, we would find all offenders and then see how many are repeat offenders. Short of this, we could observe a random sample of offenders. But in real life, we observe a sample of sexual assaults, not a sample of offenders.

In this paper, we explain how drawing naive conclusions from “act sampling”—sampling people’s actions instead of sampling the population—can make us grossly underestimate the proportion of repeat actors. We call this “act-sampling bias.” This bias is especially severe when the sample of known acts is small, as in sexual assault, which is among the least likely of crimes to be reported.¹ In general, act sampling can bias our conclusions in unexpected ways. For example, if we use act sampling to collect a set of people, and then these people truthfully report whether they are repeat actors, we can overestimate the proportion of repeat actors.

* Ayres is the William K. Townsend Professor at Yale Law School; Chwe is Professor of Political Science at the University of California, Los Angeles; and Ladd is the Founder and CEO of Callisto.

¹ David Cantor et al., Report on the AAU Campus Climate Survey on Sexual Assault and Sexual Misconduct, Westat, at iv (2015) [hereinafter AAU Survey], (noting that a relatively small percentage of the most serious offenses are reported) <https://perma.cc/5BX7-GQPU>; Nat’ Inst. Just., U.S. Dep’t of Just., Measuring Frequency (Oct. 1, 2008), <https://perma.cc/ZTD8-U9GT>.

Again, say that you are a university president and a recent undergraduate survey suggests that only 10 percent of students who are sexually assaulted make a formal complaint. Last year, your committee on sexual misconduct received 100 formal complaints for non-consensual “sexual contact involving physical force or incapacitation,”² but, given the under-reporting of sexual assault, you figure that during this time period there were likely a total of 1,000 assaults. Among the perpetrators in these 100 complaints, only five were found to have assaulted more than one person. The vast majority of offenders, 95 out of 100, were only accused by a single person.

It seems that only 5 percent of the school’s offenders are repeat offenders. But this conclusion is strikingly inaccurate. The proportion of offenders who are repeat offenders is probably much closer to 100 percent than 5 percent, and the vast majority of sexual offenders at your school may in fact be repeat offenders.

I. SHROUDING THE PREVALENCE OF REPEAT OFFENDING

A low reporting rate shrouds the prevalence of repeat offending. Recall that the probability that a given assault is reported is 10 percent, or 0.10. If the probabilities that survivors report are independent of each other,³ the probability that an offender who assaults two people will be reported by both is 0.10 times 0.10, or just 1 percent. The probability that an offender who assaults two survivors is not reported by either is 0.90 times 0.90, or 81 percent. The probability that an offender who assaults two survivors is reported by just one is the remaining probability, 18 percent.⁴ So the vast majority of reported two-time offenders will only be accused once. Or put another way, if *all* offenders committed exactly two offenses, then with a 10 percent reporting rate, we should expect that only 5.3 percent ($1\% / (1\% + 18\%)$) of reported offenders would be accused twice. In this example (with 10 percent

² AAU Survey, *supra* note 1, at ix.

³ If an initial report is made public, the probability of other reports against the same accused may be much higher than initial probability (and hence not independent).

⁴ More generally, when the reporting rate is p , we have:

Probability of being accused N times when assaulting M times = $\binom{M}{N} p^N (1 - p)^{M-N}$

reporting), even when *every* offender is a repeat offender, only 5.3 percent of reported offenders will appear to be repeat offenders.

We have built a widget, available below,⁵ that allows you to see just how large the shrouding effect will be in particular settings. Just plug in the number of repeat offenders in the population and the number of offenses they each commit, the number of non-repeat offenders, and the probability that any offense will be reported, and then click the “randomize” button. The widget will then tell you the prevalence of repeat offenders in the general population as well as in the subpopulation of reported offenders. For example, if there are 10 repeat offenders who each offend twice, 10 non-repeat offenders, and there is a 10 percent chance that offenses will be reported, then even though half the offenders are repeat offenders, only 3.4 percent of reported offenders are accused twice.

These examples dramatize how an under-reporting problem can disproportionately distort the appearance of repeat offending. When less than 50 percent of assaults are reported, we should expect that the prevalence of repeat offending in the general population is radically larger than the share of repeat offending seen in the reported cases. The following table lets us work backward from the sample of reported offenses to estimate the proportion of repeat offenders in the whole population. The values in the middle of the table are the expected proportion of offenders who are accused twice among reported offenders. If at your university, you believe that about 30 percent of sexual assaults are reported and find that about 11 percent of reported offenders are reported twice, then you should infer that about 50 percent of offenders in the entire population are repeat offenders.

