2019

Citation Stickiness

Kevin Bennardo
University of North Carolina School of Law, bennardo@unc.edu

Alexa Z. Chew
University of North Carolina School of Law, achew@email.unc.edu

Publication: *Journal of Appellate Practice & Process*

THE JOURNAL OF APPELLATE PRACTICE AND PROCESS

ARTICLES

CITATION STICKINESS

Kevin Bennardo* and Alexa Z. Chew**

ABSTRACT

This Article is an empirical study of what we call citation stickiness. A citation is sticky if it appears in one of the parties' briefs and then again in the court's opinion. Imagine that the parties use their briefs to toss citations in the court's direction. Some of those citations stick and appear in the opinion—these are the sticky citations. Some of those citations don't stick and go unmentioned by the court—these are the

*Clinical Associate Professor of Law at the University of North Carolina School of Law and Non-Resident Associate Justice of the Supreme Court of the Republic of Palau.

**Clinical Associate Professor of Law at the University of North Carolina School of Law.

Collectively, the authors wish to thank the many folks who helped them with this project. For their excellent research assistance, thank you to Taylor Carrere (UNC Law 2019), Keith Hartley (UNC Law 2018), Alexis Intriago (UNC Law 2020), and Tara Summerville (UNC Law 2019). For their assistance with study design and data interpretation, thank you to Felix S. Chew, Aaron Kirschenfeld, Guanya Liu, Annemarie Relyea-Chew, and Martin T. Wells. For their helpful comments, thank you to Ted Becker, Luke Everett, Joe Fore, Melissa Jacoby, Brian Larson, Lance Long, Pete Nemerovski, O.J. Salinas, Shaun Spencer, David Ziff, and our audiences at the Legal Writing Institute's Biennial Conference hosted by Marquette University Law School, the Legal Writing Institute's One-Day Workshop hosted by the Northeastern University School of Law, and the Festival of Legal Learning hosted by the University of North Carolina School of Law.

unsticky ones. Finally, some sources were never mentioned by the parties yet appear in the court's opinion. These authorities are endogenous—they spring from the court itself.

In a perfect adversarial world, the percentage of sticky citations in courts' opinions would be something approaching 100%. The parties would discuss the relevant authorities in their briefs, and the court would rely on the same authorities in its decisionmaking. Spoiler alert: our adversarial world is imperfect. Endogenous citations abound in judicial opinions and parties' briefs are brimming with unsticky citations.

So we crunched the numbers. We analyzed 325 cases decided by the federal courts of appeals. Of the 7552 cited cases in those opinions, more than half were never mentioned in the parties' briefs. But there's more—in the article, you'll learn how many of the 23,479 cited cases in the parties' briefs were sticky and how many were unsticky. You'll see the stickiness data sliced and diced in numerous ways: by circuit, by case topic, by an assortment of characteristics of the authoring judge. Read on!

TABLE OF CONTENTS

This article is an empirical study of what we call citation stickiness. A citation is sticky if it appears in one of the parties' briefs and then again in the court's opinion. If it helps, picture the parties tossing citations in the court's direction. Some of those citations stick and some of them don't. The ones that don't stick—that don't appear in the court's opinion—are unsticky. That covers the citations in the parties' briefs—they are either sticky or unsticky. As for the citations in a court's opinion, they are either sticky—meaning that the same source was cited in at least one brief—or they are endogenous—meaning that they appeared for the first time in the opinion. Endogenous citations spring from the court itself.

Consider a recent Tenth Circuit opinion in which the court cited thirty-three distinct cases.[1] Thirty-one of those cases were not mentioned in any of the parties' briefs. The opening brief cited twenty-nine cases,[2] the response brief cited eighteen,[3] and the reply brief cited five.[4] Out of all of the cases cited by the parties, however, the Tenth Circuit cited one from the opening brief, one from the response brief, and thirty-one that were not mentioned in any brief. On the other end of the spectrum, consider a recent opinion from the Seventh Circuit.[5] The court's opinion contains eleven unique case citations. Every single one of those eleven cases had been cited in one or more of the parties' briefs.[6]

---

1. Garling v. U.S. EPA, 849 F.3d 1289 (10th Cir. 2017).

2. Corrected Brief of Petitioners, *Garling v. U.S. EPA*, 849 F.3d 1289 (10th Cir. 2017) (No. 16-8028), 2016 WL 4010431, at ii–iii.

3. Brief of Appellee, *Garling v. U.S. EPA*, 849 F.3d 1289 (10th Cir. 2017) (No. 16-8028), 2016 WL 5899579, at ii–iii.

4. Reply Brief of Petitioners, *Garling v. U.S. EPA*, 849 F.3d 1289 (10th Cir. 2017) (No. 16-8028), 2016 WL 6247387, at ii. To be clear, there weren't fifty-two distinct cases cited in the three briefs. Some cases were cited in more than one brief (for example, a case cited in both the opening and response briefs).

5. Geiger v. Aetna Life Ins. Co., 845 F.3d 357 (7th Cir. 2017).

6. In *Geiger*, the Westlaw-generated tables of authorities reported that the opening brief cited thirty-six cases, the response brief cited sixty cases, and the reply brief cited thirteen cases. See the tables of authorities tabs associated to the following documents: Brief for Plaintiff-Appellant, 2016 WL 4254468, Brief of Defendant-Appellee, 2016 WL 5369221, and Reply Brief for Plaintiff-Appellant, 2016 WL 5461992. Our method for determining the number of cases cited in briefs and opinions is described in Part II of this article, *infra*.

For the Tenth Circuit case, we'd say that the cited cases in the opinion were 6% sticky and 94% endogenous.[7] We'd say that the cited cases in the opening brief were about 3% sticky and 97% unsticky.[8] For the Seventh Circuit case, the opinion contained 100% sticky case citations.[9] None of the cases cited in the opinion were endogenous.[10]

From those numbers alone, we cannot tell you whether the briefs in the Tenth Circuit case are better or worse than the briefs in the Seventh Circuit case. We cannot tell you whether the Tenth Circuit's decision is better or worse than the Seventh Circuit's decision. But what we can tell you is that there is wide variation in the percentage of cited cases in judicial opinions that originated in the parties' briefs.[11] The other thing we can tell you is that parties cite a lot of cases in their briefs that are never discussed in the resulting judicial opinions. Maybe you had some general sense of this already. Maybe not. But we bet you don't know the numbers. We do.

We traced the provenance of 7552 cited cases in 325 federal appellate opinions. We know how many of those case citations were borrowed from the parties' briefs and how many came from within the court itself. Read on to learn specifics, but here's a spoiler: most of the cases cited in the federal appellate opinions that we studied were not cited in either of the parties' briefs.

This result surprised us, as it surprised the participants in our online and conference-audience polls.[12] This finding is novel, as most citation-practice studies of judicial opinions do not trace the origins of the cited authorities. Studies that have tracked the communication of citations from the briefs to the resulting opinions have used smaller sample sizes and have been limited to only a few courts.

---

7. Two out of thirty-three is 6.06%.

8. One out of twenty-nine is 3.45%.

9. Eleven out of eleven is 100%.

10. Of the cases cited in that appeal's opening brief, about 17% were sticky and 83% were unsticky. (Six of the 36 cases from the opening brief were later cited in the *Geiger* opinion. Six out of 36 is 16.67%.)

11. Using our terminology, we'd say there is wide variation in the percentage of case citations in judicial opinions that are sticky and endogenous.

12. *See infra* note 110.

Additionally, we tracked the journeys of the 23,479 cases that the parties deemed worthy of citing in their briefs. We know how many of those cited cases later showed up in the resulting judicial opinions and how many did not.

Are there some characteristics that correlate with increased or decreased citation stickiness? Do some circuits tend to produce opinions with higher citation stickiness than others? Do some types of cases tend to result in higher citation stickiness? Do some judicial characteristics tend to correlate with higher citation stickiness? Fear not—we have sliced and diced the data in numerous ways.

And what does it all mean? That is a little less clear. Is it a problem that most of the cases cited in the opinions weren't mentioned by the parties? Yes, we'd say that these results indicate that something is amiss in our adversarial system. But where to point the finger? At attorneys for submitting shoddy briefs? At courts for disregarding the papers filed by the parties? Don't worry, dear reader, we devote an entire section to hypothesizing.

What follows are the results of our empirical study of citation stickiness in the federal courts of appeals. In Part I, we review the existing literature on citation studies and explain why studying citation stickiness is a worthwhile endeavor. In Part II, we lay out our research methodology. Part III reports our results, analyzing stickiness by various dimensions, such as case topic and certain characteristics of the authoring judges. Finally, Part IV hypothesizes what it all may mean and identifies some additional avenues for future research.

I. WHY [READ A] STUDY [ABOUT] CITATION STICKINESS?

So why should we study citation stickiness? Or, more saliently at this point, why should you read our study about citation stickiness? Most critically, our study is novel. It fills a heretofore unfilled gap. While filling a gap may be a necessary reason to undertake a study, it is not itself a sufficient one. There are plenty of things that haven't been studied simply because

they are not worth studying.[13] Aside from its novelty, citation stickiness is worth studying because it provides a window into judicial decisionmaking. Judges often lament the quality of attorneys' briefs. Attorneys often lament the quality of judges' decisions, especially when the opinions explaining those decisions veer away from the issues set forth in the briefs. Measuring citation stickiness will help uncover to what extent judges are conducting independent legal research. Answering that question seems foundational to determining whether judges are doing too much research, too little, or just the right amount.

This Part will proceed with a summary of the citation studies to date and identify the precise gap that our study fills. It will then discuss the utility of studying citation stickiness.

A. Review of the Citation-Study Literature

While it would be nice to say that we were experts on the citation-study literature before this project began, that would not be the whole truth.[14] We had not read every citation study out there and thoughtfully noticed a gap in the citation stickiness department. Rather, as many researchers do, we started with the question and discovered the gap. We thought citation stickiness was interesting, but when we researched it, we found little data that answered the question of whether courts generally stick to the legal authorities cited by the parties.

To be sure, there are plenty of citation-practice studies out there.[15] Many answer quite interesting questions. Given the laborious nature of the research, however, many of the studies are quite limited in scope. Many have small datasets or focus on

13. And there are a few things that have been studied even though they weren't worth studying.

14. And, yes, that is even the case for the one of us who previously authored an empirical citation study. *See* Kevin Bennardo, *Testing the Geographical Proximity Hypothesis: An Empirical Study of Citations to Nonbinding Precedents by Indiana Appellate Courts*, 90 NOTRE DAME L. REV. ONLINE 125 (2015).

15. For example, one article, which is itself more than a decade old, cites over fifty previous citation studies as background. Dietrich Fausten, Ingrid Nielsen & Russell Smyth, *A Century of Citation Practice on the Supreme Court of Victoria*, 31 MELB. U. L. REV. 733, 735–36 nn.12–20 (2007). We won't do the same here—but surely there are plenty of citation-practice studies to be found.

a particular year or narrow band of years.[16] Many are limited to studying the courts of a particular state or territory.[17] On the federal side, these citation studies disproportionately focus on the United States Supreme Court.[18] Despite the quantity of existing citation studies, there have been numerous calls for expansion of this method of research.[19]

Citation studies have largely focused exclusively on courts' opinions and ignored citation provenance. Studies have analyzed whether courts cite the same scholarship that academics do[20] and

16. *Id.* at 736 ("Because of the financial cost of collecting large datasets, most studies have focused on citation practice within a single year or a few select years.").

17. *E.g.*, A. Michael Beaird, *Citations to Authority by the Arkansas Appellate Courts, 1950–2000*, 25 U. ARK. LITTLE ROCK L. REV. 301 (2003); William H. Manz, *The Citation Practices of the New York Court of Appeals: A Millennium Update*, 49 BUFF. L. REV. 1273 (2001); Joseph A. Custer, *Citation Practices of the Kansas Supreme Court and Kansas Court of Appeals*, 7 KAN. J.L. & PUB. POL'Y 120 (1998); Fritz Snyder, *The Citation Practices of the Montana Supreme Court*, 57 MONT. L. REV. 453 (1996); James Leonard, *An Analysis of Citations to Authority in Ohio Appellate Decisions Published in 1990*, 86 LAW LIBR. J. 129 (1994); Richard A. Mann, *The North Carolina Supreme Court 1977: A Statistical Analysis*, 15 WAKE FOREST L. REV. 39 (1979); John Henry Merryman, *Toward a Theory of Citations: An Empirical Study of the Citation Practice of the California Supreme Court in 1950, 1960, and 1970*, 50 S. CAL. L. REV. 381 (1977).

18. *E.g.*, Ryan C. Black & James F. Spriggs II, *The Citation and Depreciation of U.S. Supreme Court Precedent*, 10 J. EMPIRICAL LEGAL STUD. 325 (2013); Raizel Liebler & June Liebert, *Something Rotten in the State of Legal Citation: The Life Span of a United States Supreme Court Citation Containing an Internet Link (1996–2010)*, 15 YALE J.L. & TECH. 273 (2013); Frank B. Cross, James F. Spriggs II, Timothy R. Johnson & Paul J. Wahlbeck, *Citations in the U.S. Supreme Court: An Empirical Study of Their Use and Significance*, 2010 U. ILL. L. REV. 489; Jules Gleicher, *The Bard at the Bar: Some Citations of Shakespeare by the United States Supreme Court*, 26 OKLA. CITY U.L. REV. 327 (2001).

19. *See, e.g.*, Frank B. Cross, *Chief Justice Roberts and Precedent: A Preliminary Study*, 86 N.C. L. Rev. 1251, 1277–78 (2008); Richard A. Posner, *An Economic Analysis of the Use of Citations in the Law*, 2 AM. L. & ECON. REV. 381, 381–83 (2000).

