

FEATURE

Predicting the Next School Shooting May Never Be Possible

If it's not possible, is threat assessment the best alternative?

SCHOOL SHOOTINGS are at an all-time high. That's according to the National Center for Education Statistics, which has been keeping track of the numbers for about 20 years. The center's analysis shows a recent spike in the number of incidents in which someone brandishes or fires a gun on school property. In 2000–01, the earliest school year for which data are provided, there were 30 such incidents. In the most recently documented year, 2020–21, there were 145 (see Figure 1). Numbers from other sources indicate that the increasing trend has very much continued since then. Although the number of incidents doesn't necessarily track the number of casualties, some of the deadliest shootings on record have occurred in the past few years.

What are schools to do? As a psychologist, my job is at least partly to try to predict human behavior. Is there a “profile” of the typical school shooter that could help us identify those who might commit a shooting in the future? Is there some combination of characteristics and circumstances that pushes people—specifically students—toward acts of extreme violence? And if these questions are unanswered or unanswerable, what can schools do instead to protect their students from these horrific incidents?

Although school-shooting numbers are high now compared to those from one or two decades ago, in an

ADAM NIKLEWICZ

By **STUART RITCHIE**

absolute sense the number of incidents is still tiny. Even 145 incidents in a year is a vanishingly small proportion, given the more than 130,000 schools and many millions of students in the United States (the National Center for Education Statistics notes that school shootings make up less than 3 percent of all youth homicides in the nation). For anyone who wishes to predict such incidents, this immediately throws up one of the most vexing problems in statistics: the issue of the *base rate*.

The problem is this: Imagine you had a highly accurate test for predicting future school shooters—a diagnostic interview with a very perceptive psychologist, say, or a complex machine-learning algorithm. Imagine you rolled it out and tested huge numbers of students to uncover and flag those who were potentially a danger. The fact that there are so few actual shooters in such a large population means that, despite your test’s high level

of accuracy, if it isn’t *perfect*, it will generate a large number of false positives, labeling non-shooters as shooters. As explained by Vox.com, if a 99-percent accurate test is run on a group of 100,000 people, including one genuine future shooter, it will mistakenly collar 1,000 students who have no intention of committing gun violence.

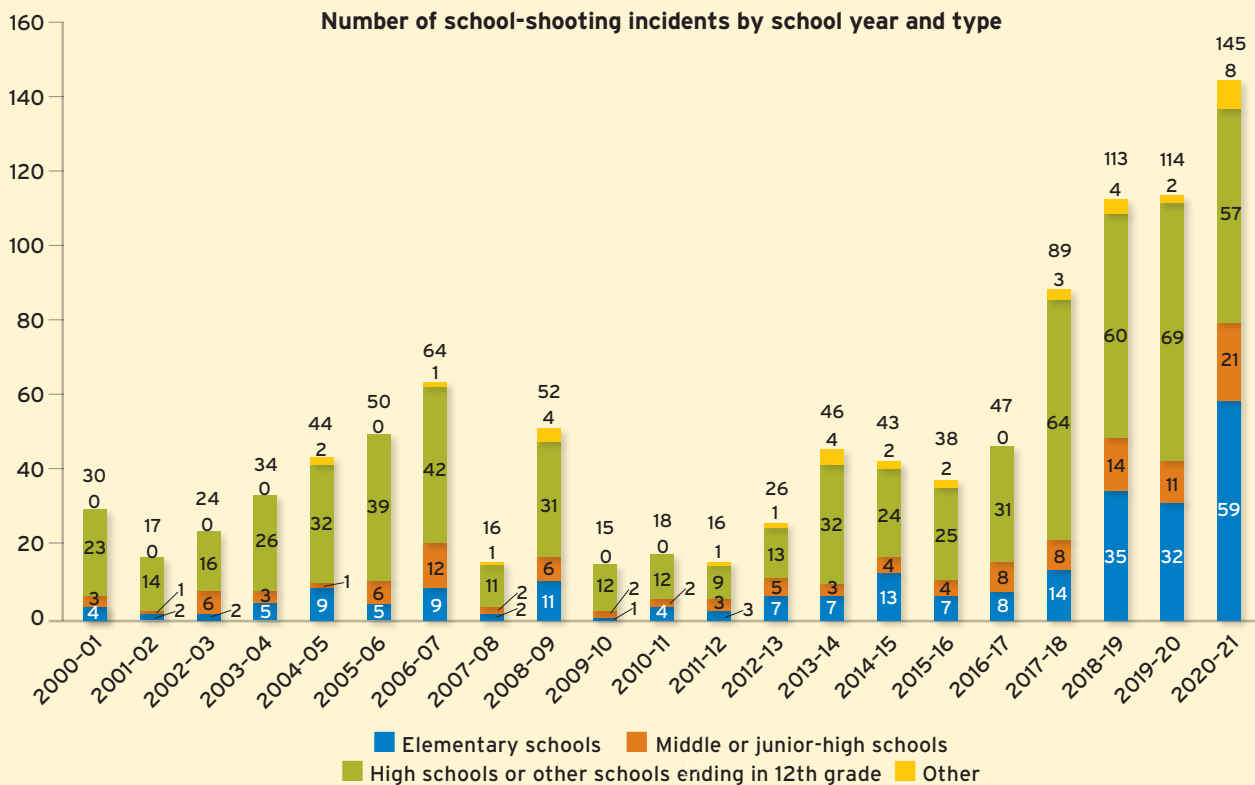
This logic is unintuitive, but it follows simply from the numbers: predicting rare, low-base-rate events is extremely difficult. With so many false positives, schools and police could never hope to identify the true shooter, let alone intervene in some way to avert their course.

