

The Problem of Biased Experts, and Blinding as a Solution: A Response to Professor Gelbach

Christopher Robertson[†]

INTRODUCTION

In a recent symposium article, Professor Jonah Gelbach discusses the problem that a litigant in the American adversarial system can consult multiple expert witnesses on a given question but only disclose the single most favorable opinion to the fact finder (a jury, judge, or arbitrator).¹ He calls this the problem of “expert mining.” In particular, Gelbach considers whether a policy that requires litigants to disclose to the fact finder the number of experts that they consulted might be a satisfactory solution to the problem.² Alternatively, Gelbach considers whether an even more radical change to the American litigation system—the exclusion of all expert opinions rendered after the first one—might be necessary.³ In doing so, Gelbach extensively discusses my own work on this problem and the third solution I developed in a 2010 article, *Blind Expertise*.⁴ There, I show that expert mining is one part of a broader problem of expert bias, and I propose a *conditional*-disclosure rule as the solution.⁵ This Essay provides some analysis of Gelbach’s framing of the problem, reviews the blinding proposal, and identifies the limits of Gelbach’s analyses.

[†] Associate Professor of Law, University of Arizona James E. Rogers College of Law, crobotson@post.harvard.edu. I thank Professors Jonah Gelbach, Jim Greiner, and David Marcus for helpful comments on prior drafts, and I thank Harvard Law School, where I served as a visiting professor during the drafting of this Essay.

¹ Jonah B. Gelbach, *Expert Mining and Required Disclosure*, 81 U Chi L Rev 131, 131–32 (2014).

² Id at 133.

³ Id.

⁴ Christopher Tarver Robertson, *Blind Expertise*, 85 NYU L Rev 174 (2010).

⁵ Id at 178–80.

I. EXPERT MINING AS ONE MECHANISM OF SELECTION BIAS

Professor Gelbach draws an analogy from the tactic of an expert conducting multiple soil-sample tests and then reporting only the favorable results (which he calls “individual data mining”) to the tactic of litigants consulting multiple experts but designating only one for testimony (which he calls “expert mining”).⁶ Gelbach argues that the two tactics have “the same statistical properties,” since either one can allow a party to produce a favorable result that may mislead a fact finder.⁷

However, these two practices have quite different epistemic, economic, and strategic properties. First, unlike the results of a valid scientific test, the likely opinions of an expert witness can often be reliably predicted by lawyers prior to consulting the expert. This is because expert witnesses have track records—from prior work with specific attorneys, attorney and expert word of mouth, advertisements in particular media (for example, plaintiffs’ bar magazines), reports of published cases and jury verdicts, publications, and professional affiliations. Accordingly, in many fields of expertise, experts tend to be identified as pro-plaintiff or pro-defendant, with reliable track records for producing opinions for one side or the other.⁸

Second, the cost profile is quite different. A given expert can presumably perform an additional soil test at little or no marginal cost. In contrast, it is much more expensive for a litigant to hire another expert to render a new opinion. Expert opinions have onerous startup costs, because experts and attorneys each charge hundreds of dollars per hour. A new expert will require several hours to learn the background facts of the case and to be directed by the attorneys toward the appropriate legal question before performing whatever research, tests, and analyses may be required to actually form an opinion, which can then be communicated back to the attorney, who must then evaluate its favorability. This process consumes both expert time and attorney time. These costs per expert impede the strategy of expert mining, limiting its practicability.

⁶ Id at 135–36.

⁷ Id.

⁸ See, for example, Aaron S. Kesselheim and David M. Studdert, *Characteristics of Physicians Who Frequently Act as Expert Witnesses in Neurologic Birth Injury Litigation*, 108 *Obstetrics & Gynecology* 273, 275 (2006) (showing that only 21 percent of frequent experts in neurologic birth injury cases “approached an even split of their caseload between plaintiffs and defendants”).

Finally, expert mining is somewhat risky. Gelbach himself notes that, with regard to expert mining, “what little case law exists is mixed.”⁹ Indeed, he argues that it is “an open question whether the Rules as written allow discovery of nontestifying experts’ identities.”¹⁰ In my view, the attorney work product protections are quite substantial.¹¹ However, there is still some risk to a litigant that an expert-mining strategy will be revealed and thereby used against the litigant at trial.