20. *See* Deborah J. Merritt & Melanie Putnam, *Judges and Scholars: Do Courts and Scholarly Journals Cite the Same Law Review Articles?* 71 CHI-KENT L. REV. 871 (1996). Academics tend to pay a disproportionate amount of attention to citations to legal scholarship relative to other types of authorities. *See, e.g.*, Derek Simpson & Lee Petherbridge, *An Empirical Study of the Use of Legal Scholarship in Supreme Court Trademark Jurisprudence*, 35 CARDOZO L. REV. 931 (2014); Brent E. Newton, *Law Review Scholarship in the Eyes of the Twenty-First-Century Supreme Court Justices: An Empirical Analysis*, 4 DREXEL L. REV. 399 (2012); David L. Schwartz & Lee Petherbridge, *The Use of Legal Scholarship by the Federal Courts of Appeals: An Empirical Study*, 96 CORNELL L. REV. 1345 (2011); Lou J. Sirico, Jr., *The Citing of Law Reviews by the Supreme Court: 1971–1999*, 75 IND. L.J. 1009 (2000); Vaughan Black & Nicholas Richter, *Did She Mention My Name?: Citation of Academic Authority by the Supreme Court of Canada, 1985–1990*, 16 DALHOUSIE L.J. 377 (1993).

how precedents are transmitted from court to court.[21] But few have compared the sources cited in the parties' briefs to the sources cited in the resulting opinions, perhaps because this is a question that is likely more interesting to practitioners than to academics.[22] With a notable exception or two, the studies that have previously compared briefs' citations to opinions' citations have been extremely limited in scope.[23] The existing studies are summarized below.

1. Study of Citation Stickiness in the State Courts

The most robust citation stickiness study to date is Thomas Marvell's state-court study from over forty years ago. The court, however, was anonymous. All we know is that it was the "supreme court of a northern industrial state," referred to by Dr. Marvell as the "focal court" of his study.[24] The Marvell dataset comprised 112 cases argued during a one-year period ending in June 1972.[25] Comparing the attorneys' submissions to the opinions, Dr. Marvell found that "[a] little less than half the legal authorities cited in the majority and minority opinions in the 112 focal cases studied here were mentioned in the parties' briefs or oral arguments, and but one-sixth of the authorities

21. *E.g.*, Iain Carmichael, James Wudel, Michael Kim & James Jushchuk, Comment, *Examining the Evolution of Legal Precedent through Citation Network Analysis*, 96 N.C. L. REV. 227 (2017); Bennardo, *supra* note 14; Russell Smyth & Vinod Mishra, *The Transmission of Legal Precedent Across the Australian State Supreme Courts over the Twentieth Century*, 45 L. & SOC'Y REV. 139 (2011).

22. *See* Cross, *supra* note 19, at 1272 (heralding briefs as "a heretofore underutilized tool of empirical research"). Even some studies that have included data regarding the sources cited in the parties' briefs and the resulting opinions have failed to compare the two datasets to determine the extent to which the parties' citations influence the court's citations. *See, e.g.*, Jack Knight & Lee Epstein, *The Norm of* Stare Decisis, 40 AM. J. POL. SCI. 1018, 1025 tbl.2, 1030 tbl.4 (1996). Other studies have compared the language used in briefs and opinions, but specifically excluded the citation from comparison. *See, e.g.*, Pamela C. Corley, *The Supreme Court and Opinion Content: The Influence of Parties' Briefs*, 61 POL. RES. Q. 468, 471 (2008) (noting methodology of skipping citations in a study using plagiarism software to compare language in briefs and opinions).

23. This statement isn't meant to be critical. Most of these studies were primarily studying other things.

24. THOMAS B. MARVELL, APPELLATE COURTS AND LAWYERS: INFORMATION GATHERING IN THE ADVERSARY SYSTEM 6 (1978).

25. *Id.*

mentioned by the attorneys were cited in the opinions."[26] As a sort of control study, Dr. Marvell also compared the citations in the published opinions and briefs of thirty civil cases from the United States Court of Appeals for the Sixth Circuit.[27] His results were fairly consistent with the focal court study: 55% of the authorities cited by the Sixth Circuit in those cases were first mentioned in the briefs.[28]

Combining his data with interviews with judges and court staff, Dr. Marvell concluded that courts generally do a lot of independent legal research, although the exact amount seemed to vary quite a bit from chambers to chambers.[29] He noted that several of the focal court justices' law clerks "said that they used the briefs hardly at all or only as a place to begin the research when writing draft opinions or memorandums. The law clerks or, increasingly, the staff attorneys do the great bulk of the research."[30]

The other existing citation-stickiness study of a state court focused on the citation of New Jersey cases in the appellate courts of New Jersey. In an effort to bring attention to the deficiencies of the state bar, the Chief Justice of the New Jersey Supreme Court asked the sergeants-at-arms of the state supreme court and the appellate division "to go through our opinions and the briefs on every appeal that has been decided over the past year and to note the New Jersey decisions in the opinions that are not cited in the briefs."[31] Unfortunately, the Chief Justice did not report the data in a way that would allow us to calculate the actual stickiness rate. Instead, he highlighted the fact that 82% of the opinions of the state supreme court and 41% of the opinions of the intermediate appellate court cited to at least one New Jersey case that was not mentioned in the briefs. In the

26. *Id.* at 132 (endnote omitted). Oral argument did not add much in the way of new authorities—Dr. Marvell found that only 1% of the parties' authorities were mentioned at oral argument and not in briefs. *Id.* at 133. We note here that Dr. Marvell served up the data several different ways, *see, e.g.*, *id.* at 132–36, and we recommend reading his study in full.

27. *Id.* at 134–36, 135 n.18.

28. *Id.*at 134–35.

29. *Id.* at 135–36.

30. *Id.* at 135.

31. Arthur T. Vanderbilt, *Our New Judicial Establishment: The Record of the First Year*, 4 RUTGERS L. REV. 353, 361 (1950).

Chief Justice's view, this was enough to show "how deficient a large portion of the briefs filed in our appellate courts are in point of law and what a burden of independent research they impose on the judges."[32]

2. Study of Citation Stickiness in the Federal Courts

Aside from the Marvell study mentioned in the last subsection, the next most sizable study of citation stickiness to date was William Manz's study of the Supreme Court. Mr. Manz compared the decisional authorities cited in the briefs to those cited in the Supreme Court's eighty majority opinions during its 1996 term.[33] On the issue of citation stickiness, he found that 74.5% of the decisional authorities cited in the Court's opinions were also cited in one or more of the briefs.[34] Mr. Manz thus surmised that "roughly one-quarter of the Court's case citations resulted from its own research."[35] Relatively few of the cases from the briefs were later cited by the majority opinion—only about 25%.[36]

In a self-described "brief" study, Professor Cross assessed Chief Justice Roberts's approach to precedent by examining the opinions that the Chief Justice authored in his first term.[37] As part of the examination, Professor Cross looked at how often Chief Justice Roberts cited the same cases that had been cited in the parties' briefs.[38] The sample size was small: only nine

32. *Id.*

33. William H. Manz, *Citations in Supreme Court Opinions and Briefs: A Comparative Study*, 94 L. LIBR. J. 267, 267–68 (2002). The Manz study included citations to judicial opinions and administrative decisions, but excluded citations to constitutions, statutes, and regulations. *Id.* at 268.

34. *Id.* at 271 tbl.5 (reporting that, of the 1915 authorities cited in the Court's opinions, 1427 were first cited in a brief). Of the 1915 citations in the Court's opinions, 146 (or 7.6%) appeared first in an amicus brief and not in any of the parties' briefs. *Id.* at 272 tbl.7.

35. *Id.* at 271.

36. *See id.* at 271, 272 tbl.6. The Manz study includes a number of other worthwhile data points, including citation-stickiness data for numerous types of secondary authorities, and we commend it to you in full.

37. Cross, *supra* note 19, at 1251. All the cases were published in 2005 and 2006. *See id.* at 1274 nn.146–54.

38. *Id.* at 1273.

cases.[39] In the briefs for the nine cases, he found that 168 cases had been cited by both parties.[40] Professor Cross also noted that "[o]ne might think that if both parties relied on the case, it would be an unavoidable citation for the Court's opinion."[41] That did not turn out to be so. Of the 168 cases cited by both parties, only seventy-eight of them (a little over 46%) appeared in Chief Justice Roberts's opinions.[42] By aggregating some of the data reported in the Cross study, it appears that the nine Roberts opinions cited a total of 305 cases,[43] of which 124 (over 40%) were not mentioned in either party's briefs. Professor Cross concludes that "[i]t seems plain that Justice Roberts exercised considerable discretion in choosing which precedents to cite."[44]

While much narrower in scope than the Manz study, Professor Cross's observations from Chief Justice Roberts's opinions demonstrate a much lower rate of citation stickiness than did Manz's study of the broader court. Other researchers have attempted to measure the stickiness of specific types of authorities at the Supreme Court, including its citations to legal periodicals,[45] its rate of "interpretation" of legal authorities,[46]

---

39. *Id.* at 1274 tbl.3 ("Opinion Citations Compared with Briefs").

40. *Id.*

41. *Id.* at 1273.

42. *Id.* at 1274 tbl.3.

43. We arrived at this number by summing the middle four columns of Table 3 in the Cross article. *See id.* This result is at odds, however, with Professor Cross's earlier claim that Chief Justice Roberts cited an average of twenty-seven cases per opinion, *id.* at 1268, which would result in the citation of only 243 cases in nine opinions. In any event, Chief Justice Roberts got a substantial proportion of his cited cases from somewhere other than the parties' briefs.

44. *Id.* at 1274.

45. Professor Newland's study focused on the citation of legal periodicals by individual Supreme Court Justices from 1924 through 1956. Chester A. Newland, *Legal Periodicals and the United States Supreme Court*, 7 U. KAN. L. REV. 477, 477 (1959). In one portion of his study, he identified the thirteen Justices who most frequently cited legal periodicals. *Id.* at 480 tbl.3 ("*Totals of Articles Cited by 13 Justices Who Have Most Frequently Cited Legal Periodicals*"). Although Professor Newland himself did not total the data, the upshot is that of the 1453 articles cited by the Justices, only 262 of them had appeared in the briefs (approximately 18%). In majority opinions, the percentage of articles cited in the opinion that had appeared in the briefs was a little over 20% (199 of 958). In concurring opinions, the percentage was highest at a little over 23% (22 of 94). And for dissenting opinions, the percentage was lowest at only a little over 10% of the cited articles coming from the briefs (41 of 401). Professor Newland did not draw many conclusions from this data, but did note that Justice Brandeis's citations of many articles not mentioned in the briefs "reflects

and its use of legislative history.[47] Because the Supreme Court is a judicial body that is uniquely not bound by the traditional

Brandeis's well-known practice of completing considerable original research in preparation of his opinions." *Id.* at 480. The Newland study shows that Justice Brandeis cited 127 articles, only five of which appeared in the briefs. *Id.* at 480 tbl.3.

46. Professors Spriggs and Hansford studied how the Supreme Court "chose to legally interpret the set of available Supreme Court precedents" in the 1991 and 1995 terms. James F. Spriggs II & Thomas G. Hansford, *The U.S. Supreme Court's Incorporation and Interpretation of Precedent*, 36 L. & SOC'Y REV. 139, 139 (2002). Importantly, the authors focused on which precedents the Court elected to "interpret," not on the Court's mere citation of precedents. *Id.* at 146. The authors noted that the Court may deal with a precedent in "three basic ways": positively interpret it (by relying on it), negatively interpret it (by distinguishing, limiting, or overruling it), or not legally interpret it. *Id.* at 141. To identify the world of precedents that the Court could potentially interpret, the authors "assumed that the available set of precedents in a case consisted of the Supreme Court cases referred to in its briefs." *Id.* at 145. The authors noted all of the Supreme Court cases that the parties had cited in their briefs and found that there were approximately sixty potential "precedents" that the Court could interpret in each case. *Id.* Thus, the authors' focus was not comparing the cases cited in the parties' briefs to the cases cited in the Court's opinions; rather, their focus was comparing the cases cited in the parties' briefs to the cases actually "interpreted" in the Court's opinions. The authors relied on *Shepard's Citations* to determine whether a cited case was actually "interpreted" by the Court. *Id.* at 146.

The results were quite low: the Court "interpreted" only 2.3% of the cases cited in the parties' briefs. *Id.* at 149–50 (reporting that the Court interpreted 250 out of 10,842 possible precedents). The authors opined that this result may be the product of attorneys' adopting a "scattershot" approach of citing many precedents in their briefs, "many of which are not particularly relevant to the case at hand." *Id.* at 150 n.16. Stripping out the irrelevant precedents, the authors found that the Court "interpreted" a little over 15% of the legally relevant precedents. *Id.* at 151. One of the more interesting (to us) tidbits was relegated to a footnote: the Court analyzed twenty-six precedents that were absent from the parties' briefs, meaning that something like 10% of the cases that the Court "interpreted" were not even mentioned by the parties. *Id.* at 145 n.4.

The authors later updated the study using a revised methodology. *See* THOMAS G. HANSFORD & JAMES F. SPRIGGS II, THE POLITICS OF PRECEDENT ON THE U.S. SUPREME COURT 93–108 (2006). In the revised study, the authors conceded that their earlier approach of defining the world of potential precedent as only the authorities cited in the parties' briefs was "underinclusive," as evidenced by the fact "that such a research design misses approximately 10% of all cases actually interpreted by the Court." *Id.* at 95–96. In the revised study, the authors defined the world of potential precedent much more expansively: as all of the cases orally argued at the Supreme Court since 1946. *Id.* at 96.

47. Professor Parrillo examined the use of legislative history as a tool of statutory interpretation in judicial decisions. Nicholas R. Parrillo, *Leviathan and Interpretive Revolution: The Administrative State, the Judiciary, and the Rise of Legislative History, 1890–1950*, 123 YALE L.J. 266 (2013). Specifically, he sought to track the path by which legislative history went from being a permissible tool of statutory interpretation to a "normal, routine, and expected" one. *Id.* at 274. Using statutory-interpretation cases decided by the Supreme Court from 1940 to 1945, he compared citations to legislative history in the briefs to those in the resulting opinions. *Id.* at 281. Although Professor Parrillo's focus was demonstrating the quantity of citations to legislative history that the

system of precedents, studying it provides a skewed perspective into judicial decisionmaking writ large.