Problems with Profiling

After a conference in 1999 where criminal-profiling experts debated this issue, the FBI released a report that argued against

School Shootings on the Rise (Figure 1)

During the 2000–2001 school year, there were 30 school shootings. In 2020–2021, the most recently documented school year, there were 145. Shootings occur most often in high schools or other schools ending with the 12th grade.



SOURCE: National Center for Education Statistics. (2022). Violent Deaths at School and Away From School and School Shootings. *Condition of Education*. U.S. Department of Education, Institute of Education Sciences, <https://nces.ed.gov/programs/coe/indicator/a01>.

attempting to produce a profile of the “typical” school shooter. The problem of the low base rate means that almost any profile that might predict the likelihood to commit a school shooting—the stereotype would perhaps be a white male loner with anger-management issues, suicidal tendencies, and heavy involvement in a niche subculture—will also be shared by a considerable number of other students. Tarring them all with the brush of “future school shooter” would be unfair, to say the least—for

though in an absolute sense the correlation between past and future antisocial behavior was no more than moderate. Meta-analyses are infamously only as good as the studies they include, but this result is highly plausible: disagreeable and violent tendencies are relatively stable across a lifetime. Yet it is dispiriting that our best, most comprehensive, and most up-to-date efforts at prediction produce such obvious answers (“earlier violent behavior predicts later violent behavior”).

More than two decades after the FBI cautioned against profiling, the agency’s advice remains true: attempting to profile a school shooter produces far more questions than answers.

reasons that will be familiar to readers of Philip K. Dick’s “The Minority Report”—and would also be useless in a practical sense.

What’s more, to rely on this profile would be to overlook a large number of potential school shooters. The impression from the media, understandably focused on mass-casualty spree killings like Columbine, Sandy Hook, Parkland, and Uvalde, is that school shooters typically fit the profile I sketched above. But a 2021 study by Sarah Gammell and colleagues looked at all incidents in which a gun was fired on school property, during or around the school day, between 1970 and 2020—of which there were 785—and found that many school shootings are not at all like the most sensational, highly publicized incidents. Thirty-seven percent of shooters were adults; just 14 percent turned the gun on themselves. Across the subset of 276 cases in which the race of at least one of the shooters was reported, a minority of shooters (44 percent) were white. Moreover, and despite the intense media focus on these weapons, only 8 percent of the shootings involved a rifle; more than 75 percent involved a handgun. Finally, just over half, 52 percent, occurred outside the school building—though incidents that occurred inside tended to be deadlier.

