In non-piracy copyright infringement cases, especially those involving music, it's common for plaintiffs to argue that the similarities between the parties' creative works are so great that it is simply implausible that the defendant's work was created without copying from the plaintiff's work. Not all copying is unlawful, but this argument is often deployed both to prove the fact of copying and also to convince a factfinder that the copying was significant enough to constitute infringement. And courts are sometimes sympathetic to these arguments.

The trouble is that, in its present form, the argument is mathematically illiterate: It assumes, without any underlying evidence, that we *know* how likely it is that a song with similarity level X to another, earlier song was created without copying from the earlier song. Our argument is modest: Until the state of the underlying art changes, it is reasonable for experts to testify about similarities between works, but it is unsupported and unreasonable for them to testify about the likelihood that those similarities came about from copying. We don't know that likelihood in the absence of evidence about base rates: how common is it for a song to have similarity level X with some other song in the corpus of existing songs, and how common it is for that similarity to come from copying or from independent creation (or from both copying a shared antecedent). Although it is conceivable that we could develop metrics to measure the first probability, and we could likely improve ways to think about the second one, that work has not yet been done. Until it is, testimony about the likelihood that copying occurred is likely to mislead factfinders.

The problem is well-understood in other contexts. For example, where a serious disease is very rare, and there is a very accurate diagnostic test for it, one might think that everyone should be tested for it. But with universal screening for a rare disease—a low base rate, say 1 in 5 million—and a diagnostic test that is even slightly imperfect—say, it generates "false positives" 1% of the time—it is mathematically certain that most positive results on the test will be false positives. Thus, a person who tests positive for the disease is more likely than not to *not* have the disease, even though the test is almost perfectly accurate. It is for this reason that public health professionals generally don't recommend universal screening for very rare diseases. Many people who hear that a 99% accurate test was positive for a fatal fetal anomaly understandably believe that the result means that it is overwhelmingly likely that the fetus is afflicted. But it is not, unless screening is limited to those already known to be at risk.

Our argument is not that we know the base rate of similarity or the accuracy of experts' assessments of how often similarity results from copying (false positives). To the contrary, it is that we *don't* know. Until experts are able to say more about base rates and rates of copying, they should be saying less.