Thus, for all these reasons, a competent attorney often does not draw randomly from the field of potential expert witnesses—planning to simply draw again if the first one is unfavorable—as that strategy would be wasteful and slightly risky. Instead, she handpicks the one expert that is most likely to render a favorable opinion, either because that expert is known to have strong priors that tend to favor the litigant’s side or because the expert is known to be the most malleable to that litigant. As long as a biased, malleable, or fair-but-luckily-favorable expert can be found on the first try, a litigant can proceed to mislead the jury.

In this way, the analogy from the real world of expert witnesses to data mining is rather weak, or perhaps applicable only in a specialized set of cases in which the foregoing issues do not arise. Accordingly, the potential reforms Gelbach considers—disclosure of the number of experts consulted or a proscription on subsequent opinions—are both rather weak solutions. Under either regime, a competent attorney could still handpick a favorable and unrepresentative expert on her first draw.

Instead, as I argued in 2010, the broader concept of selection bias provides a stronger framework for thinking about the epistemological and strategic issues in this domain.¹² Regardless of whether it is through a mechanism of handpicking and shaping one expert or mining through multiple opinions, the impact on the fact finder is the same. The ultimate judgment will be based on the disclosed sample of unrepresentative expert opinions rather than the population of all opinions that were rendered or could have been rendered through an unbiased sampling process. A fact finder may thus be misled by the revealed opinions, making false inferences about the state of knowledge in a given field.

⁹ Gelbach, 81 U Chi L Rev at 132 (cited in note 1).

¹⁰ Id at 135.

¹¹ See Robertson, 85 NYU L Rev at 210, 213 & n 178 (cited in note 4).

¹² Id at 184–85.

To be sure, the adversarial process does not solve this problem, because in any given case it reveals two disagreeing experts, which overrepresents whichever side is in the minority of expert opinion. A more robust solution is required.

II. THE BLINDED-EXPERT PROTOCOL

The blinded-expert protocol solves the selection bias problem and the malleability problem. (In *Blind Expertise*, I explained the malleability problem in terms of “affiliation bias” and “compensation bias.”¹³). Under this protocol, either litigant may unilaterally and confidentially decide whether to utilize an intermediary between itself and its potential expert witness. If the litigant chooses to pay in advance for such a blinded-expert opinion, the blinding intermediary would select an expert witness in an unbiased way. One way to perform this selection would be to choose randomly from a prescreened list of competent experts. To further prevent the expert from being manipulated by the litigant, the intermediary would ensure that the expert renders an opinion on the case without knowing which litigant requested the opinion. This would solve the problem of expert malleability.

Laboratory experiments have shown that this reform improves experts’ credibility with the fact finder, yielding more favorable outcomes at trial for litigants that choose this option.¹⁴ A litigant’s use of a blinded expert is thus likely to drive a more favorable settlement, making it a rational strategy, which also happens to improve litigation accuracy and legitimacy.

Most pertinent to Professor Gelbach’s analysis, the blind-expert protocol includes a *conditional*-disclosure rule. If the opinion turns out to be unfavorable, the litigant need not disclose that opinion at all, and may instead proceed with a traditional unblinded expert (or settle the case). Allowing litigants to retain discretion about whether to disclose the expert opinion minimizes the risk of blinding to the litigant, making it more likely to be utilized in a system dependent on the choices of rational actors.

Nonetheless, *if* a litigant receives a favorable blinded-expert opinion *and* proposes to use it at trial, the litigant would be

¹³ Id at 184–88.

¹⁴ Christopher T. Robertson and David V. Yokum, *The Effect of Blinded Experts on Juror Verdicts*, 9 J Empirical Legal Stud 765, 786 (2012).

required to disclose all the other blinded opinions that it procured on the given question. A primary advantage of this reform is that it does not require any changes to rules of evidence or procedure; the disclosure mandate is already implicit in the waiver provisions around the attorney work product doctrine.¹⁵ In short, litigants would not be allowed to claim that their expert was unbiased while also exploiting a selection bias.