In another notable study that touched on citation stickiness, a trio of researchers investigated the effectiveness of using automated content analysis as a research methodology in legal scholarship.[48] Although they were "focused primarily on validating [their] methodology, rather than on the results it generates," they noted that they managed to "generate[] intriguing results that suggest avenues for further study."[49] In other words, the citation-stickiness data was but a happy collateral byproduct of the actual focus of the study.

One area at which the trio aimed their automated content analysis was a citation study of First Circuit opinions.[50] As a way to measure "judicial responsiveness," the researchers "assessed the relationship between the briefs and the opinions in terms of authorities upon which both relied."[51] The researchers stated that "a court's resort to the same authorities as relied upon by the parties seems almost necessarily to be coextensive with a responsive analysis."[52] With a sample size of thirty First Circuit opinions, the authors found that only 35% of the authorities cited by the court were cited in either party's brief.[53] Conversely, only about 16% of the authorities cited in the

---

Court received in briefs filed by the federal government, *id.* at 281–82, he uncovered some data regarding citation stickiness in this slender area. In the cases involving briefs from the federal government, his research showed that

> 22% of the citations [to legislative history in the Court's opinions] matched both the federal brief and at least one non-federal brief (such as the brief of a private party or state government); 33% of the citations matched the federal brief and no other brief; 10% of the citations matched at least non-federal brief *but not* the federal brief; and 24% of the citations matched no brief (suggesting they arose from the Court's own research).

Id. at 317. In cases with no federal-government briefs, Professor Parrillo found that "45% of the citations to legislative history appeared in at least one of the briefs, while 55% did not (suggesting they came from the Court's own research)." *Id.* at 318. Parrillo then hypothesized about the factors that gave rise to the Court's ability to conduct so much internal research into legislative history. *Id.* at 361–64.

48. Chad M. Oldfather, Joseph P. Bockhorst & Brian P. Dimmer, *Triangulating Judicial Responsiveness: Automated Content Analysis, Judicial Opinions, and the Methodology of Legal Scholarship*, 64 FLA. L. REV. 1189, 1189 (2012).

49. *Id.*

50. *Id.* at 1232–39.

51. *Id.* at 1232.

52. *Id.* at 1234.

53. *Id.* at 1238.

parties' briefs were cited by the court.[54] Although the researchers were focused on methodology rather than results, they found the citation-analysis results "intriguing in their own right" and suggested further study.[55] Similar to Professor Cross's observation that Chief Justice Roberts did not confine himself to the sources cited in the parties' briefs, these researchers observed that "judges have, and exercise, a considerable amount of discretion in choosing which precedent to follow."[56]

B. The Utility of Studying Citation Stickiness

As chronicled above, previous comparison studies of citations in briefs and resulting opinions are spotty, scattered, and often aged. They are limited in scope. Some deal only with citations to authorities other than cases.[57] Many focus on the United States Supreme Court.[58] Most have very small sample sizes. Thus, there is a gap in the literature. But is it a gap worth filling?

We think it is. More accurately, we thought it was. And then we partially filled it with this article. Now we'll attempt to convince you that we spent our efforts wisely.

First, there is the incomplete, but interesting, story told by the data reported above. Although the previous citation stickiness studies are spotty and limited, there is one common thread: they consistently indicate that a substantial proportion of the authorities cited in courts' opinions were not cited in the parties' briefs. A larger and more comprehensive study was needed to validate those findings.

Second, data on judicial decisionmaking aids brief writers. There have been an increasing number of calls for more empirical research into judicial decisionmaking. Although there

---

54. *Id.*

55. *Id.* at 1238–39 (noting it was "striking how little overlap there is between the parties' citations and the court's").

56. *Id.* at 1239. Additionally, the Marvell state-court study contains a limited citation-stickiness analysis of Sixth Circuit opinions from the 1970s. *See supra* Section I.A.1.

57. *See* Newland, *supra* note 45; Parrillo, *supra* note 47.

58. *See, e.g.*, Newland, *supra* note 45; HANSFORD & SPRIGGS, *supra* note 46; Spriggs & Hansford, *supra* note 46; Cross, *supra* note 19; Parrillo, *supra* note 47.

is much conventional wisdom (some of which conflicts with other conventional wisdom), there is relatively little evidence on what affects judicial decisionmaking.[59] For example, there is a foundational perception that the parties' briefs are important, but we don't actually know to what extent that is true, nor do we have a good sense of what makes some briefs more persuasive than others.[60] Identifying the factors that increase judicial responsiveness can help attorneys write briefs that are more likely to prompt relevant discussion by the court.[61]

Third, data on citation stickiness can help shape debates over the process of resolving disputes in our judicial system.[62] On the one hand, some judges have vocally expressed a belief that attorneys' briefs are largely deficient and generally unhelpful.[63] They complain that attorneys do not write well and, worse, fail to discuss the controlling precedents.[64] On the other hand, some attorneys and commentators have decried shortcomings in the quality of judicial decisions, and particularly have complained about judicial attempts at decisionmaking without the benefit of the parties' input.[65] The most robust literature debates to what extent judges may or

---

59. *See, e.g.*, Ted Becker, *What We Still Don't Know About What Persuades Judges—And Some Ways We Might Find Out*, 22 J. LEGAL WRITING INST. 41, 41–43 (2018).

60. Attempts at empirically assessing the effect of briefs have increased in recent years. *See, e.g.*, Shaun B. Spencer & Adam Feldman, *Words Count: The Empirical Relationship Between Brief Writing and Summary Judgment Success*, 22 J. LEGAL WRITING INST. 61 (2018); Adam Feldman, *A Brief Assessment of Supreme Court Opinion Language, 1946–2013*, 86 MISS. L.J. 105 (2017); Adam Feldman, *Counting on Quality: The Effects of Merits Brief Quality on Supreme Court Decisions*, 94 DENV. L. REV. 43 (2016).

61. *See* Oldfather et. al, *supra* note 48, at 1218–19 ("One could imagine, for example, large-scale analysis of the relationships among briefs and opinions generating information about the relative utility of briefing practices and approaches.").

62. *See id.* at 1217–18.

63. *See, e.g.*, Vanderbilt, *supra* note 31, at 361 ("[F]our out of five of all the briefs submitted to us are of inferior quality."); Stephen L. Wasby, *As Seen From Behind the Bench: Judges' Commentary on Lawyers' Competence*, 38 J. LEGAL PROF. 47, 61–68 (2013) (reporting judges' negative reactions to briefs).

64. *See, e.g.*, Vanderbilt, *supra* note 31, at 361; Wasby, *supra* note 63, at 61–62 (recounting judicial reaction to the omission of a key case from the briefing).

65. *See, e.g.*, Michael Abramowicz & Thomas B. Colby, *Notice-and-Comment Judicial Decisionmaking*, 76 U. CHI. L. REV. 965, 972 (2009) ("Other times, the court will resolve the case by employing legal reasoning and citing legal authorities not suggested by the parties—which means that the parties were never able to challenge or criticize the legal reasoning that drove the court's decision. . . . This can lead to mistakes that the parties might have caught if given a chance." (footnote omitted)).

should engage in independent factual research outside of the record.[66] But some commentators are equally critical of sua sponte rulings on issues not briefed by the parties, and of courts basing their decisions on precedents that the parties did not brief.[67] If judges are restricted from independently researching facts, should they be similarly restricted from independently researching law? If not, why not?[68]

In short, judges are skeptical of attorneys' ability to be helpful, but attorneys are equally skeptical of judges' ability to make sound decisions without their help. Trust is lacking on both sides. This article won't resolve this issue. Instead, its contribution is data. This article will tell you what proportion of

---

66. The Model Code of Judicial Conduct prohibits independent judicial research of "facts." MODEL CODE OF JUDICIAL CONDUCT R. 2.9(C) (AM. BAR ASS'N 2010); *see also* ABA Comm'n on Ethics & Prof'l Responsibility, Formal Op. 478 (2017); CHARLES GARDNER GEYH ET AL., JUDICIAL CONDUCT AND ETHICS § 5.04, 5-24 (5th ed. 2013) ("Independent factual investigation impairs the function of an adversarial system by allowing a judge to craft decisions on the basis of facts that may be unknown to one or both of the parties and therefore indisputable by them regardless of their accuracy or relevance." (footnote omitted)). Of course, that leaves open to debate what is "fact" and what is not. *See, e.g.*, Allison Orr Larsen, *Confronting Supreme Court Fact Finding*, 98 VA. L. REV. 1255, 1264–71 (2012); Elizabeth G. Thornburg, *The Curious Appellate Judge: Ethical Limits on Independent Research*, 28 REV. LITIG. 131, 148–57 (2008); David L. Faigman, *Normative Constitutional Fact-Finding: Exploring the Empirical Component of Constitutional Interpretation*, 139 U. PA. L. REV. 541, 551–56 (1991); Laurens Walker & John Monahan, *Social Facts: Scientific Methodology as Legal Precedent*, 76 CAL. L. REV. 877, 880–82 (1988); Peggy C. Davis, *"There is a Book Out. . .": An Analysis of Judicial Absorption of Legislative Facts*, 100 HARV. L. REV. 1539, 1539–42 (1987); John Monahan & Laurens Walker, *Social Authority: Obtaining, Evaluating, and Establishing Social Science in Law*, 134 U. PA. L. REV. 477, 479–95 (1986).

67. *See* Michael J. Donaldson, *Justice in Full is Time Well Spent: Why the Supreme Court Should Ban Sua Sponte Dismissals*, 36 QUINNIPIAC L. REV. 25, 43 (2017) ("When a court decides sua sponte, it is deciding without input of the people who know the most about the case—the parties and their counsel. This increases the likelihood that the court will miss some relevant statute, precedent, fact, or argument in making its decision." (footnotes omitted)); Adam A. Milani & Michael R. Smith, *Playing God: A Critical Look at Sua Sponte Decisions by Appellate Courts*, 69 TENN. L. REV. 245, 313 (2002) (arguing that decisions rendered without input from the parties should carry less precedential weight than dicta); *but see* Amanda Frost, *The Limits of Advocacy*, 59 DUKE L.J. 447 (2009) (defending the practice of federal judges injecting new legal issues into the cases before them).

68. One answer may be that legal precedent is passed from case to case, but factual determinations are not. *See* Frost, *supra* note 67, at 493 ("Just as it is important for courts to respect stare decisis, it is essential that litigants not be allowed to slip its bonds simply by refusing to cite established precedent."); *see also* GEYH ET AL., *supra* note 66, at § 5.04, 5-25 ("Whereas judges are not presumptively experts on questions of fact, they are experts on matters of law who are charged with the duty of declaring what the law is.").

cited cases the federal courts of appeals are getting from the parties' briefs. We will not tell you whether that number is too few, too many, or just right. But future debate should be grounded in data, not in anecdote and perception.

II. OUR METHODOLOGY

Now let's talk about how we did it. For our dataset, we targeted recent published opinions of the federal courts of appeals. We selected the federal courts of appeals for a few reasons. First, we wanted to focus on federal rather than state courts. We figured that briefs filed in federal courts would be more easily accessible than many briefs filed in state courts, and, to be frank, we did not want to limit our audience by focusing on any particular state.

Second, we knew that we did not want to focus our study on the United States Supreme Court. It is unique and therefore unrepresentative of courts in general. It is also over-studied relative to other courts, particularly considering its small caseload. And because it is not bound by precedent, it makes a poor subject for a study that touches on the communication of precedent.

So that left federal district courts and federal courts of appeals. Because appellate cases follow a more consistent briefing lifecycle, it is simply easier to construct a consistent dataset out of appellate cases than trial cases. Appellate cases often progress along the same path: opening brief, response brief, and (maybe) reply brief. These briefs tend to be formal and contain tables of authorities. Trial cases can involve numerous types of motions, some of which are briefed and some of which are not, and perhaps even a trial. In the federal district courts, briefs tend to vary more in length, consistency, and formality. Also, many lack tables of authorities. Thus, in order to more easily construct a consistent dataset, we opted for the federal courts of appeals.

To create a broad sample and avoid focusing on a single circuit that might turn out to be an outlier, we sampled cases from each of the thirteen federal courts of appeals. We also thought it would be interesting to have some data from every circuit because we could then compare each circuit to the others.

Beyond that, we wanted to capture a random assortment of recent cases. Why recent? Because we wanted our data to be as current as possible.

We wanted to capture cases in which a full opinion resulted from adversarial briefing. Thus, we excluded unpublished opinions, per curiam opinions, and memorandum opinions. We figured that these opinions were more likely to be short, to cite few authorities, or to contain sections that were cut and pasted from a court's stockpile of generic language. Instead, we limited our dataset to only authored, published opinions.

As for briefs, we limited our dataset to cases in which the briefs were available on Westlaw. We also excluded cases in which there were supplemental briefs or amicus briefs. We sought to capture truly mine-run cases: those that progressed along the traditional pathway of opening brief, response brief, and (maybe) reply brief.[69]

We generated our dataset by identifying the first twenty-five cases from each circuit to meet our research criteria.[70] First we would verify that the briefing in the case met our criteria and that the briefs were available on Westlaw.[71] If the briefs met our

69. Cases with amended briefs were not disqualified because an amended brief is not an additional brief.

70. In this context, "we" means "us," not "our research assistants." We created separate spreadsheets for each of the thirteen circuits, which we called "circuit spreadsheets." Here are the steps to creating a circuit spreadsheet. First, from the Westlaw main screen at https://www.westlaw.com/, select "Cases" under the "All Content" tab. Next, under "Federal Cases by Circuit," select the desired circuit (e.g., "1st Circuit"). On the next page, select the court of appeals (e.g., "First Circuit Court of Appeals"). This leads to a database of that court's cases. Within that database, restrict the results to those starting in 2017 by entering the following search: "advanced: DA(aft 12-31-2016 & bef 01-01-2018)." Then, under "Reported Status," click the filter for only "Reported" cases. Then sort the results by date and go to the end of the search results to begin the screening process with the oldest opinion, which will be the opinion closest in time to January 1, 2017. (Note that these instructions were created during the latter half of 2017 and may no longer hold true after the conversion to Westlaw Edge.)