The wide variety of school shooting events—spree incidents, gang-related violence, escalated personal rivalries, and more—further complicates any attempt to profile a typical shooter. For some types of shootings, the reliable predictors might be similar to those that predict all kinds of youth violence: after all, it is well understood that antisocial behavior generalizes across contexts, and so its predictors should, as well. These predictors include low academic achievement, deviant peer groups, poor social skills, or substance abuse, among many others.

The best predictor of future school violence, though, is an obvious one: prior antisocial behavior. This was the headline result of a 2022 meta-analysis by Jillian Turanovic and colleagues, who systematically reviewed the entirety of the scientific literature—761 studies—on predicting many different kinds of school violence. No other predictor came close,

There is also a worrying detail in this meta-analysis. The results described above pertain to school violence *in general*. When the study authors restricted their analysis to the prediction of students bringing a weapon to school, the correlation grew far weaker. The meta-analysts concluded that “weapon carrying may be more difficult to predict based on youths’ past behaviors or participation in other forms of school aggression or delinquency.”

Thus, the more-specific, extreme act of bringing a weapon to school—let alone using it in a shooting—is simply not something we can predict satisfactorily, even in the more abstracted context of a study. For making real-life predictions, or predictions about specific individuals, all bets are off. More than two decades after the FBI cautioned against profiling, the agency’s advice remains true: attempting to profile a school shooter produces far more questions than answers.

Do Other Strategies Work?

So, if profiling is a fool’s errand, what kinds of policies should schools adopt to reduce the risk of a mass-shooting incident?

One option is “target hardening”—the incorporation of metal detectors, locked doors, security guards, cameras, and other means to make the school a more difficult place for a would-be shooter to assault. As Bryan Warnick and Ryan Kapa have argued in *Education Next*, target hardening also has a meager evidence base supporting its effectiveness, and might, in at least some cases, have adverse consequences (see “Protecting Students from Gun Violence,” *features*, Spring 2019). The authors argued that schools with more target hardening have higher levels of student fear, anxiety, and alienation from the school in general—though it should be noted that the evidence for these effects is, in a now-familiar pattern for the research in this area, itself rather thin.

Despite the lack of evidence in favor of target hardening, the federal government and some individual states together committed \$800 million in 2018 to increase target-hardening measures.

Another option is a “zero-tolerance” approach—the decision that even minor acts of rule breaking that could potentially relate to future violence should be punished harshly and similarly to much more severe infractions. An analogy is to the “broken-windows” policing used in New York City, among other places, in the 1990s. The broken-windows theory holds that visible signs of violence and misbehavior in a neighborhood often incite more violence and misdeeds. In policing, the idea was that Draconian punishment for misdemeanors like jaywalking contributed to a climate where other, more-serious crimes also began to decline (though the evidence examining this proposition is mixed).

Zero tolerance has been the default discipline policy in many schools since the mid-1990s. Many parents favor the approach,



At Ecorse High School in Michigan, an example of target hardening: students go through a metal detector and have their bags searched by security as they enter the school building.

since they perceive that potential threats will be swiftly dealt with in their children’s school. Critics, though, point to the very high level of suspensions and expulsions that arise under such a policy and argue that not only is the evidence on the policy’s effectiveness unclear, but also, in some cases, this approach might even backfire. One example is the 1998 Thurston High School shooting in Springfield, Oregon, where an expulsion seemed to have been among the primary triggers for the perpetrator’s decision to act. The policy is also by design rigid and inflexible, and there are many examples of students being suspended or expelled for naïve but ultimately innocent acts, such as bringing toy guns or camping utensils to school.

What’s more, the evidence implies that prior acts of violence are not strong predictors of later antisocial behavior that

involves bringing (real) weapons to school. We might thus expect that a zero-tolerance policy, even if it did work to keep general school-violence rates low, would still fail to prevent the rarer, more-extreme bursts of violence that characterize school shootings.