Ultimately, Gelbach agrees that this disclosure mechanism would work: “a well-crafted cross-examination would undermine the credibility of the adversary’s expert evidence when many experts are consulted. This would at least reduce the incentive to use expert mining.”¹⁶ Notice that there are two dynamics at play here. Ex post, the disclosure rule prevents the fact finder from being misled by the expert-mining tactic, because the fact finder can discount the favorable blinded opinion in light of the unfavorable blinded opinions. Thus, the reliability of the blinding mechanism is preserved as it prevents a selection bias. As Gelbach writes, “To the extent that additional *fully disclosed* expert testimony increases the fact finder’s information, we can expect a beneficial increase in accuracy.”¹⁷

Ex ante, the disclosure rule also reduces the incentive for litigants to undertake expert mining within the blinding system. Conditional upon getting an unfavorable initial blinded-expert opinion, the litigant faces a choice between (a) reverting back to the traditional litigation process with an unblinded expert (hiding the unfavorable blinded opinion in attorney work product protections); or (b) taking another unbiased draw from the blinded-expert pool, in the hope that the second blinded expert will disagree with the first (if so, both opinions will then be presented to the fact finder). Option (b) would generally be unattractive to litigants, because the second expert opinion has a cost, and it is unlikely to disagree with the first opinion (assuming that the blinded experts have a modicum of accuracy).¹⁸ Most importantly, even if the second blinded expert disagrees with the first, a pair of counterpoised blinded experts would not be very persuasive to

¹⁵ See Robertson, 85 NYU L Rev at 211–12 & nn 175, 176 (cited in note 4).

¹⁶ Gelbach, 81 U Chi L Rev at 145 (cited in note 1).

¹⁷ Id at 133.

¹⁸ The first blinded opinion has epistemic content for the litigant himself. Assuming that issues in litigation have truth values, which can be detected reliably by blinded experts, the first blinded-expert opinion is the litigant’s best evidence of what a second blinded expert would say. Thus, a first unfavorable opinion portends a second unfavorable opinion, making the expert-mining strategy unattractive.

the fact finder, particularly compared to a single unblinded expert who could be hired instead under option (a).

Thus, if a litigant's first blinded expert provides an unfavorable opinion, the litigant is likely to revert back to option (a)—the traditional litigation process with unblinded experts—or settle the case. Expert mining is unlikely to occur within the blinding system given the conditional-disclosure rule that I have specified. Thus, the testimony of a blinded expert will be a strong signal to the truth, one that largely solves the expert-mining problem as an instance of the broader problem of expert bias.

III. THE RELIABILITY OF BLINDED EXPERTS

Professor Gelbach takes issue with some of this analysis.¹⁹ Analogizing to the false positives and false negatives of soil tests, Gelbach supposes that each expert witness likewise has *two* rates of error, depending on whether he renders positive or negative opinions.²⁰ For such scientific tests, the rates of error can be determined by running the tests repeatedly on cases with known truth. Gelbach further seems to assume that the litigation fact finder has prior knowledge about divergent error rates and thus will discount certain opinions as less reliable, depending on their substance.²¹ Gelbach argues that, if a litigant knew about the fact finder's prior belief about this disparity in the two error rates, then the conditional-disclosure rule might not deter

¹⁹ Gelbach asserts that "Robertson is mistaken in suggesting that there is 'No Signal' sent to the fact finder when each party introduces blinded-expert evidence in its favor." Gelbach, 81 U Chi L Rev at 141–42 (cited in note 1). Gelbach here cites to another part of my paper where, for exegetical purposes, I postulated that a given sort of blinded expert might have a (single) error rate (say, 5 percent), and then demonstrated how adversarial use of the blinding procedure could dramatically reduce the rate of errors. *Id.* at 142 n 42, citing Robertson, 85 NYU L Rev at 217 (cited in note 4). In rare cases in which each side procures a blinded opinion that is favorable to that side, the fact finder would receive two disagreeing blinded-expert opinions. From the perspective of the analyst, the blind procedure thus produces no signal on net and thus cannot lead the fact finder astray. This adversarial use of the blinded procedure is one of its virtues compared to a single court-appointed expert.