71. The first criterion used to disqualify cases was whether the briefs were available on Westlaw because we found through experience that this was the most likely piece to be missing. So, after accessing the court's opinion on Westlaw, we would click on the "Filings" tab to see which filings were available. Often the filings would lead to disqualification either because there were too few briefs available on Westlaw or because there were supplemental, amicus, or other additional briefs. If the briefing was very straightforward (e.g., Westlaw displayed only an opening, responsive, and reply brief) we ended our briefing investigation there. If the briefing was potentially within the bounds of our parameters but the filings available on Westlaw raised some suspicions (e.g., there was an opening and responsive brief but no reply brief or there were duplicate or amended

criteria, we would turn our attention to the court's opinion and verify that it was an authored, published opinion. For each circuit, we'd take the first twenty-five cases of 2017 that met those criteria.[72] Although not randomized from a larger dataset, this approach comports with the methodology used previously by others.[73] For some circuits it was relatively easy to find cases that met our criteria. For others, we had to assess hundreds of opinions to find twenty-five. Given the variation in the number of published cases issued by each court of appeals and the wide variation in the coverage of briefs available on Westlaw, each circuit's dataset has a unique span of dates.[74] For eight of the thirteen circuits, we were able to assemble our twenty-five cases from those issued by the end of March 2017. Thus, our dataset overwhelmingly comprises cases from the first half, and mostly the first quarter, of 2017.[75] In all, our dataset comprises 325 cases—far more than any previous study of citation stickiness.[76]

---

versions of briefs), then we investigated the court's docket using the "Dockets Search" feature of Bloomberg Law's "Litigation Intelligence Center." *See Litigation Intelligence Center*, BLOOMBERG LAW ((June 19, 2018), https://www.bloomberglaw.com. For example, we would check to verify that there really was no reply brief filed in the case. If it turned out that there was a reply brief, but Westlaw did not have a copy, then the case would be disqualified from our dataset. Additionally, in the cases that aroused our suspicions, we verified whether the parties filed additional briefs (and, if so, disqualified the case from our dataset).

72. In the Seventh Circuit, two cases that otherwise met our criteria were excluded because no cases were cited in the court's opinions. *See* Hart v. Amazon.com, Inc., 845 F.3d 802 (7th Cir. 2017); United States v. Gibbs, 845 F.3d 804 (7th Cir. 2017). While Judge Posner authored both opinions, one was accompanied by two concurring opinions that also cited no decisional authority. *See Gibbs*, 845 F.3d at 806 (Sykes & Kanne, JJ, concurring in the judgment, & Kanne, J, concurring). Because our primary research focus was determining whether the court's case citations come from the briefs or elsewhere, we decided that it was sensible to exclude opinions that cited no cases.

73. *See, e.g.*, MARVELL, *supra* note 24, at 337 n.18 (describing the methodology of selecting thirty Sixth Circuit cases to study).

74. The Eighth Circuit has the shortest range of dates—we found our twenty-five Eighth Circuit cases between January 1, 2017, and February 2, 2017. The Third Circuit took the largest range of dates for us to get our twenty-five cases: January 1, 2017, to July 25, 2017.

75. We considered staggering the start date of each circuit (e.g., starting the First Circuit with cases decided on January 1, starting the Second Circuit with cases decided on February 1, and so on). Ultimately, however, we failed to imagine any way in which starting every circuit on the same date would skew our results. Therefore, we opted to start every circuit on the same date for the sake of simplicity and consistency.

76. Previously, Dr. Marvell's study from the 1970s contains the most comprehensive comparison of citations in briefs and opinions. *See supra* Section I.A. The Marvell study contained a dataset of 112 cases. *See generally* notes 24–30, *supra*, and accompanying text.

Having assembled our dataset, we then turned to collecting the actual citation data so that we could compare the citations in each of the briefs with the citations in the resulting opinions. For this task, we primarily relied on one of our research assistants. For each case listed in each circuit spreadsheet, she used Westlaw to download the tables of authorities of each opinion and brief (opening, responsive, reply if any) in .docx format. At the time we conducted this study, Westlaw's table-of-authorities feature captured cases cited in the opinion and briefs but not statutes, regulations, or legislative history. Only very rarely did the tables of authorities capture any non-decisional authority.[77] Thus for each appeal listed in a circuit spreadsheet, we had three or four lists of case citations—one list for each brief or opinion. We then were able to compare the lists of case citations from each brief's table of authorities to the case citations in the other briefs' tables of authorities and to the cases cited in the resulting opinions.[78]

---

77. Other than judicial opinions, the materials that turned up in the Westlaw tables of authorities in our dataset were decisions of administrative agencies (classified as "Administrative Decision & Guidance"), patents (classified as "Intellectual Property"), and a lone American Law Report annotation (classified as "Secondary Sources"). We manually deleted the patent and ALR citations from our dataset, but we kept the citations to administrative adjudications because they were decisional in nature. For examples, click on Westlaw's Table of Authorities tab for the briefs associated with the following database identification numbers: 2016 WL 1466312, 2014 WL 5421879, and 2016 WL 7435951.

78. Although we summed the process in a single tidy sentence, the process itself was not so streamlined. Here's what our team did:

To mechanize the task of comparing the lists of cases, our research assistant used a markup program to convert each .docx document into a plain text document. This process stripped out the Westlaw formatting (including KeyCite flags), leaving each plain-text citation on the list in its own paragraph.

Next, we used Excel spreadsheets to compare the plain-text lists of cited cases to one another. We wrote the necessary formulas in the spreadsheets, which compared the lists of citations to one another. Each appeal had its own spreadsheet that included five tabs: a tab for each of the three briefs, one for the opinion, and a summary sheet tabulating citation counts from the first three. Our research assistant pasted the lists of citations into the spreadsheets. For example, she pasted the list of case citations from the opening brief into column A of the first tab, then the list of case citations from the responsive brief into column A of the second tab, and so on.

Because Westlaw's table-of-authorities feature formatted each case citation identically, Excel's comparing function could search for matches across the first four tabs in each spreadsheet. For example, the formulas in column B of the opening-brief sheet compared each case citation in column A with all of the case citations listed in the responsive brief, reply brief, and opinion sheets. If Excel found a match, it filled the corresponding cell with a "1." If it didn't find a match, it filled the corresponding cell with

A few words on sample size. Our 325 judicial opinions contained 7552 unique case citations; the briefs in those 325 cases contained 23,479 unique case citations. It is, then, useful to think in terms of citations—7552 in the opinions and 23,479 in the briefs—rather than in terms of 325 cases. Because we tracked citations by the courts and the parties, the number of citations is the number that matters. Our sample size wasn't twenty-five in each circuit, as if we looked at the briefs and opinions in twenty-five appeals. Rather, it was the hundreds of court-cited cases and the thousands of party-cited cases in each set of briefs and resulting opinion. With this background in mind, read on to learn what we found.[79]

III. OUR RESULTS

This part summarizes the results of our study and visualizes the most important results using charts. The results focus on measures of citation stickiness because that was our primary question, but our study measured a few other things like average number of cases cited per opinion, so we included non-citation-stickiness results where they seemed interesting. In addition to figuring out what the overall stickiness percentage was, we were looking for characteristics that might have affected the number of sticky citations in a particular opinion—perhaps the type of appeal or the authoring judge's level of experience. To that end, we calculated the percentage of sticky citations of different

---

a "0." The summary sheets contained formulas that took those counts and tabulated how many cases were in each brief or opinion, how many were in briefs but not the opinion, how many were in one or two briefs only, how many cases were cited in both parties' briefs and the opinion, and—of course—how many cases were cited in the opinion that were also cited in at least one of the briefs (i.e., how many citations were sticky).

Finally, we collected all of the summary counts into a single master spreadsheet that contained data about all 325 appeals in our datasets. With this dataset we were able to calculate not only stickiness percentages for the appeals as a whole, but also stickiness percentages for specific parts of the dataset by using different characteristics of the appeals and authoring judges.

79. Although we were primarily interested in citation stickiness, we amassed data on a variety of other fronts because we thought the results were likely to be interesting. We coded for the information that was available on Westlaw at the outset (for example, Westlaw's categorization of each case's topic and the name of the authoring judge). This information allowed us to slice and dice the stickiness data based on the various characteristics of each appeal.

subsets of our data, divided by characteristic. Just to manage your expectations, our results are primarily counting, percentages, and averages (means). For some of the percentages and means, we calculated 95% confidence intervals using the Exact test in Stata.[80]

Our most common calculation was for citation stickiness, which is a proportion. We calculated the proportion of sticky cites in various subsets of briefs and opinions. A convenient feature of citation stickiness as a measurement is that it's easy to calculate, as Figure 1 shows: for opinions, tally the number of cases in an opinion also cited in at least one party's brief, and divide by the number of cases cited in the opinion. For briefs, tally the number of cases in a brief also cited in the opinion, and divide by the total number of cases cited in the brief.

Figure 1
Equations for Calculating Citation Stickiness and Super-Stickiness

$$\text{Sticky cites in opinion} = \frac{cases\ cited\ in\ both\ opinion\ and\ at\ least\ one\ brief}{cases\ cited\ in\ opinion}$$

$$\text{Sticky cites in brief} = \frac{cases\ cited\ in\ brief\ and\ opinion}{cases\ cited\ in\ brief}$$

$$\text{Super-sticky cites in opinion} = \frac{cases\ cited\ in\ opinion\ and\ both\ parties'\ briefs}{cases\ cited\ in\ opinion}$$

This part first describes our "big picture" results: aspects of citation stickiness within our whole dataset, including by circuit.[81] It then compares the stickiness of cases cited in winning briefs and the stickiness of cases cited in losing briefs, using the appeals in our dataset with clear winners and losers. And last, it summarizes the stickiness of opinions authored by judges with particular characteristics, including political affiliation and experience.

80. These confidence intervals are noted in footnotes when we first mention a particular result. A confidence interval expresses the percentage probability that data lies between two limits. *See, e.g.*, ALAN R. JONES, PROBABILITY, STATISTICS AND OTHER FRIGHTENING STUFF 102 (2019).

81. If you aren't going to read this whole article, but you are looking for a little something to read (or cite), go for the "big picture" results.

A. The Big Picture: Stickiness Percentages and Number of Cases Cited

Using the nomenclature we invented,[82] the overall stickiness percentage in our 325-opinion dataset was 49%.[83] This means that 49% of the 7552 cases that were cited in the courts' opinions had been cited by at least one party in a brief. The other 51% were endogenous—they originated from somewhere else, most likely the courts' own research.[84] Of all the cases cited in the 325 opinions, only 21% were cited in both parties' briefs. To coin some more nomenclature, we have referred to cases that were cited in both parties' briefs and then again in the opinion as super sticky. Figures 2 and 3 illustrate these overall stickiness and super-stickiness percentages of the opinions in our dataset.

Another way of looking at the stickiness data is from the perspective of the briefs, rather than the opinions. In our 325-case dataset, the parties cited 23,479 cases. Of those, only 16% were later cited by the courts in their opinions—or to use our nomenclature, were sticky. And of the 4276 cases cited by both parties in their briefs, only 38% were sticky. In other words, if a party cited a case in a brief, there was only a 16% chance that the court would also cite that same case—and an 84% likelihood that the court would not mention the case. And even for a case that both parties cited—the universe of cases that the parties agreed were worth discussing—there was only a 38% chance that the court would cite the case—and a 62% likelihood that the court wouldn't. Figure 4 illustrates the overall stickiness of citations from the briefs in our dataset.

Note that we follow the same graphical conventions for all our figures: solid dark gray for analysis of sticky citations in opinions (like Figure 2), black with white dots for analysis of super-sticky citations in opinions (like Figure 3), and white with

82. A case citation in a brief is *sticky* if it later appears in the court's opinion; if it does not appear there, then it is *unsticky*. A case citation in an opinion is *sticky* if it first appeared in any party's brief; if it appeared for the first time in the court's opinion, then it is *endogenous*. If you need more of a refresher, turn back to the first page of this article.

83. The 95% confidence interval = 47.8%–50.1%.

84. *See infra* Part IV. That's where we speculate.

black squares for analysis of sticky citations in briefs (like Figure 4).

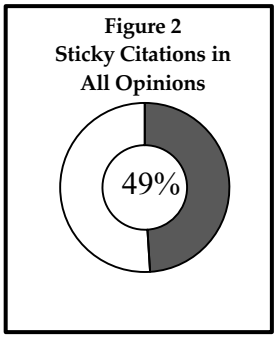

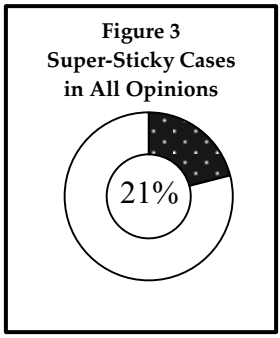

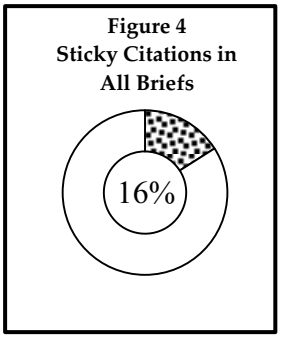

The distribution of our stickiness results, displayed in Figure 5, is mostly bell-shaped, with the center of the bell at 51 to 60% (with 55 cases in that range) and an overall peak at 31 to 40% (with 67 cases in that range). The two cases with stickiness of 10% or less involved briefs that cited relatively few cases (11 in one, 42 in the other). On the other end of the chart, nine opinions contained 100% sticky citations—all of the case cites in the opinions originated in the parties' briefs. Those opinions tended to include few case citations, eight citing ten cases or fewer,[85] and the last eleven. For comparison, the mean number of cases cited per opinion in all cases was 23.2 and the mean number cited in the parties' briefs was 72.2.