Threat Assessment

A final option is to *wait for the potential shooter to make the first move*. That is, instead of attempting to predict which students might commit a school shooting or other form of violence, and instead of expelling every minor or major rule-breaker, schools can wait for students to *threaten* to commit violence and respond immediately. A 2002 report from the Secret Service noted that, in 81 percent of the school-violence cases they analyzed, “at least one person had information that the attacker was thinking about or planning the school attack.” That percentage needs updating with data from the past two decades, but even if it turns out to be considerably lower, it remains true that, in a substantial number of cases, we have prospective information—not mere “he-was-always-very-suspicious” hindsight. This is information that schools could act upon.

Having information on a threat immediately narrows the field of possible shooter candidates and gets around the base-rate needle-in-a-haystack problem of predicting who will become a shooter in a very large pool of individuals. Indeed, the entire premise of the “threat assessment” perspective, as it is known, is that prediction *in general* is not a viable option. Professor Dewey Cornell of the University of Virginia, the most prominent proponent of threat assessment and the author of its most commonly used variant, the Comprehensive School Threat Assessment Guidelines (CSTAG), argues that “the barriers to conducting rigorous research on the prediction of mass violence seem insurmountable.”

Instead of proactive prediction, then, Cornell advises schools to be reactive—but in a methodical, structured way. Upon discovering a threat of school violence, the CSTAG system recommends that a team of experts in each school follow these steps:

- 1) Interview everyone involved: the student who made the threat, the target(s) of the threat, and any witnesses.
- 2) Decide whether the threat is credible. Some threats may be attempts at dark humor, online trolling, or cries for help that don’t reveal a true underlying danger. These are

Despite the lack of evidence in favor of target hardening, the federal government and some individual states together committed \$800 million in 2018 to increase target-hardening measures.

considered “transient” and are handled with a light touch.

3) For more-substantive, credible threats, warn and protect the targets, make an attempt to resolve any conflict, and discipline the threatener.

4) For the most serious threats—for example, specific threats to carry out a shooting or to murder an individual—contact law enforcement and/or mental health services.

5) Make a “safety plan” and keep monitoring the student in the months after the threat is made.

Although this policy has the strong advantage of highlighting potential cases that might explode into violence, and discourages overreaction and unnecessary intervention, it does have some obvious limitations. As noted above, not all shootings are accompanied by a threat. And even *if* at least one person knows about the threat before the shooting, in many cases the person may never report it, perhaps because they fail to take it seriously, or because of intimidation (one commonly discussed element of threat assessment is how to encourage these witnesses to come forward). One can also imagine a future scenario where more-savvy potential shooters are aware that threats are taken seriously and so keep quiet about their plans.

Nevertheless, the research on threat assessment is far more detailed and substantive than that on other approaches. There are correlational studies, quasi-experimental studies, and even a randomized controlled trial that compared threat assessment to business-as-usual zero-tolerance discipline. But once again, we hit the inherent limitations of conducting research on rare occurrences like school shootings: the researchers, in reporting the outcomes of their studies, have often had to use a variety of proxies rather than measures of actual violence.

For example, there are studies showing students in schools that use threat assessment are more likely to seek help; more likely to feel fairly treated; and have lower levels of suspensions. The randomized controlled trial, similarly, showed lower likelihoods of suspensions in 100 students who made threats in threat-assessment schools compared with 101 students in other schools. But since meting out

fewer suspensions is explicitly how the policy works—after all, it is the alternative to the zero-tolerance approach—these results seem more like confirmation that schools are implementing the policy properly; they do not directly answer the question of whether the policy reduces levels of violence (or indeed, reduces the likelihood of a school shooting). As Cornell himself notes, there have been no K–12 school shootings in Virginia since 1998, “but this cannot be attributed to the widespread use of threat assessment, a practice that was initiated in 2001 and was legally mandated statewide in 2013.”