²⁰ The analogy is to a "soil test on uncontaminated soil." Gelbach, 81 U Chi L Rev at 136 (cited in note 1). Gelbach is correct to observe that "[e]ven two test results pointing in opposite directions *can* be very informative, *if* the [known] probabilities of false negatives and false positives are sufficiently different." *Id.* at 141 (emphasis added). The bracketed insert is essential for Gelbach's epistemological point to be true.

²¹ *Id.* The blinded-expert protocol does not require such knowledge of the fact finder. It simply requires that the fact finder regard blinded experts as more persuasive than unblinded experts.

a rational litigant from drawing again from the blind-expert pool after receiving a first unfavorable opinion.²² Instead, he would go back to the pool of blinded experts (that is, he would engage in expert mining), hoping for a favorable opinion that would not be fully offset by the unfavorable, discounted opinion.²³

It bears emphasis that Gelbach and I agree that the conditional-disclosure rule in the blinded-expert protocol would nonetheless prevent the tactic from misleading the fact finder.²⁴ As all blinded opinions will be disclosed to the fact finder, there is no selection bias.

However, we apparently disagree on whether expert mining would be likely to occur within the blinded-expert protocol at all. Perhaps there are expert opinions like those that Gelbach has in mind, in which the expert opinion itself has known rates of error that diverge depending on the substance of the opinion (that is, its positive or negative character). These cases would be distinct from those in which the expert has made all sorts of discretionary judgments about methods, datasets, samples, and criteria, or is otherwise simply rendering an all-things-considered opinion based on his experience and judgment.²⁵ For such an expert opinion, we do not have reference cases or a history of repeated opinions, which could be used to discern the differential rates of positive and negative errors.

The challenge, then, is to discern whether Gelbach's objection is more than a theoretical possibility and, if it is real, whether it constitutes the exception or the rule. Gelbach provides no estimates of the prevalence of the sorts of cases he has in mind. Indeed, a casual survey of common types of expert testimony suggests that it is difficult to find such a case. Consider, for example, the physician in a medical malpractice case who renders an opinion about whether the defendant met the requisite standard of care. Or, consider the physician in the ubiquitous auto-accident case or disability-benefits case, who must evaluate whether the plaintiff is partially or permanently disabled. With these sorts of medical-expert opinions, we have already accounted for a huge proportion of real world cases, and

²² Id at 141–42.

²³ Id at 141 (“While Robertson suggests that, in general, required disclosure would be sufficient to eliminate mining of blinded experts, this is too strong a claim.”).

²⁴ See Gelbach, 81 U Chi L Rev at 145 (cited in note 1).

²⁵ See id at 147–48 (acknowledging “many other important contexts . . . [in which there is] variation in experts’ good-faith opinions regarding subjective matters”).

there is no suggestion of known differential rates of false positives and negatives.²⁶

Or, consider the economist or accountant who opines about the total amount of damages suffered, scaling them up for inflation and down for the time value of money.²⁷ Although we have moved into the realm of the numerical, these estimates of a continuous variable (money) also do not seem amenable to Gelbach's notion of positive- and negative-error testing.

There are some domains in which we might expect, or at least hope, that there are known rates of false positives and false negatives.²⁸ For example, fingerprint experts are ubiquitous in criminal litigation and some civil litigation.²⁹ In principle, it would be easy enough to use known cases to create error-rate estimates for fingerprint examiners as a whole, or for a particular examiner in a particular sort of case. (Note the difficulty in even specifying the relevant basis for computed error rates.) However, the profession of fingerprint examiners has largely refused to adopt such scientific methods, preferring instead to insist on infallibility.³⁰ Even if the discourse surrounding this sort of expert becomes more sophisticated, there will likely remain dispute about the very validity of error rates themselves.³¹ Thus, it is not clear how a fact finder would implement the sort of discounting function that Gelbach has in mind.

²⁶ See Samuel R. Gross, *Expert Evidence*, 1991 Wis L Rev 1113, 1119 (“Half of the experts in our data were medical doctors, and an additional 9% were other medical professionals—clinical psychologists, rehabilitation specialists, dentists, etc.”).