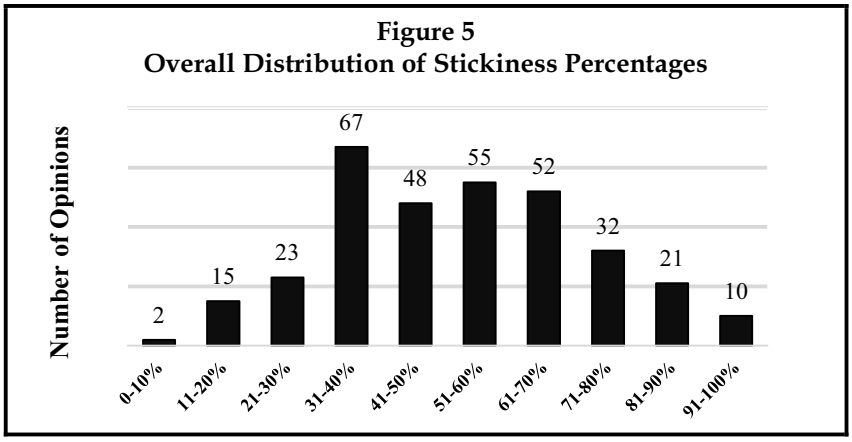

85. Indeed, two of the opinions cited only a single case. Both opinions were authored by Judge Posner.

We hoped to gain a bit more insight into our overall dataset by comparing the top and bottom quintiles of all of our stickiness data (the opinions in the top and bottom 20% in terms of citation stickiness). Each quintile was composed of sixty-five opinions. The opinions in the top quintile ranged from 70% to 100% sticky, with a mean stickiness of 79%.[86] The opinions in the bottom quintile ranged from 6% to 33% sticky, with a mean stickiness of 26%.[87] On average, the opinions in the bottom quintile cited about two-thirds the number of cases as those in the top quintile. The bottom quintile opinions also cited more than twice as many sticky cases as the top-quintile opinions. Thus the least-sticky opinions cited about one and a half times the number of cases as the stickiest opinions, but that larger number of case cites included fewer than half the number of cases cited by the parties as the stickiest opinions. Or, in simpler terms: the least-sticky opinions cite a lot of cases but very few from the briefs.

Looking at the top and bottom quintiles by number of cases cited revealed the same pattern. The quintile of opinions in our dataset that cited the most cases cited an average 47.5 cases and were 45% sticky. But the quintile of opinions in our dataset that cited the fewest cases cited an average of only 7.9 cases and were 60% sticky.

B. Stickiness by Circuit

To refresh your recollection, we analyzed citation data from twenty-five cases in each of the thirteen circuits. Because each circuit exists as a separate solar system of precedent, we hypothesized that there might be some variation in the stickiness rate among the circuits. For example, circuits may have different cultures regarding how they approach parties' briefs and the authorities cited in them or different norms regarding independent legal research. Moreover, each circuit has a non-homogeneous pool of binding precedent. Factors such as the age of the circuit, its case load, and its rate of publication would influence how much binding case law is available in each

86. The 95% confidence interval = 76.3%–81.2%.

87. The 95% confidence interval = 24.3%–28.6%.

circuit.[88] Differences in the quantity of available binding case law may affect citation stickiness.

Twenty-five cases per circuit resulted in an average of 581 case citations in the opinions and an average of 1806 cited cases in the briefs. Those data are summarized in Figures 6 and 7, but here are the highlights: stickiness percentages ranged from 42%[89] in the Third Circuit to 57%[90] in the Eighth. Super-stickiness percentages ranged from 16% in the Third Circuit to 31% in the D.C. Circuit. The Third Circuit also cited the highest number of cases per opinion (31.8), double that of the Seventh (15.9). A case cited in a party's brief had the lowest chances of being cited by the court in the Seventh Circuit (11%) and the highest in the First (20%). Even though the Third Circuit produced opinions with the lowest stickiness percentages and the largest number of case citations, an above-average percentage of cases cited to the Third Circuit were cited in the opinions: 18%. (The average for all circuits and cases was 16%.)

That brings us to the average number of cases cited in the parties' briefs. Overall, the average set of briefs cited 72.2 cases, more than three times the average number of cases cited by the court (23.2). The parties in the Federal Circuit cited the fewest cases (57.2), and parties in the Second Circuit cited the most (92.5). For more detail by circuit, see Appendix A.

Figure 6
Percentage of Sticky Citations in Opinions, by Circuit

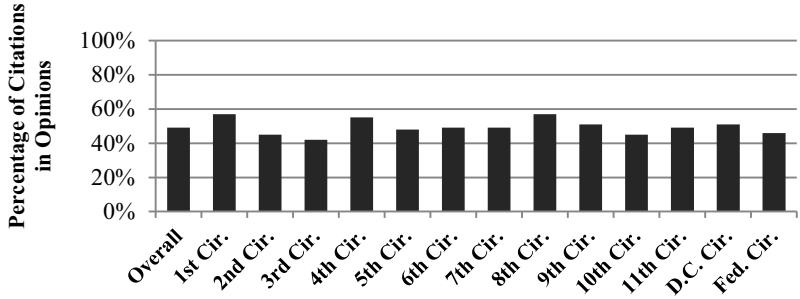

88. For example, the Eleventh Circuit and the Federal Circuit date back only to the early 1980s, although decisions of their predecessor courts were adopted as binding. *See* South Corp. v. United States, 690 F.2d 1368, 1370–71 (Fed. Cir. 1982); Bonnor v. City of Prichard, 661 F.2d 1206, 1207–11 (11th Cir. 1981).

89. The 95% confidence interval = 38.4%–45.3%.

90. The 95% confidence interval = 51.9%–61.4%.

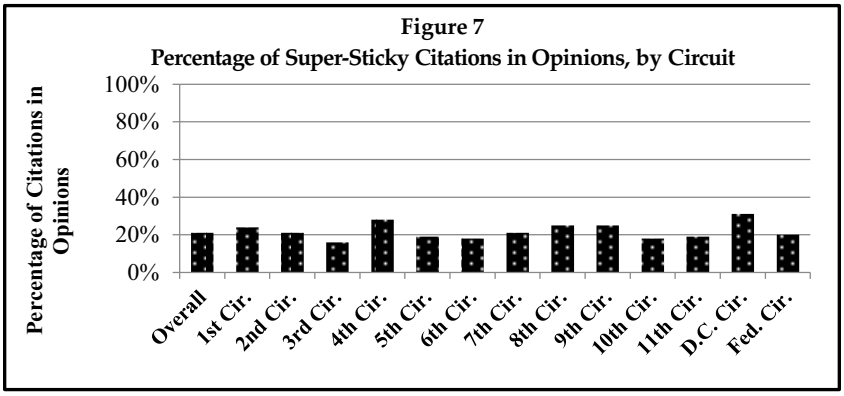

C. Stickiness by Case Characteristics

Of our 325 opinions, about two-thirds were civil (213) and one third were criminal (112).[91] Within the broad civil-criminal divide, the 325 opinions also covered many different topics. We hypothesized that case type could affect citation stickiness for a variety of reasons, including the fact that the quantity of available binding precedent would differ depending on case type. For example, the circuits may be relatively richer in binding precedent for civil matters (or particular types of civil matters) when compared to criminal matters. Variations in the amount of available binding precedent may affect citation stickiness. Moreover, criminal prosecutions involve government attorneys on one or both sides, whereas civil cases often involve private attorneys on both sides. Briefs by government and private attorneys might have different stickiness percentages, so we hypothesized that looking at broad categories of cases could provide a very rough proxy for this information.

As it turns out, the two categories of cases had similar stickiness percentages, with 48% of the cases cited in civil

91. For our purposes, the distinction between civil and criminal cases was based on the Westlaw-designated topic. Westlaw designates most opinions with both a topic and sub-topic. We categorized cases as "Criminal" if their main topic was "Criminal Justice"; all other cases were categorized as "Civil." (There is no "civil" main topic in Westlaw; examples of other main topics include "Copyright," "Environmental Law," and "Labor and Employment.") Using this method, habeas corpus cases are categorized as criminal.

opinions being sticky[92] and only a slightly higher 51% of the cases cited in criminal opinions being sticky,[93] as shown in Figure 8. The super-stickiness percentages were an identical 21%. Civil and criminal opinions also cited a similar number of cases, civil opinions citing an average of 22.4 cases and criminal opinions citing an average of 24.8. The parties in criminal cases brought fewer cases to the courts' attention, on average citing 65.9 cases in each set of briefs. On average, each set of civil briefs cited 74.9 cases. A case cited in a criminal brief had a higher likelihood of getting cited in the resulting opinion (19%) than a case cited in a civil brief (14%) as shown in Figure 9, which makes sense given similar stickiness percentages.

Figure 8
Percentage of Sticky Citations in Civil and Criminal Opinions

Figure 9
Percentage of Sticky Citations in Civil and Criminal Briefs

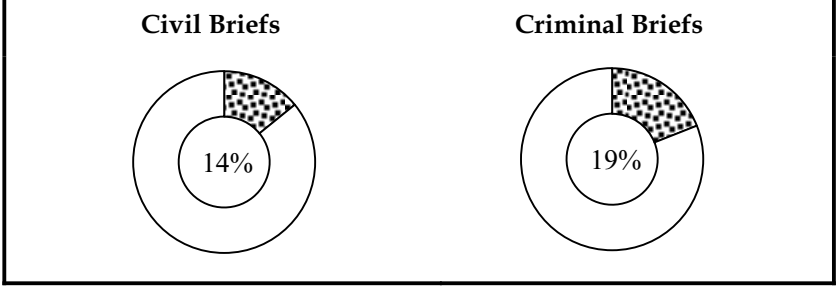

92. The 95% confidence interval = 46.4%–49.2%.
93. The 95% confidence interval = 49.2%–52.9%.

Figure 10
Opinion Stickiness by Westlaw Topic

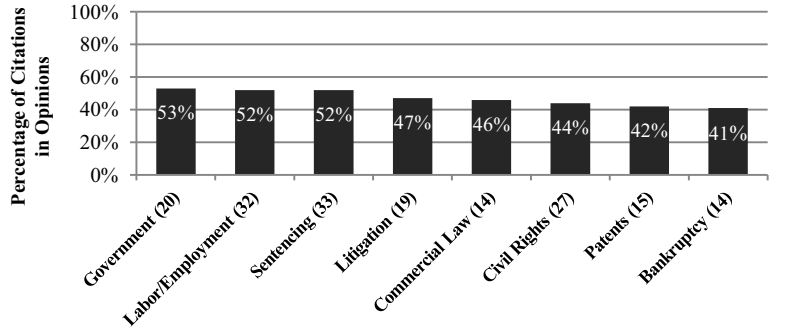

Our 325 opinions covered seventy-one different Westlaw "topics," ranging from Administrative Practice to Weapons. Many of those topics were assigned to only one (Espionage) or a few (Fraud) opinions in our dataset. We identified eight Westlaw topics that appeared at least ten times in our dataset—Bankruptcy, Civil Rights, Commercial Law, Government, Labor and Employment, Litigation, Patents, and Sentencing—and calculated stickiness percentages for each. These percentages are summarized in Figure 10, but here are the high points. The stickiest topic[94] was Government (53%) and the least sticky was Bankruptcy (41%). Government and Bankruptcy opinions also had the highest and lowest percentage of super-sticky cites (28% and 16%, respectively). But cites in Patents and Government briefs were the least sticky (12%), and cites in Sentencing briefs were the most sticky (21%). Patents and Sentencing briefs also tended to cite the fewest cases (55.3 and 52.7), while Bankruptcy and Litigation briefs cited the most (78.5 and 83.7).

D. Stickiness by Type of Brief

Although so far we've largely reported citation-stickiness data from the opinions perspective, we also looked at stickiness from the perspective of the briefs. This section describes how

94. Yes, we realize that the topic itself isn't actually sticky. Rather, it was citations in the briefs and opinions associated with the topic that were sticky. But hopefully you'll allow us a little latitude; it is our nomenclature after all.

sticky the citations were in the briefs we studied.[95] All 325 appeals in our dataset included an opening brief filed by the appellant and a responsive brief filed by the appellee.[96] Most of them—295—also included a reply brief. We hypothesized that differences in stickiness rates among the briefs could shed some light on how much weight the briefs carried with the court; for example, a very low stickiness percentage for reply briefs could indicate that reply briefs are rarely worth filing.

The three types of briefs had similar percentages of cited cases stick to the resulting opinions: 19% cited in opening briefs, 21% cited in responsive briefs, and 21% cited in reply briefs were sticky. Taking account of overlapping cases cited by appellants in their opening briefs and reply briefs, 17% of cases cited by appellants were sticky. Of cases cited by both parties' briefs, 38% were cited in the opinions (these are the super-sticky citations). Figure 11 shows these stickiness percentages.

Figure 11
Stickiness by Type of Brief

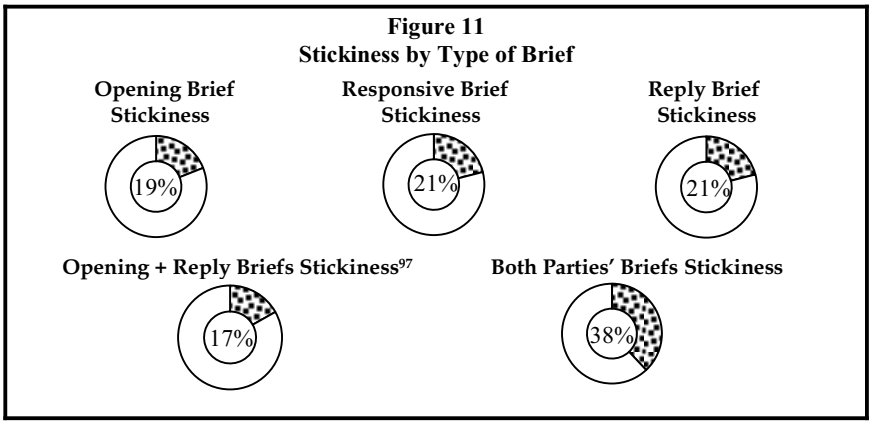

95. We also calculated the average number of cases cited per brief, which is potentially interesting to some readers: responsive briefs cited the most cases, an average of 40.5. Opening briefs cited about six fewer, an average of 34.4 cases. And reply briefs cited the fewest cases, an average of 19.8.

