



ARTICLE

Through the Looking Glass: Reflections on a Pedagogical Experiment

Abigail L. Perdue

Professor of Law

Founding Director, D.C. Summer Judicial Externship Program

Wake Forest University School of Law

Published: August 2025

Have you ever wished that you could more deeply assess the efficacy of a new teaching exercise you developed, an innovative pedagogical method that you tried, or a novel course you taught? Too often, standardized course evaluations do not capture the data necessary to determine the success (or failure) of a pedagogical experiment.

During the spring of 2025, I found myself in exactly this position. After learning more about generative artificial intelligence (AI) at the 2024 Legal Writing Institute (LWI) Biennial and reading some of the numerous articles on generative AI that our exceptional colleagues have published, I attempted to integrate generative AI into an upper-level *Litigation Drafting* course that I was teaching for the very first time. The course met once per week for a marathon two-hour session.

I created a dedicated unit on drafting with generative AI that consumed roughly six instructional hours—four at the outset of the semester and two at the end. Among other things, this unit featured a guest lecture from Professor Dyane O’Leary, Director of Suffolk’s Legal Innovation & Technology Center. A product expert from a commercial

vendor also visited our class to demonstrate how to use the vendor's new AI assistant to complete tasks that students would be performing in the course. Our class also discussed the ethical implications of using generative AI for legal drafting, reviewed court opinions involving AI misuse and case hallucinations, and studied a recent survey measuring attorneys' use of AI in the practice. After we completed the first four hours of this unit, the students were permitted (but neither required nor encouraged) to use the vendor's AI tool in any capacity to complete the subsequent graded exercises in the course. However, they were forbidden to tell me whether AI was used to avoid any inference that its use had impacted my grading, consciously or not.

As the semester progressed, I became increasingly curious. Were students using generative AI, and if so, then how? On which exercises were they employing it and to what extent? Was it helpful or problematic? Did it increase their efficiency or slow their progress?

Before the class began, I had initially planned to seek answers to some of these questions by conducting a reflection exercise in our final class. But after we completed the first part of our unit on generative AI, it became clear that a short in-class reflection would be insufficient to learn the answers to my ever-growing list of questions. Nor could I present my findings outside my institution or publish them unless I attained Institutional Review Board (IRB) approval. So that's what I did. These are the things I wish I'd known before jumping down the rabbit hole of empirical research.

Empirical research involves deriving conclusions from observing (and often measuring) behavior or phenomena. It can encapsulate everything from analyzing recent Supreme Court opinions¹ to conducting a clinical trial to assess a drug's efficacy. Quantitative research aims to collect *numerical* data, such as statistics on how often female Supreme Court Justices are interrupted compared to their male counterparts.² By comparison, qualitative research involves gathering *non-numerical* data that is descriptive.

After you identify the specific, narrow question(s) you want to explore, think deeply about the data you must collect to answer it and which research method is best suited to accomplish that goal. Vet your research question and planned *study* with seasoned empirical researchers. Participate in scholarship incubation workshops and similar

¹ See, e.g., JILL BARTON, *THE SUPREME GUIDE TO WRITING* (2024) (relying on a long-term study of U.S. Supreme Court opinions to suggest best practices in legal writing).

² See, e.g., Tonja Jacobi & Dylan Schweers, *Female Supreme Court Justices Are Interrupted More by Male Justices and Advocates*, HARV. BUS. REV. (Apr. 11, 2017), <https://hbr.org/2017/04/female-supreme-court-justices-are-interrupted-more-by-male-justices-and-advocates>.

opportunities. Reach out to your technology experts to see if there are programs or apps that can streamline your vision for collecting information.

If this is your first adventure in empirical research—or if you are exploring an interdisciplinary question involving doctrine outside your expertise—consider collaborating, but choose your collaborators wisely. My first empirical project³ benefited from a productive partnership with an experienced sociologist and psychologist who not only navigated the IRB approval process for our team but also selected our survey tool.

To help ensure a fruitful collaboration, discuss important questions at the outset. Consider the delegation of responsibilities and how the data will be used. Must every paper produced from the data be *jointly* vetted and published? Must everything always be co-presented? Do the collaborators have veto power on publication or presentation of any results? What is the anticipated timeframe for completion? How will grant funding be allocated among collaborators? These are just a few of the many questions that potential collaborators should thoughtfully consider *before* submitting a study for IRB approval. Although technology has made remote, inter-campus collaboration less onerous, selecting individuals on your campus is often still more convenient.

If, like me, you plan to conduct a survey or collect interviews, you must attain IRB approval because both research methods constitute research involving human subjects. The modern IRB regulatory scheme may seem daunting, but it was created to prevent the kinds of horrific research abuses perpetrated in the not-too-distant past.⁴

In response to these and other abuses, the United States enacted the National Research Act,⁵ which created the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research (Commission).⁶ The Commission later published the Belmont Report, a document highlighting core ethical principles that still guide IRB

³ See generally Abigail L. Perdue, *Transforming “Shedets” Into “Keydets”: An Empirical Study Examining Coeducation through the Lens of Gender Polarization*, 28 COL. J. GENDER & LAW 371, 392 (2014) (discussing findings from an anonymous survey of students at a formerly all-male military college).

⁴ See, e.g., Mark A. Rothstein & Leslie E. Wolf, *National Research Act at 50: An Ethics Landmark in Need of an Update*, HASTINGS CTR. (July 12, 2024), <https://www.thehastingscenter.org/national-research-act-at-50-it-launched-ethics-oversight-but-it-needs-an-update/> (last visited July 1, 2025); Allan Brandt, *Racism and Research: The Case of the Tuskegee Syphilis Study*, 8 HASTINGS CTR. RPT. 21-29 (1978).