²⁷ See *id.* (noting that experts involved in “various aspects of business and finance” accounted for 11 percent of the experts in the sample).

²⁸ I have suggested elsewhere that in certain fields (such as radiology), the facts of an actual litigation case could itself be reviewed by multiple experts and compared against known cases, thereby allowing estimation of something like an error rate for the particular case. Even this method is, however, distinct from suggesting that the expert conveying this information has a knowable error rate. See Daniel J. Durand, et al, *Expert Witness Blinding Strategies to Mitigate Bias in Radiology Malpractice Cases: A Comprehensive Review of the Literature*, 11 J Am Coll Radiology *3–4 (forthcoming 2014), online at [http://www.jacr.org/article/S1546-1440\(14\)00248-8/abstract](http://www.jacr.org/article/S1546-1440(14)00248-8/abstract) (visited July 30, 2014).

²⁹ See 36 Am Jur Proof of Facts 2d 285 § 1 (2014).

³⁰ See *United States v Mitchell*, 365 F3d 215, 228 (3d Cir 2004) (discussing the testimony of Simon Cole, who was critical of the field of fingerprint examination in part because it “did not recognize error rates” but instead attributed all errors to the incompetence or corruption of particular examiners). See also *id.* at 231 (mentioning testimony about “the limited studies performed specifically to establish an error rate for fingerprint identification”).

³¹ *Id.* at 239 (“[T]he existence of any error rate at all seems strongly disputed by some latent fingerprint examiners.”).

A more profound point is that, even when an expert uses an objective test with a known rate of error, the human factor often intervenes in ways that are difficult to observe, much less quantify.³² In the fingerprint domain in particular, research has shown that the contextual information that police routinely provide to their examiners can color their decisions.³³ The same fingerprint examiner has been shown to render inconsistent opinions on the same task depending on the contextual information at hand.³⁴ Similarly, we might suppose that forensic DNA testimony would be the paradigmatic case of objectivity. However, it turns out that this domain is also susceptible to cognitive biases, as human factors intervene.³⁵ Perhaps most analogous to Gelbach's soil-sampling example is the forensic testing of contraband, which is used to determine whether the contraband is an illegal drug. Even here, chemists have been known to violate the test assumptions by sampling nonhomogeneous populations.³⁶ In those contexts, the optimal "known" error rates are therefore unreliable.

One might hope that rigorous thinking about error rates would come from econometricians or statisticians, such as those in a class action case, who must determine whether racialized block voting exists in a jurisdiction, or determine whether minority workers are suffering a disparate impact from their employer's policy. Here, at least, we might hope that the expert will produce confidence intervals or, even better, a probability distribution across all possible values from which the fact finder could compute something like a rate of false positives or false negatives. However, these error-rate estimates are themselves likely to be biased by the expert. Professor Jim Greiner explains three problems: "[F]irst, an analyst fitting a regression sees the litigation answer before she assesses goodness of fit; second, deciding whether a model is adequate for the data requires judgment;

³² See Itiel E. Dror, David Charlton, and Ailsa E. Péron, *Contextual Information Renders Experts Vulnerable to Making Erroneous Identifications*, 156 *Forensic Sci Intl* 74, 76 (2006).

³³ See Itiel E. Dror, et al, *Cognitive Issues in Fingerprint Analysis: Inter- and Intra-expert Consistency and the Effect of a "Target" Comparison*, 208 *Forensic Sci Intl* 10, 11 (2011).

³⁴ *Id.*

³⁵ Itiel E. Dror and Greg Hampikian, *Subjectivity and Bias in Forensic DNA Mixture Interpretation*, 51 *Sci & Just* 204, 205 (2011).

³⁶ Andrea Widener and Carmen Drahl, *Forcing Change in Forensic Science*, 92 *Chem & Eng News* 10, 14 (May 12, 2014).

and third, adding or removing variables from a regression can result in wholesale changes to the results.”³⁷

In theory, many of these human factors could be resolved through robust blinding processes. And remaining factors could be incorporated into sophisticated error-rate estimates with significant investments in error-rate testing with known cases. However, those estimates will be accurate only to the extent that they are calibrated to the particular expert rendering an opinion on a particular case in a particular context. Thus, Gelbach’s epistemological point would seem to be correct in a vanishingly small set of cases.