Of the 40.5 cases cited in the average responsive brief, an average of 30.8 cases were not cited in the opening briefs. Of those 30.8 cases, an average of 3.7 were cited in the reply brief. Thus, an average of 27.1 cases introduced by the appellee were not cited by the appellant in the reply. However, of the 3.7 cases introduced by the responsive brief and then cited in the reply brief, 26% were later cited in the opinion.

96. In some cases, the parties were petitioners and respondents rather than appellants and appellees, but this did not affect our analysis.

97. The "Opening + Reply Briefs" pie chart in Figure 11 shows a combined result based on unique citations. If the appellant cited a case in both the opening brief and the reply, we counted it only once.

E. Stickiness by Winning and Losing Briefs

A prime candidate for Top Stickiness Influencer is winningness. We hypothesized that opinions would cite more cases from the winners' briefs than the losers' because the opinions would be more likely to align with the winning briefs' reasoning (although the opinions should also include some explanation of why the losing briefs' reasoning is flawed). Moreover, judges may use the winning briefs as a jumping off point when drafting the opinion. Our calculations showed that our prediction was correct, but the effect was relatively small.

Of the 325 appeals in our dataset, 297 had a clear winner and a clear loser; the others were more muddled (for example, reversals- and affirmations-in-part). Of those 297 clear-result cases, the appellant won eighty-six and the appellee won 201. Forty-one percent[98] of the cases cited in those 297 opinions were also cited by the winning party, while only 32%[99] were cited by the loser, as illustrated by Figure 12. So our research reveals a 9% difference between the likelihood that a case cited in the opinion came from the winning party rather than from the losing party. Most of the cases cited by the winner were also cited by the loser, and vice versa: 22%[100] of the cases in the opinions were cited by both parties.

Figure 12

Percentage of Sticky Citations First Cited in Winners' and Losers' Briefs

Considering stickiness from the briefs perspective, 21% of cases cited by winners in their briefs made it into the opinions, while 17% of cases cited by losers made it into the opinions, as illustrated by Figure 13. Of the smaller universe of cases cited

98. The 95% confidence interval = 39.3%–41.8%.

99. The 95% confidence interval = 31.3%–33.7%.

100. The 95% confidence interval = 21.4%–23.5%.

by both parties in these 297 cases, 35% made it into an opinion. Winners and losers cited about the same number of cases in their briefs: an average of 40.2 in winners' briefs and 39.7 in losers' briefs.[101] The appellant also cited about the same number of cases in winning and losing reply briefs: 19.9 cases in winning reply briefs and 18.2 cases in losing reply briefs.

Figure 13
Percentage of Sticky Citations from Winners' Briefs and Losers' Briefs

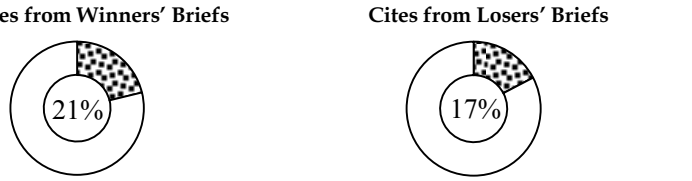

Comparing opinions in which the appellant won to opinions in which the appellee won, there isn't much difference in stickiness. Opinions in which the appellant won contained 48%[102] sticky citations, and opinions in which the appellee won contained 52%[103] sticky citations. Considering stickiness from the briefs perspective, 15% of the cases cited by winning appellants made it into an opinion, and 16% of the cases cited by winning appellees made it into an opinion. We did notice a marked difference in the stickiness of cases introduced by the appellee and then cited by the appellant in the reply brief: when the appellant won, only 19% of those cases were cited in the opinion, but when the appellee won, 29% of those cases were cited in the opinion.

F. Stickiness by Judicial Characteristics

One reason we chose to include only authored opinions in our dataset was so that we could look at various characteristics of the authoring judges to see if they had any effect on citation

101. These calculations include the total number of cases cited in both the opening and reply briefs, if both were present.

102. The 95% confidence interval = 45.4%–50.0%.

103. The 95% confidence interval = 50.1%–53.1%.

stickiness. We hypothesized that judges likely have varying levels of enthusiasm for outside legal research and varying approaches to opinion writing. We further hypothesized that some trait or combination of traits may correlate with increased or reduced stickiness.

Our full 325-opinion dataset included thirty-four appeals that produced multiple opinions, meaning a majority opinion plus at least one dissenting or concurring opinion. To keep things clean, we omitted those thirty-four opinions from our study of the authoring judges. That left 291 opinions with single-judge authorship.

For the judges authoring those 291 opinions, we looked at the following characteristics: political affiliation of the appointing President; the appointing President; the law school that awarded the judge's primary law degree; whether the judge was sitting by designation;[104] and various measures of experience, including the numbers of years since the judge's birth, since the judge received her law degree, since the judge became a judge, and since the judge received her current commission.[105]

1. Stickiness by Political Affiliation of Appointing President

We hypothesized that political ideology may affect a judges' approach to relying on independent legal research when drafting opinions.[106] We followed the approach of previous researchers and used the political affiliation of the President who appointed the judge as a proxy for the judge's ideology.[107]

---

104. Twenty opinions in our judge-characteristic dataset were authored by judges sitting by designation—meaning they were hearing cases outside their home courts. Most were district court judges, but four of the opinions in our dataset were authored by appellate judges sitting by designation in circuits other than their own.

105. We gathered the judicial-characteristics data from the Federal Judicial Center's website. *See Biographical Directory of Article III Federal Judges, 1789–present*, FED. JUDICIAL CTR. (Aug. 27, 2018), https://www.fjc.gov/history/judges.

106. Judicial ideology is the subject of numerous studies, notwithstanding the fact that it is neither easy to define nor easy to measure. *See generally* Joshua B. Fischman & David S. Law, *What Is Judicial Ideology, and How Should We Measure It?* 29 WASH. U. J.L. & POL'Y 133 (2009) (describing the theoretical problems inherent in defining "judicial ideology" and the methodological problems inherent in measuring it).

107. *See, e.g.*, Frank B. Cross & Emerson H. Tiller, *Judicial Partisanship and Obedience to Legal Doctrine: Whistleblowing on the Federal Courts of Appeals*, 107 YALE

Of the 291 opinions in our judge-characteristic dataset, 139 were authored by judges appointed by Democratic Presidents, and 152 were authored by judges appointed by Republican Presidents. The stickiness of the two sets of opinions was similar: opinions authored by Democratic-affiliated judges contained 48%[108] sticky citations, and opinions authored by Republican-affiliated judges contained 52%[109] sticky citations, as illustrated by Figure 14.[110] Twenty percent of the citations by Democratic-affiliated judges had been cited by both parties (i.e., were super-sticky), and 23% of citations by Republican-affiliated judges were super-sticky. Cases cited by both parties had about the same likelihood of being cited by judges of either affiliation, with 16% of cases cited by the parties appearing in opinions by Democratic-affiliated judges and 15% appearing in opinions by Republican-affiliated judges, as illustrated by Figure 15. Democratic-affiliated judges cited more cases per opinion—an average of 23.6 cases—to Republican-affiliated judges' 20.8.

Figure 14
Percentage of Sticky Citations Used in Opinions by Democratic-Affiliated Judges and Republican-Affiliated Judges

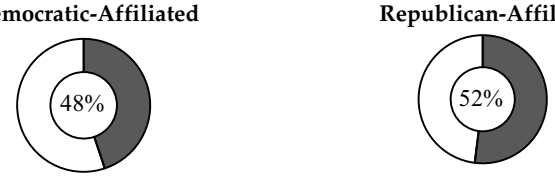

L.J. 2155, 2168 (1998); Richard L. Revesz, *Environmental Regulation, Ideology, and the D.C. Circuit*, 83 VA. L. REV. 1717, 1718–19 (1997); *see also* Fischman & Law, *supra* note 106, at 167 ("The most popular proxy for a judge's ideology, however, has been the party of the official who appointed the judge."); *but see* Fischman & Law, *supra* note 106, at 169–72 (noting potential limitations of using the political party of the appointing president as a proxy for judicial ideology).

108. The 95% confidence interval = 46.7%–50.2%.

109. The 95% confidence interval = 50.3%–53.8%.

110. Here's a fun aside: before presenting our initial research results at an academic conference, we conducted two informal polls asking what characteristics folks thought would reveal the biggest variance in stickiness. We posted one on Twitter the day before our presentation and conducted the other by a show of hands at the beginning of our presentation. One of the poll options was political affiliation of the appointing president. Given the rhetoric around the supposedly different judging habits of Democratic-affiliated and Republican-affiliated judges, we expected political affiliation to be a popular answer. We were wrong. Quite wrong. Zero of the thirty Twitter users who responded selected this answer. And zero of the thirty or so presentation attendees selected it either. Zero!

Figure 15
Percentage of Sticky Citations from Briefs Used in Opinions by Democratic-Affiliated Judges and Republican-Affiliated Judges

Next, we drilled down to the actual appointing president, whether Barack Obama (82), George W. Bush (73), Bill Clinton (48), George H.W. Bush (37), Ronald Reagan (39), Jimmy Carter (9), Gerald Ford (2), or Richard Nixon (1).

Figure 16
Percentage of Sticky Citations in Opinions, by President Who Appointed Judge

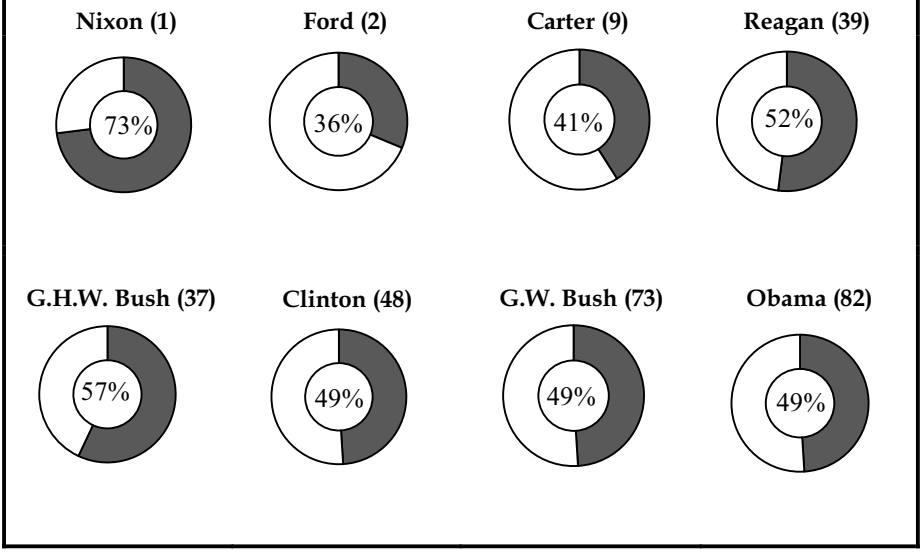

We didn't see an obvious pattern of citation stickiness that matched political ideology, illustrated in Figure 16. The single opinion authored by the judge appointed by President Nixon had the highest stickiness percentage: 73%.[111] And the two opinions

111. The 95% confidence interval = 56.1%–85.4%.

authored by judges appointed by President Ford had the lowest stickiness percentage: 36%.[112] The nine opinions authored by judges appointed by President Carter were a bit higher at 41%.[113] The number of cases cited in these three sets of opinions also varied widely, with the Nixon appointee citing forty cases, the Ford appointees citing an average of 42.5 cases, and the Carter appointees citing an average of 24.7 cases. However, one, two, and nine are pretty small as sample sizes go.

Looking at the next grouping of opinions, those authored by judges appointed by Presidents Reagan, George H.W. Bush, and Clinton, the stickiness percentages get a bit closer to the 49% overall average (and the average of 50% for opinions in the judge-characteristic dataset). The citations in the thirty-nine opinions authored by Reagan appointees were 52% sticky,[114] the citations in the thirty-seven opinions authored by George H.W. Bush appointees were 57% sticky,[115] and the citations in the forty-eight opinions authored by Clinton appointees were 49% sticky.[116] The number of opinions cited in these three sets of opinions varied, with the Reagan appointees citing an average of 15.8 cases per opinion,[117] the George H.W. Bush appointees citing an average of 21.5 cases per opinion, and the Clinton appointees citing an average of 26.1 cases per opinion.

The last two groups of opinions were authored by the most recently appointed judges, those appointed by Presidents George W. Bush and Obama. There was no stickiness difference between these two groups of opinions. The cases cited in the seventy-three opinions authored by George W. Bush appointees were 49% sticky,[118] and the opinions cited an average of 22.3 cases. The cases cited in the eighty-two opinions authored by

112. The 95% confidence interval = 26.2%–47.6%. A possibly interesting tidbit is that despite the 37% difference in stickiness between the Nixon appointee's opinion and the Ford appointees' opinions, the likelihood of a case cited in the briefs later being cited in the opinion was nearly identical for these judges. Twenty-three percent of the cases cited by the parties made it into the Ford appointees' opinions and 24% made it into the Nixon appointee's opinion.

113. The 95% confidence interval = 34.9%–48.2%.

114. The 95% confidence interval = 48.3%–56.4%.

115. The 95% confidence interval = 53.3%–60.4%.

116. The 95% confidence interval = 46.0%–51.7%.

117. The Reagan-appointee opinions include four by Judge Posner, who cited a total of thirteen cases.

118. The 95% confidence interval = 46.7%–51.6%.

Obama appointees were 49% sticky,[119] and the opinions cited an average of 22.0 cases. The super-stickiness percentages were also identical—21% for both George W. Bush and Obama appointees. As was the likelihood of a case cited in a brief later being cited in the opinion—15% for both.