⁵ Act-Sampling Bias Widget, <http://chwe.net/repeat/>

	Report Probability											
	3.4%	1%	5%	10%	15%	20%	25%	30%	35%	40%	45%	99%
Repeat offender	1%	0.0%	0.0%	0.1%	0.1%	0.2%	0.2%	0.3%	0.3%	0.4%	0.4%	1.0%
share of offender	10%	0.1%	0.5%	0.9%	1.4%	1.9%	2.3%	2.8%	3.3%	3.8%	4.3%	9.9%
population	20%	0.2%	0.8%	1.7%	2.6%	3.4%	4.3%	5.3%	6.2%	7.1%	8.1%	19.8%
	30%	0.2%	1.2%	2.4%	3.6%	4.8%	6.1%	7.4%	8.8%	10.2%	11.6%	29.6%
	40%	0.3%	1.4%	2.9%	4.5%	6.1%	7.7%	9.4%	11.1%	12.9%	14.8%	39.4%
	50%	0.3%	1.7%	3.4%	5.3%	7.1%	9.1%	11.1%	13.2%	15.4%	17.6%	49.3%
	60%	0.4%	1.9%	3.9%	6.0%	8.1%	10.3%	12.7%	15.1%	17.6%	20.3%	59.0%
	70%	0.4%	2.1%	4.3%	6.6%	9.0%	11.5%	14.1%	16.8%	19.7%	22.7%	68.8%
	80%	0.4%	2.3%	4.7%	7.1%	9.8%	12.5%	15.4%	18.4%	21.6%	25.0%	78.6%
	90%	0.5%	2.4%	5.0%	7.6%	10.5%	13.4%	16.6%	19.9%	23.4%	27.1%	88.3%
	99%	0.5%	2.6%	5.2%	8.1%	11.0%	14.2%	17.5%	21.1%	24.8%	28.8%	97.0%

Table 1: The proportion of repeat offenders among reported cases as a function of the proportion of repeat offenders in the entire population and the probability that a given assault is reported (assuming repeat offenders assault 2 times).

These examples are based on the assumption that people who are assaulted by repeat offenders have the same probability of reporting as people who are assaulted by non-repeat offenders.⁶ This might not be the case; it is possible that repeat offenders assault people who they think are less likely to file a formal complaint, for example. In any case, the tendency of many schools to keep investigations of (and even sanctions for) sexual misconduct private will keep potential complainants in the dark about whether another complaint has already been filed against a particular offender.

Indeed, the shrouding of repeat offending may reinforce the reluctance of some survivors to report. If survivors think that reports of sexual assault are more likely to be taken seriously if corroborated by another independent report accusing the same person, then they, like the college president, might also infer from the sample of reported cases that there is a relatively small chance that their accused was a repeat offender. Thus, the shrouding effect decreases the probability that an assault is reported, which in turn further worsens the shrouding effect, in a negative feedback loop.

The foregoing estimates are also based on an assumption that all repeat offenders offend exactly twice. The shrouding effect becomes less pronounced as the average number of offenses committed by repeat offenders increases. For example, with a 10 percent reporting rate, if all offenders are repeat offenders who each offend

⁶ And as noted above, it also assumes that each survivor of a repeat offender has an independent and equal probability of reporting.

three times, then there is a 10.3 percent chance that reported offenders will be reported at least twice (instead of the 5.3 percent chance when the repeat offenders committed just two offenses). We have created an Excel spreadsheet, available below,⁷ where you can vary the assumptions of the proportion of the offending population that commits one, two, three, or four offenses and see the resulting expected proportion of repeat offenders among the subgroup that are reported.

Schools should take into account the shrouding of repeat offenders when devising their reporting protocols. Survivors of sexual assault are at times reluctant to be the first person to bring a complaint that will launch an investigation against the person who attacked them.⁸ If many survivors are attacked by repeat offenders, a reporting mechanism like Callisto, which allows survivors to deposit secure, time-stamped descriptions of their attack into encrypted escrow and also allows survivors to choose to automatically alert authorities if and when a second accusation is made against the same person, can powerfully respond to first-mover reluctance.⁹ With a program like Callisto, survivors are assured that their claim will indirectly be corroborated by another person's claim.¹⁰

II. ACT-SAMPLING BIAS

Decades of research have shown that it is common for people to make predictions based on how representative something is rather than how likely it is.¹¹ An oft-cited example of this bias is a study in

⁷ Act-Sampling Bias Spreadsheet, <http://chwe.net/repeat/censoringbias.xlsx>. The excel spreadsheet downloads automatically.