Even there, while it may be theoretically possible for a Bayesian fact finder to integrate the two estimates for each expert’s error rate into a single posterior probability, it seems doubtful that a jury has the capacity or inclination to actually integrate the expert opinions in this way.³⁸ For practical purposes, a litigant may have to assume that the fact finders will take expert opinions at face value, such that a negative and positive opinion roughly cancel each other out, with unpredictable noise directing the outcome.

Nonetheless, perhaps one could make Gelbach’s point about litigation strategy in a less formal way. Rather than think about divergent error rates, let us just recognize the obvious point that some expert opinions will be more persuasive to the fact finder than others, and this will remain true for blinded experts. Therefore, two disagreeing blinded-expert opinions will not necessarily cancel each other out in the eyes of a fact finder. Thus, when a rational litigant receives an unfavorable blinded opinion, it may sometimes perceive that—for whatever reason—the opinion is not likely to be very persuasive to the fact finder. If so, the rational litigant might perceive disclosure of that opinion to be a rather small cost that does not deter a second draw from the blinded-expert pool.

However, even in this looser formulation, this analysis overlooks the alternative strategy of hiding the unfavorable opinion

³⁷ D. James Greiner, *Causal Inference in Civil Rights Litigation*, 122 Harv L Rev 533, 544 (2008).

³⁸ Even highly intelligent experts have consistently been shown to fail to properly integrate information about divergent error rates. See generally A.K. Ghosh, K. Ghosh, and P.J. Erwin, *Do Medical Students and Physicians Understand Probability?*, 97 Q J Med 53 (2004).

and proceeding with a traditional litigation expert instead.³⁹ Even if a litigant accurately predicts that a fact finder will tend to discount an unfavorable blinded opinion, the question is whether that unfavorable blinded opinion will hurt the litigant more than the blinded-expert protocol helps the litigant. Outside of convoluted hypothetical situations in which the litigant knows that the fact finder loves the blinding procedure but hates blinded opinions favoring one particular side, it will usually be better for a litigant to hide the unfavorable blinded opinion and proceed with a traditional unblinded expert instead.

The key point is that traditional litigation remains an outside option. Because the disclosure rule is *conditional* on introducing a blinded expert at trial, it protects the integrity of the blinded-expert procedure and thereby preserves it as an epistemic signal. In contrast, Gelbach's critique is applicable to other non-conditional-disclosure mandates, such as that of Judge Richard Posner.⁴⁰

CONCLUSION

Professor Gelbach's analogy to data mining sheds light on one part of the problem of expert bias. In particular, Gelbach has identified a theoretical situation in which expert mining might occur within the blinded-expert procedure. Yet we agree that, even there, it would be largely harmless.⁴¹ As the protocol requires that all blinded opinions be disclosed to the fact finder, if a litigant discloses any opinion, then the protocol averts selection bias.

Much more commonly, it remains true that a rational litigant that receives an unfavorable blinded opinion will instead opt back into the regime of traditional litigation. In those cases, it is likely that the rational adversary will hire a favorable blinded expert, who will guide the fact finder to the true answer.⁴² Blinding remains a valuable epistemic device, a rational strategy for litigants, and a promising solution for the multifaceted problem of expert bias.

³⁹ Gelbach acknowledges elsewhere in his paper that he "will focus on [only one] aspect" of the blinded-expert proposal. Gelbach, 81 U Chi L Rev at 133 n 10 (cited in note 1). Such a narrow focus may be the source of any confusion.

⁴⁰ See *id.* at 140, discussing Richard A. Posner, *An Economic Approach to the Law of Evidence*, 51 Stan L Rev 1477, 1541 (1999).

⁴¹ Gelbach, 81 U Chi L Rev at 141–42, 149 (cited in note 1).

⁴² See Robertson, 86 NYU L Rev at 217 (cited in note 4).