2. Stickiness by Law School Attended

Next, we turned to the very serious question of whether the law school attended by the authoring judge affected the stickiness percentages in the opinions.[120] We limited our analysis to law schools that appeared at least ten times[121] in our judge-characteristic dataset. The eight law schools that met this criterion—Boston University (10 opinions), Georgetown (13), Harvard (43), Michigan (14), Texas (11), Tulane (10), Virginia (16), and Yale (23)—are shown in Figure 17. The least sticky law school was Georgetown, with alumni-authored opinions yielding an average stickiness of 40%.[122] The stickiest was Boston University, with alumni-authored opinions yielding an average stickiness of 61%.[123] The schools with the most opinions in our dataset, Harvard and Yale, were close to the 50% mean for the judge-characteristic dataset, with 51% and 49% stickiness, respectively. Opinions authored by Yale graduates cited the most cases, an average of 30.2 per opinion, nearly double the average of 16.8 cited in opinions by Virginia graduates.

119. The 95% confidence interval = 46.7%–51.3%.

120. Okay, maybe this one was not such a very serious question, but it is measurable and folks seem interested in knowing the answer. Moreover, differences in the approaches that various schools take to training their law students may affect how the schools' graduates approach opinion-writing years later when they ascend to the federal bench.

121. This does not mean law schools attended by at least ten judges in our dataset. It means that at least ten opinions were authored by judges who attended these law schools; for example, the judge-characteristic dataset includes four opinions authored by Judge Posner, which would count as four opinions for Harvard.

122. The 95% confidence interval = 34.0%–45.5%.

123. The 95% confidence interval = 53.9%–68.3%.

Figure 17
Percentage of Sticky Citations in Opinions,
by Authoring Judge's Law School

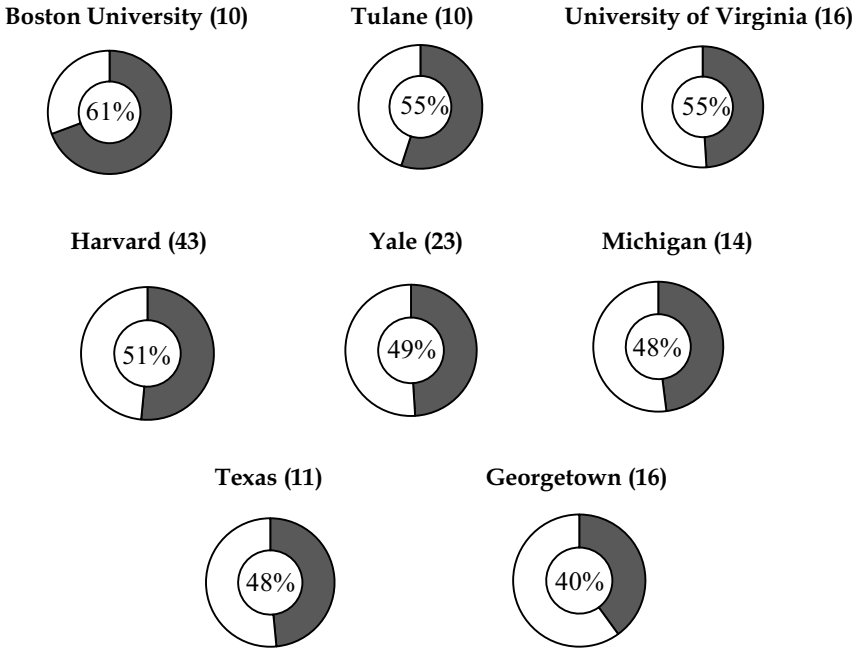

3. Stickiness by Whether Judge Sat by Designation

Of the 291 cases in our judge-characteristic dataset, twenty were authored by judges sitting by designation. Of those opinions, four were by appeals court judges and sixteen were authored by district court judges. We had expected judges who were sitting by designation to write opinions with notably different stickiness percentages than opinions authored by judges sitting with their home courts. However, it was less clear to us which way the results would cut. A judge sitting away from her home court is likely less familiar with the governing law she is called upon to apply. That may drive her to rely more heavily upon the briefs or to engage more heavily in independent research. Either way, we figured it would have some effect on citation stickiness.

We were wrong. Judges sitting by designation wrote opinions that were 49% sticky,[124] and judges sitting with their home courts wrote opinions that were 50% sticky,[125] as illustrated by Figure 18. The percentages were similarly close for trial and appellate judges: district court judges wrote opinions that were 51% sticky, and appellate judges wrote opinions that were 50% sticky, as shown in Figure 19.

Figure 18
Percentage of Sticky Citations in Opinions,
by Judges Sitting with Home Court or Sitting by Designation

Figure 19
Percentage of Sticky Citations in Opinions,
by Designated District Judges and Designated Appellate Judges

Judges sitting by designation did tend to cite six more cases per opinion than did judges sitting with their home courts. The difference between district and appellate judges in the group of designated judges was smaller; they cited an average of 25.1 and 22.0 cases per opinion, respectively.[126]

124. The 95% confidence interval = 44.5%–53.1%.

125. The 95% confidence interval = 48.9%–51.5%.

126. The four opinions authored by appellate judges sitting by designation contained a much higher number of citations on average: 36.3.

4. Stickiness by Judicial Experience

Finally, we looked at various aspects of judicial experience: years since primary law degree, years since current judicial commission, years since birth,[127] total number of judging years,[128] and senior or active status. We hypothesized that judicial experience might correlate with increased or decreased citation stickiness. Perhaps inexperienced judges would rely more heavily on the authorities cited by the parties, or perhaps their lack of experience would drive them to thoroughly research the law in an effort to educate themselves. Perhaps more experienced judges would give less weight to the parties' briefs and arguments and therefore be more likely to diverge from the authorities cited by the parties. Or perhaps age and experience would correlate with a sort of judicial lethargy in which the judge becomes less likely to look beyond the parties' briefs.

With the exception of the binary categories of senior and active status, we calculated citation stickiness for all the experience categories by dividing the data into decade-long buckets: one to ten years, eleven to twenty years, and so on. We expected judicial experience to affect stickiness, but overall we didn't observe any consistent pattern. The results are summarized visually in Figures 20 through 25, but if you prefer to read the results in text form, buckle up for the next five paragraphs.

Let's start with years since birth. The youngest judges in our judge-characteristic dataset were, in 2017, between 41 and 50 years out from their births. As Figure 20 shows, the youngest cohort authored seventeen opinions in our dataset and had the second-lowest stickiness: 48%. This cohort also cited the fewest cases in their opinions, an average of 18.9. The next four cohorts (51 to 60, 61 to 70, 71 to 80, and 81 to 90) had stickiness percentages hovering around the 49% mean, but gradually increasing as their age increased (respectively, 49%, 49%, 49%,

127. To calculate the number of years since birth, we subtracted the year of the judge's birth from the year 2017, the year from which we collected data. We did not use age because the benefits seemed minimal compared to the cost of calculating each judge's age at the time the opinion was issued.

128. Many federal judges have prior judging experience on a state court, another federal court, or both.

52%, and 52%). The oldest cohort of judges, who were between 91 and 100 years out from their births, authored three opinions. These opinions had the lowest stickiness—45%. The five age-range cohorts cited a similar number of cases in their opinions (respectively, 21.5, 23.4, 20.9, 24.2, and 22.0).

After being born, the next step to becoming a judge in our dataset is attaining a law degree. Figure 21 shows that there were four opinions in our dataset by judges who had attained their law degrees between eleven and twenty years before 2017. Those four opinions had the lowest stickiness percentage—43%. The next five cohorts (21 to 30, 31 to 40, 41 to 50, 51 to 60, and 61 to 70) had similar stickiness percentages (respectively, 51%, 50%, 49%, 51%, and 55%). Those who had been law school graduates the longest did cite notably more cases per opinion: 31.6 as compared to 22.8, 20.3, 22.2, 22.6, and 21.1.

At some point after attaining their law degrees, the judges in our dataset were of course commissioned as federal appellate judges.[129] The most newly minted of them had received their commissions between one and ten years before 2017. These judges authored opinions that contained citations that were 49% sticky. In fact, as you can see from Figure 22, years since commission didn't seem to move stickiness much one way or the other. The next four cohorts (11 to 20 years, 21 to 30 years, 31 to 40 years, and 41 to 50 years) had similar stickiness percentages (respectively, 48%, 55%, 49%, and 48%). The greatest variation in cases cited was between the two longest-commissioned cohorts: opinions by judges who had been commissioned for between 31 and 40 years cited an average of 17.5 cases, and opinions by judges who had been commissioned for between 41 and 50 years cited an average of 41.7 cases. However, there were only three opinions in the last group.

129. When we mention a judge's years since commission, what we mean is the years since the judge was commissioned on whatever court she was on when she authored the opinion in our dataset. For judges not sitting by designation, that is the same court that rendered the opinion. As an example, for a First Circuit opinion authored by a First Circuit judge, the judge's years since commission would be the number of years since her commission on the First Circuit. Prior judging experience, including prior judging experience as a federal district court judge, was not included. However, for a district court judge sitting by designation, we measured the years since the judge was commissioned as a district court judge.

Our penultimate category was total judging experience, which is reported in Figure 23. We added up each judge's years of judging—whether as a state judge or on any federal court. The least-experienced cohort of judges had been judging for between one and ten years. These judges authored forty-six of the opinions in our judge-characteristic dataset, and these opinions contained citations that were 48% sticky. The next four cohorts (11 to 20 years, 21 to 30 years, 31 to 40 years, and 41 to 50 years) also had stickiness percentages around the mean (respectively, 49%, 54%, 47%, and 49%). The most-experienced cohort, judges who had been judging for between 51 and 60 years, authored opinions with the least sticky citations, having an average of 44%. However, there were only two opinions in that cohort. The number of cases cited per opinion didn't vary much across the cohorts: 20.7, 25.0, 21.8, 20.9, 20.4, and 26.0.

At last, we reach senior and active status. Federal judges may take senior status based on their age and years of service; thus, it is a proxy for experience.[130] Within our judge-characteristic dataset, 224 opinions were authored by active-status judges and 67 opinions were authored by senior-status judges. We didn't see a great difference in stickiness: opinions authored by senior judges contained citations that were 48% sticky[131] and opinions authored by active judges contained citations that were 51% sticky.[132] Active judges cited fewer cases in their opinions, an average of 20.9 to senior judges' 26.2.

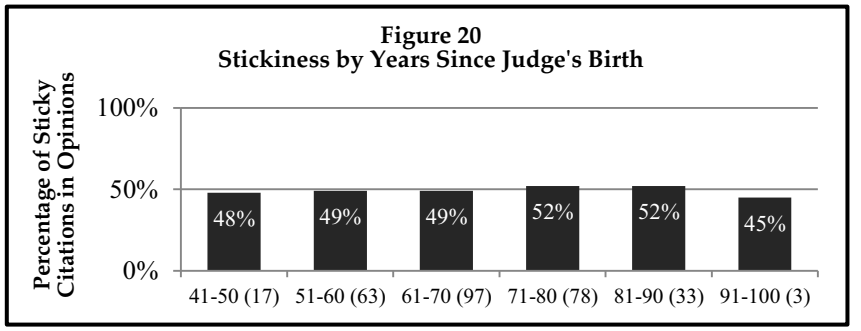

130. *See* 28 U.S.C. § 371(b)(1), (c); *see also* RICHARD A. POSNER, THE FEDERAL COURTS: CHALLENGE AND REFORM 32 (1996) (explaining the criteria for senior status); Stephen B. Burbank, S. Jay Plager & Gregory Ablavsky, *Leaving the Bench, 1970–2009: The Choices Federal Judges Make, What Influences Those Choices, and Their Consequences*, 161 U. PA. L. REV. 1, 11, 19–20 (2012) (same).

131. The 95% confidence interval = 49.5%–52.3%.

132. The 95% confidence interval = 45.3%–50.0%.

Figure 21
Stickiness by Years Since Judge's Law Degree

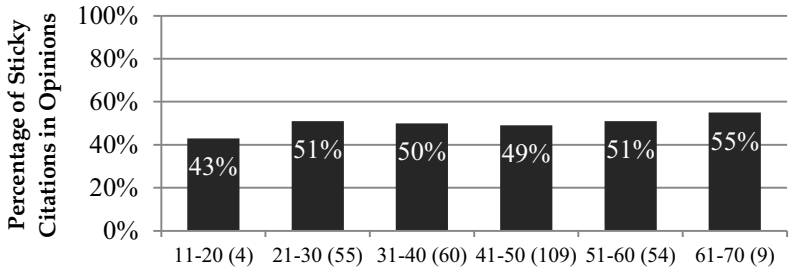

Figure 22
Stickiness by Years Since Judge's Commission

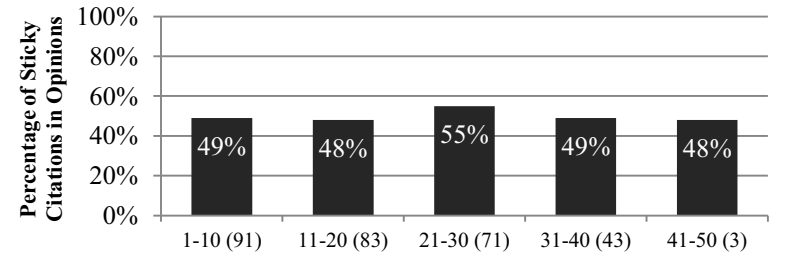

Figure 23
Stickiness by Judge's Total Years of Service on All Courts

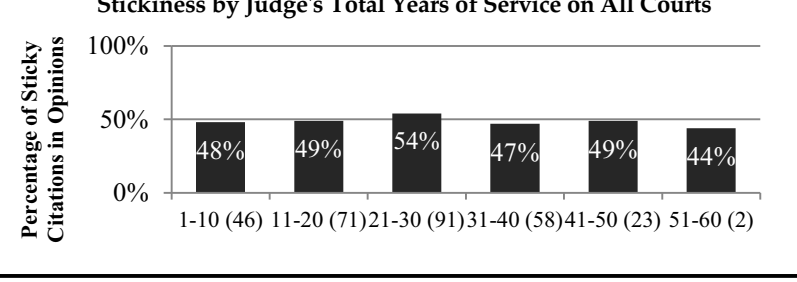

Figure 24
Stickiness by Judge's Active or Senior Status

The text and figures in the section that you just read describe an overview of our results. You can find a more detailed presentation of the results in Appendices A, B, C, and D.