⁸ Ian Ayres, Meet Callisto, the Tinder-like Platform that Aims to Fight Sexual Assault, *Wash. Post* (Oct. 9, 2015),

⁹ Project Callisto, <https://www.projectcallisto.org>.

¹⁰ Ayres, *supra* note 8.

¹¹ Amos Tversky & Daniel Kahneman, Evidential Impact of Base Rates, *in* *Judgment Under Uncertainty: Heuristics and Biases* 153, 153 (Daniel Kahneman et al. eds., 1982); *see* Amos Tversky & Daniel Kahneman, Judgments of and by Representativeness, *in* *Judgment Under Uncertainty: Heuristics and*

which participants were presented with the hypothetical of a 31-year old woman, Linda, who is “single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.”¹² When presented with the question whether it is more likely that Linda “is a bank teller” or “a bank teller and is active in the feminist movement,” a majority of participants said it was more probable that “Linda is a bank teller and is active in the feminist movement” even though any person who is a bank teller and active in the feminist movement automatically also falls into the category of being a bank teller.¹³

Our example regarding sexual assault statistics points to a similar kind of error. People mistakenly tend to think that the proportion of repeat offenders in a sample of offenses indicates the proportion of repeat offenders in the entire population. As noted above, we call this “act-sampling bias” because it results from people incorrectly making a conclusion based on the proportion in a sample of *acts*, instead of a sample of people. This leads them to incorrectly conclude that if in a sample of offenses, only 5.3 percent are from repeat offenders, then only 5.3 percent of the offenders in the entire population are repeat offenders.

We are misled because of the disproportionate absence of repeat offenders in the sample relative to the population. For a one-time offender to appear in the sample, he has to be reported only once. For a two-time offender to appear in the sample as a two-time offender, he has to be reported twice, which is much less likely. Put another way, a one-time offender “escapes” the sample if one survivor does not report. A two-time offender “escapes” appearing in the sample as a two-time offender if either one of two survivors does not report.

Biases 84, 84–89 (Daniel Kahneman et al. eds., 1982) [hereinafter “Tversky & Kahneman, Judgments”].

¹² Tversky & Kahneman, Judgments, *supra* note 11, at 84, 91–93.

¹³ *Id.*; see also Amos Tversky & Daniel Kahneman, Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment, 90 *Psychol. Rev.* 293, 297 (1983) (providing another description of the representative bias).

In other contexts, extrapolating from the sample to the population is warranted. For example, if 52 percent of a representative sample of U.S. voters support Hillary Clinton, it is appropriate to infer that within a small margin of error, about 52 percent of the true voting population is in support of Clinton.¹⁴ However, when sampling is based on people's actions, then the very process of sampling disproportionately affects the likelihood that people who take that action will appear in the sample.

What we are calling the "act-sampling bias" joins a host of previously recognized representativeness "fallacies," including the base rate, conjunction, disjunction, regression, and gamblers' fallacies, that can lead to false inference.¹⁵ The act-sampling bias occurs whenever sampling is tied to the actions of the subjects instead of to the subjects themselves. Sampling on the basis of the act instead of the actor distorts the representativeness of the sample based on how often the subject engages in the activity.

For example, imagine that McDonalds is interested in learning what proportion of high school students visit its restaurants more than once a week. If it went to high schools and surveyed a random 10 percent of students, the sample should capture one-time visitors and multiple-time visitors with no bias. If instead, McDonald's went to its restaurants at different times and randomly sampled, by noting who was present, 10 percent of the high school patrons eating at the restaurant, then the proportion of subjects who appeared repeatedly in the sample would, as in our sexual assault

¹⁴ When drawing inferences regarding the population from a representative sample, sampling weights are usually used by social scientists to adjust for the probability of an individual or a household being sampled from the population. These weights, therefore, account for the degree of the representativeness of the sample when extrapolating from the sample. For a review of the construction and use of sampling weights to draw inferences for the true population, see Paul P. Biemer & Sharon L. Christ, *Constructing The Survey Weights*, in *Sampling of Populations: Methods and Applications* 489, 489-94 (Paul S. Levy & Stanley Lemeshow eds., 2008); Danny Pfeffermann, *The Role of Sampling Weights When Modeling Survey Data*, 61 *Int'l Stat. Rev.* 317, 329-34 (1993).