Certainly, other correlational studies find evidence that schools that choose threat assessment tend to have lower levels of proxy measures such as student-reported bullying. But correlational studies are vulnerable to confounding—that is, it’s possible that schools that adopt the threat-assessment approach differ from business-as-usual schools on important variables (such as parents’ income), and that these other factors are the real drivers of the differences. The researchers attempted to control for this and other factors statistically, but controlling for variables after the fact can often, as statisticians have noted, yield misleading results.

Thus, what we need are more studies where schools (or students) are randomized to different conditions to provide *causal* evidence on the impact of these policies. For that reason it’s dismaying that the only randomized trial of threat assessment occurred as long ago as 2012 (there is evidence that people disapprove of randomized policy-based experiments in general, which might partly explain the apparent lack of interest in further research). What’s more, also essentially all the research on threat assessment has been carried out by Cornell, who wrote the guidelines. It is no criticism of the quality of the existing work to say that one hopes that other, independent researchers will run their own studies, to provide the detailed and varied evidence required to fully evaluate the policy.

Now is perhaps a good time to perform this research. The Bipartisan Safer Communities Act, signed into law by President Biden in June 2022, includes—among various provisions relating to mental health and gun safety—\$200 million to follow up on



Professor Dewey Cornell of the University of Virginia is a proponent of threat assessment.

Having information on a threat immediately narrows the field of possible shooter candidates and gets around the needle-in-a-haystack problem of predicting who will become a shooter.

2019's STOP School Violence Act, which focused on support and training for threat-assessment policies in schools. One hopes that rigorous evaluation of such policies will be part of this funding initiative.

Where to from Here?

In the meantime, should we give up on prediction entirely? In a recent working paper, University of Michigan economist Sara Heller and colleagues used new statistical techniques on a large dataset provided by the Chicago Police Department to build a model predicting who is most likely to become a shooting victim. Their result was pithily described in the title of their paper: "Machine Learning Can Predict Shooting Victimization

posts, and more—it's not inconceivable. In a recent study, University of Pennsylvania researchers Richard Berk and Susan Sorenson used novel statistical methods to predict intimate-partner violence. Someday, we might see advances in scientific methodology that allow us to surmount the fundamental problems in predicting much rarer events, such as school shootings. But just look at how even Berk and Sorenson talk about their new methodology. Even though they are optimistic, they argue that their results should not:

be construed as a powerful justification for our procedures nor for proceeding quickly toward policy implementation. There are many moving parts, each of which needs to be skeptically examined and subjected to much more empirical testing.



Students run from Robb Elementary School in Uvalde, Texas, during the mass shooting there on May 24, 2022. Two adults and 19 children were murdered.

Well Enough to Help Prevent It." Their model identified a small subset of people who were at disproportionately high risk of being shot: 0.02 percent of the population who accounted for 1.9 percent of shooting victims. They argued that their classification was accurate enough to justify government spending on social-service interventions for this population, in order to reduce their risk of being shot.

Could a similar model be developed in the future for the perpetrators of shootings, and then for mass-shooting perpetrators? In a world where ever-more data is collected on individuals—via cellphones, GPS, cameras, internet-search histories, social-media

It is refreshing to see scientists with a cutting-edge method engaging in the opposite of hype, urging caution about their research. And the acknowledgement of the many moving parts reminds us that, when it comes to preventing gun violence in schools, school policy is just one part of a complex picture that also includes failures of law-enforcement procedure (most recently and tragically highlighted at the Uvalde shooting) and the myriad firearm laws and loopholes across different states, among other factors. All of these aspects deserve the attention of policymakers and researchers.

Predicting rare events, at least for the moment, hits up against fundamental limitations of psychology in particular and statistical analysis more generally. Despite some tiny glimmers of hope, education leaders and social scientists may never be able to anticipate school shootings with any useful degree of accuracy. For that reason, responding to threats as they arise might be our best bet. But to be sure about this, we urgently need a better, fuller set of research evaluating the threat-assessment perspective and any specific variations on its basic themes.

School shootings are on the rise. Research on how to prevent them should be, too.

Stuart Ritchie is a senior lecturer at the Institute of Psychiatry, Psychology and Neuroscience at King's College London. This article benefitted from additional research by Anna Wood.