⁵ National Research Act, Public Law 93-348, 88 Stat. 342 (1974); see also Rothstein & Wolf, *supra* note 4.

⁶ Rothstein & Wolf, *supra* note 4 (explaining that the Commission was “directed to ‘identify the basic ethical principles which should underlie the conduct of biomedical and behavioral research involving human subjects [and to] develop guidelines...to assure that it is conducted in accordance with such principles.’”).

decision-making.⁷ The first principle—respect for persons—primarily involves protecting participants’ autonomy and privacy by ensuring that they are fully informed of the risks and agree to participate voluntarily without coercion. The second principle—beneficence—requires the study to minimize the risk of harm to participants while maximizing the possible benefits of the study. The third principle—justice—requires the study to allocate its benefits and burdens as equitably as possible to better protect vulnerable groups. Prior to adoption of these principles, the participants who bore the risks typically came from marginalized groups, while the knowledge gained disproportionately benefited non-marginalized groups.

IRBs review research proposals to ensure that they adequately protect participants from similar abuses. Thus, as you design your study, think deeply about these paramount ethical principles. For example, the IRB is less likely to approve a study that will not fully disclose the risks to participants at the outset or that involves participants with diminished autonomy, like children or prisoners. Interestingly, in some circumstances, even law students can be considered a potentially vulnerable population, as in my study where the respondents were students in my class.

Not surprisingly, there are three levels of IRB review: exempt, expedited, and full. If you are simply working with a de-identified data set that another organization or agency has produced, your project is likely exempt, meaning IRB approval will not be required. In most cases, however, if you are interacting with humans, your project will require review. Although my study was eligible for expedited review in part because it involved a minimal risk of harm to subjects, full review is usually required for research that involves vulnerable subjects and/or poses more than a minimal risk of harm.

If you are considering human subjects research, the first step is to schedule a meeting or phone call with your IRB. Vet your research question and incorporate the IRB’s constructive feedback. You will likely be required to complete an IRB certification, which usually involves completion of an intensive online course. You will also create an eIRB Profile, which is the platform through which you will submit your proposal.

The IRB process is more onerous than requesting a summer research grant. For example, if you are conducting interviews, you will be required to submit the consent form you will provide participants, the process and script you will use to recruit participants, a confirmation of your IRB certification, your current C.V., as well as the list of questions you plan to ask in the order you plan to ask them. You will also need to complete a lengthy online application, which asks a diverse array of questions from

⁷ *Id.*

how you plan to protect the privacy of the information your participants provide to the names and contact information of your collaborators (*i.e.*, anyone that will be handling the data). However, you will likely not be required to submit your course syllabus or class materials, only materials that directly pertain to the study. And if any of that information changes after submission or approval, you must amend your application.

Accordingly, you should build in plenty of time for IRB approval. IRB applications are rarely accepted as is, and the review process often takes much longer than anticipated. The investigator must usually engage in a back-and-forth dialogue with the IRB, which will flag its “concerns” and request revisions. For instance, each time you revise your consent form, you must delete the erroneous form and upload the revised one.

Likewise, when you revise an answer on the digital form, you may also be prompted to briefly explain the change. This wearying exchange might take days, weeks, or even months, depending on the complexity and volume of the required revisions. Persevere.

Remember that the IRB is your partner, not your nemesis. The IRB will answer your questions, provide you with helpful resources, and guide you through the (sometimes maddening) process one step at a time. Read IRB resources carefully and follow their suggested templates very closely. Failing to use the template’s exact wording, organization, or formatting, even when it seems discretionary, could prompt rejection of your proposal.

Whenever possible, do not submit a draft that you plan to later revise. If you make even minor changes after attaining IRB approval, such as adding or removing a question, changing answer choices, reordering questions, altering your collaborator, etc., you will likely have to amend your proposal and undergo additional review.

For these reasons, conducting empirical research to assess the efficacy of your pedagogical experiment may take as much or more time than creating the exercise, developing the approach, or designing the course. Indeed, empirical research is arguably more time-consuming than conventional modes of legal scholarship. But it is also often more rewarding and illuminating. Unfortunately, most law schools don’t recognize its complexity, so you will likely be given the same amount of time and funding for an empirical project as if you were simply critiquing a recent opinion. For this reason, it’s imperative that you get the most traction out of your research. Rather than taking a one study/one paper approach, consider publishing multiple pieces from a single study. Create subsets of questions that each explore a different theme. While a short, scholarly article in an online journal might discuss results from three to five questions, a longer, more traditional law review article might analyze findings from ten to fifteen completely different questions from the same study.

While pedagogical science generally refers to the *science* of learning, empirical research projects like the one I conducted this spring contribute to the growing body of knowledge on the *art* of teaching. Evaluating the efficacy of your pedagogical experiments will improve your teaching, contribute to the discipline, and demonstrate that legal writing professors are not just dedicated teachers but also curious scholars. Notably, the results of my study contradicted several of my hypotheses about student use of generative AI. For instance, it revealed that students used generative AI less than I expected and that their use of generative AI for legal analysis actually diminished somewhat as the semester progressed. Much to my surprise, most students also believed that generative AI use should not be permitted in first-year legal writing courses.

The findings from my study have forced me to rethink how I might integrate generative AI into this or other courses in the future and cemented my decision to not permit my students to use generative AI in my first-year legal writing course. But absent the study, none of this valuable knowledge would have been captured. In conclusion, you never know what incredible knowledge you'll discover when your curiosity takes the lead.