¹⁵ Amos Tversky & Daniel Kahneman, *Judgment Under Uncertainty: Heuristics and Biases*, in *Utility, Probability, and Human Decision Making* 141, 141-62 (Dirk Wendt & Charles Vlek eds., 4th ed. 1973).

example, understate the proportion of multiple visitors in the population.

Similarly, act sampling combined with self-reports about repeated actions can *overstate* the seeming prevalence of repeated action. Imagine that half of McDonald's high school patrons visit once a week and half visit twice a week. If McDonald's surveys 10 percent of its high-school patrons eating at the restaurant and they each report truthfully whether they eat there more than once a week, then the sample response is likely to be biased upward: 10 percent of one-time visitors will be sampled, but 19 percent of the two-time visitors will be sampled, because two-time visitors have roughly twice the chance of being sampled.¹⁶ This act-based sample would suggest that approximately two-thirds of McDonald's high school patrons are two-time visitors even though their actual population prevalence is only one-half.¹⁷

Published studies at times fail to account for the ways in which act-sampling bias positively or negatively distorts estimates of reoffending. For example, the long-standing concern of "undetected recidivism" among sex offenders might be at least partially explained by act-sampling bias.¹⁸ If only a fraction of sex offenses are successfully prosecuted, then among those who seem to be one-time offenders—based solely on their convicted offenses—many will be offenders who offended before or after the convicted offense.

Other studies which combine act sampling with offender surveys are at risk of overstating the repeat offending rate. For example, Weinrott and Saylor report that in their sample of institutionalized sex offenders, the median offender reported assaulting six different

¹⁶ One percent of the two-time visitors will be sampled twice ($0.10 \times 0.10 = 1\%$), 81 percent of the two-timers will be sampled zero times ($0.9 \times 0.9 = 81\%$), and the remaining 18 percent of the two-time goers ($100\% - 1\% - 81\%$) will be sampled once. Thus, 19 percent of all two-time visitors will be sampled ($1\% + 18\%$).

¹⁷ We have $19\% / (19\% + 10\%) = 65.52\%$. Hence, 65.52 percent of the sample will consist of two-time goers even though only 50 percent of the true population consists of two-time goers.

¹⁸ A. Nicholas Groth et al., *Undetected Recidivism among Rapists and Child Molesters*, 28 *Crime & Delinq.* 450, 450–51, 453 (1982).

people.¹⁹ But the authors make no allowances for the prospect of a *positive* bias in the self-reported statistics drawn from act sampling. The subjects in this sample came to be institutionalized because they had been convicted of at least one offense or act. But as in the previous example of truthful self-reports following the act sampling of McDonald's patrons, there is probably a disproportionately higher prevalence of repeat offenders in this sample than in the true population.²⁰

The criminal justice system is replete with other examples of act-sampling bias. Imagine, for instance, a series of police speed traps that occur at quasi-random times and places on a stretch of the local interstate. Because sampling is based on the act of speeding, the sample of people caught speeding will tend to understate the prevalence of repeat speeders in the general population.

These examples show that act-sampling bias is not limited to sexual assault reporting. When a sample is based upon a subject's acts instead of the subjects themselves, then the sample will systematically understate the prevalence of repeat actors in the population. Representative act sampling will produce an unrepresentative pool of subjects. We need to take the additional step of working back from the sampling percentage and the proportion of repeat actors in the sample, as in Table 1 above, to produce informed conclusions about how many offenders are repeat offenders in a society.

¹⁹ Mark R. Weinrott & Maureen Saylor, *Self-Report of Crimes Committed by Sex Offenders*, 6 *J. Interpersonal Violence* 286, 291 (1991).

²⁰ Cf. Gene G. Abel et al., *Self-Reported Sex Crimes of Nonincarcerated Paraphiliacs*, 2 *J. Interpersonal Violence* 3, 21–23 (1987) (noting that self-reported number of offenses was significantly greater than the number of convictions and that many sex crimes are not reported).