IV. WHAT IT ALL MIGHT MEAN

So what do these results mean? Depending on what you want to know, they mean either a great deal or not very much at all.

If you want to know how often federal courts of appeals cite to the same authorities as the parties, then our results will tell you a lot. Forty-nine percent of the cases cited in published, authored federal appellate opinions were previously cited in at least one of the parties' briefs. Fifty-one percent of the cases cited were not mentioned in any of the briefs. To us, that 49% figure is notable because it is lower than we expected it would be.[133]

133. To be clear, we did not distinguish among which parts of the opinions (or briefs) cited the cases. For example, it may well be that judges or courts have generic procedural passages that they frequently cut and paste into opinions to recount well-worn standards of review. Thus, the cases cited in the procedural passages may have stickiness rates below 49% and the cases cited in the "substantive" legal analysis may have stickiness rates above 49%. This is certainly an additional avenue for exploration, although the Marvell study provides some relevant information here. *See supra* notes 24–30 and accompanying text. In his study, Dr. Marvell found that only one-sixth of the cases cited in the parties' briefs were mentioned by the court. MARVELL, *supra* note 24, at 132. He surmised that this low percentage may have been the product of the parties pressing arguments that the court did

But if you want to know why the stickiness is 49%, we cannot tell you that. We can speculate though. Where do the other 51% of cited cases come from? Intuition—and our own clerkship experiences—point to independent research by the courts. Either the judges or members of the judges' staffs are likely doing their own independent research and locating cite-worthy cases that the parties never mentioned.[134]

In an ideal world, the parties would present the court with the relevant cases and the court would discuss those same relevant cases. Thus, the stickiness percentage of the cases cited in the briefs would be at or near 100%. Courts would not need to engage in independent legal research to locate relevant authorities and there would be few or no endogenous case citations in the opinions.[135]

That being said, a stickiness percentage of 100% would not necessarily mean that the system was functioning well. A 100%

---

not ultimately address. Indeed, he found that the court did not reach one quarter of the legal arguments that the parties made in their briefs. *Id.* However, after omitting these undecided legal arguments, Dr. Marvell still found that only one-fifth of the cases cited in the parties' briefs were mentioned by the court. *Id.* Thus, looking only at the citations relating to the substantive issues decided by the court still yielded a rather low percentage of citations transmitted from the briefs to the opinion.

Additionally, our 49% stickiness figure does not account for depth of treatment. Our study did not distinguish between cases that were cited only once in the parties' briefs and cases that were cited ten times. Presumably, the stickiness percentage is higher for cases that garnered multiple citations in the parties' briefs. This is another avenue for additional research.

134. Perhaps some small percentage of cases cited in the opinions in our dataset were mentioned by the parties somewhere other than the briefs—like at oral argument—and then picked up by the court. We didn't review the oral argument transcripts, so we can't say for sure, but our experience and Dr. Marvell's previous research indicates that this is a rare occurrence. *See* MARVELL, *supra* note 24, at 133 (finding that only 10% of the authorities cited in the briefs were later mentioned at oral argument and only 1% of the total authorities cited in the opinions were mentioned at oral argument but not mentioned in the parties' briefs). Indeed, one reason that we didn't deem it necessary to review the oral-argument transcripts is because it is so rare for a case to be mentioned for the first time at oral argument.

Another possibility is that some percentage of the cases cited in the appellate opinions in our dataset were cited in the underlying trial-level litigation—including the district court's opinion that was the focus of the appeal—yet were not mentioned in either of the parties' appellate briefs. Again, this seems likely to be rare.

135. Some endogenous case citations may be acceptable, however, because some legal rules are identically repeated in numerous cases. Thus, for well-worn legal rules like standards of review, it may matter little that the parties are citing one case and the court is citing a different case as long as everyone is applying the identical rule.

stickiness average could be achieved if the parties cited a bunch of irrelevant cases and the court was too idle to do any research outside of the authorities cited in the briefs. On the other hand, a stickiness percentage near zero would necessarily mean that the system was working poorly. The parties and the court would necessarily be talking past each other, either because the parties were citing wholly irrelevant authorities or because the court went rogue and decided the matter on some spontaneous ground without the benefit of the parties' input, or both.

So if we know that something approaching 100% stickiness is our aspiration (although not in itself a sufficient indicator of a properly functioning appellate process) and we know that something approaching 0% stickiness is necessarily a sign of major dysfunction in the system, what do we make of 49%? A 49% stickiness figure tells us that we don't live in an ideal world, but it doesn't tell us precisely where the breakdowns are occurring. Perhaps the critics are correct and either attorneys or courts or both are producing less-than-stellar work product. Perhaps attorneys' briefs routinely feature numerous irrelevant citations and not enough relevant ones. Maybe courts are ignoring relevant cases and lines of argument that were raised by the parties.[136] It could be that courts are deciding cases based on their own reasoning without the benefit of the parties' input, perhaps driven by the desire to rule in a particular party's favor regardless of whether that party has the better legal argument.[137] Or it could be that judges—or their law clerks—strive to be deliberately unsticky in an attempt to outdo the parties and impress the world with garish displays of legal citation.[138]

---

136. Or, at the very least, courts may be failing to explain why the cases and lines of argument that they treat as irrelevant are in fact irrelevant.

137. *See* RICHARD A. POSNER, HOW JUDGES THINK 74 (2008) (noting that certain judges may be "more likely to focus on the 'equities' of the individual case—the aspects of the case that tug at the heartstrings—and less on its precedential significance"). We would think, however, that such judicial behavior—if in fact it occurs—would be more likely to be relegated to unpublished opinions. *See* Richard S. Arnold, *Unpublished Opinions: A Comment*, 1 J. APP. PRAC. & PROCESS 219, 223 (1999) (pointing out that "if, after hearing argument, a judge in conference thinks that a certain decision should be reached, but also believes that the decision is hard to justify under the law, he or she can achieve the result, assuming agreement by the other members of the panel, by deciding the case in an unpublished opinion and sweeping the difficulties under the rug").

138. Excessive citation is a well-documented plague in legal writing. *See, e.g.*, Alexa Z. Chew, *Stylish Legal Citation*, 71 ARK. L. REV. 823, 854, 856–57 (2019) (describing

Indeed, one trend that we observed was that the judicial opinions that cited more cases tended to have lower citation-stickiness percentages than opinions containing relatively fewer case citations.[139] Looking at number of cases cited, the top quintile of opinions in our dataset cited an average of 47.5 cases. The bottom quintile of opinions in our dataset cited an average of 7.9 cases. The citation stickiness varied meaningfully between these two groups of opinions: the top-quintile opinions were 45% sticky, but the bottom-quintile opinions were 60% sticky. Perhaps judges who do more independent research feel compelled to show their work by citing more cases.

Another research-based factor driving down citation stickiness may be the variety of research tools and methods. Perhaps courts and attorneys use different research techniques or platforms and are exposed to different spheres of research results. If a court relies on Westlaw and the parties rely on Lexis, similar searches could simply be returning different results.[140] Or even more innocuously, perhaps many cases are simply interchangeable. If ten or twenty cases all state the same proposition, then the parties and the court may cite to different cases while discussing the same legal rules or lines of reasoning. In that scenario, the court's independent research turned up the same legal principles as the parties' research, which does not suggest dysfunction.

Through comparison to Dr. Marvell's study, one thing that we do know is that stickiness percentages haven't changed much in recent decades. Our results are extremely similar to the results

---

hypercitation as citing more authorities than are needed to support the proposition asserted).

139. *See supra* notes 85–87 and accompanying text.

140. *See* Susan Nevelow Mart, *Results May Vary: Which Database a Researcher Uses Makes a Difference*, 104 A.B.A. J. 48, 48–49 (Mar. 2018) ("In a comparison of six legal databases—Casetext, Fastcase, Google Scholar, Lexis Advance, Ravel and Westlaw—when researchers entered the identical search in the same jurisdictional database of reported cases, there was hardly any overlap in the top 10 cases returned in the results."); Susan Nevelow Mart, *The Relevance of Results Generated by Human Indexing and Computer Algorithms: A Study of West's Headnotes and Key Numbers and LexisNexis's Headnotes and Topics*, 102 LAW LIBR. J. 221, 241, 240–44 (2010) (comparing the results of using West and Lexis features for research and concluding that "each one will find some relevant cases on your legal topic" but "it does not look like any one of them can be used to find all relevant cases on your legal topic").

of that study from the 1970s.[141] Dr. Marvell found that 48% of the authorities cited in the opinions in his dataset had previously been cited in the parties' briefs.[142] We found that 49% of the cases cited in the opinions in our dataset had previously been cited in the parties' briefs.[143] Dr. Marvell found that 17% of the authorities cited in the parties' briefs were later mentioned in the opinions.[144] We found that 16% of the cases cited in the parties' briefs were later mentioned in the opinions.[145]

While Dr. Marvell studied a state supreme court in the early 1970s,[146] we studied the federal courts of appeals in 2017.[147] Yet the similarity of our findings is notable, indicating that the evolution of the research process from books to digital databases[148] has had little effect on citation stickiness. Given the consistency between our findings and those of the Marvell study, future researchers can use our findings—49% sticky citations in opinions and 16% sticky citations in briefs—as baselines to which to compare their results.[149]

---

141. *See supra* notes 24–30 and accompanying text.

142. MARVELL, *supra* note 24, at 134.

143. *See supra* Part III.

144. MARVELL, *supra* note 24, at 134.

145. *See supra* Part III. Sixteen percent was a direct match with what another team of researchers found when studying thirty First Circuit cases. *See* Oldfather, et al., *supra* note 48, at 1238.

146. MARVELL, *supra* note 24, at 6.

147. *See supra* Part II.

148. Over the past quarter century, legal research has shifted from a book-bound endeavor to one that is often entirely digital. *See, e.g.*, KENT C. OLSON, PRINCIPLES OF LEGAL RESEARCH § 1.2, 5 (2009) (acknowledging that "[m]ost research these days is conducted on the Internet"); *see also* Ellie Margolis & Kristen Murray, *Using Information Literacy to Prepare Practice-Ready Graduates*, 39 U. HAW. L. REV. 1, 3 (2016) ("Today's students do not remember a time when research could only be done in books, when the only way to access cases was through a digest, when near-instantaneous access to information did not exist." (footnote omitted)); Ellie Margolis & Kristen E. Murray, *Say Goodbye to the Books: Information Literacy as the New Legal Research Paradigm*, 38 U. DAYTON L. REV. 117, 121–26 (2012) (describing the migration from print to digital sources); John Palfrey, *Cornerstones of Law Libraries for an Era of Digital-Plus*, 102 L. LIBR. J. 171, 172–77 (2010) (describing trends in legal information use).

149. For example, future research could uncover whether citations in amicus briefs are similarly sticky. That would shed light on whether amicus briefs are worthwhile. Likewise, future research could uncover whether particular types of cases—such as nonbinding cases in general or unpublished cases in particular—include citations that are similarly sticky. Again, this could inform briefwriting decisions.

Our results also demonstrate that judicial characteristics play little role in citation stickiness. The stickiness percentage was similar for judges appointed by Democratic (48%) and Republican (52%) presidents.[150] Whether the authoring judge was a district court judge (51%) or an appellate judge (50%) also mattered little, as did whether the authoring judge was sitting by designation (49%) or sitting with her home court (50%).

None of our metrics of judicial experience yielded significantly different stickiness rates. Active-status judges (51%) were similarly sticky to senior-status judges (48%). While there were some isolated differences in stickiness when it came to age and judging experience, there were almost no observable trends. For example, the oldest judges were the least sticky (45%), but judges in the two next oldest age brackets were the most sticky (52%).[151] Setting aside two very small sample sizes, there was almost no variation in stickiness based on how long it had been since the authoring judge graduated from law school.[152] Mild variations were exhibited based on the number of years since the authoring judge had been

150. While there wasn't a meaningful difference between the stickiness percentages of opinions authored by appointees of Democratic and Republican presidents, there was some variation among judges appointed by individual presidents. The most noteworthy variation was the 57% stickiness of opinions authored by judges appointed by President George H.W. Bush. This stickiness percentage was notable because it was the product of a not-insignificant sample size (thirty-seven opinions) and was noticeably higher than other recent presidents with not-insignificant sample sizes (all of whom clustered around stickiness rates between 49 and 52%). Perhaps judges appointed by President George H.W. Bush possess an approach to decisionmaking that is uniform to a unique extent. Another researcher has found that judges appointed by President George H.W. Bush are more partisan than judges appointed by any other President from Carter to George W. Bush. *See* Corey Rayburn Yung, *Beyond Ideology: An Empirical Study of Partisanship and Independence in the Federal Courts*, 80 GEO. WASH. L. REV. 505, 539 (2012). Perhaps some relationship exists between partisanship and citation stickiness.

151. *See* Figure 20, *supra* page 103. Setting aside judges older than ninety because there were only three results in that category, the other five decades showed almost no differences in stickiness. In ascending order, the stickiness rates by decade were 48%, 49%, 49%, 52%, and 52%.

152. *See* Figure 21, *supra* page 104. In ascending order, the stickiness rates by decade for the other four brackets were 51%, 50%, 49%, and 51%.

commissioned a federal judge[153] or had started judging,[154] but again without any consistent trend.

Likewise, the type of case, broadly speaking, had little effect on stickiness. Civil cases (48%) and criminal cases (51%) were similarly sticky. And, although there were some differences in stickiness rates when the cases were sorted by Westlaw topic, we simply don't have a big enough dataset to determine whether those differences are meaningful.[155]

If judge and case characteristics don't matter much, potential avenues for further research could involve isolating characteristics of attorneys or their briefs. We observed that only 16% of the cases cited in the parties' briefs later appear in the resulting opinions. But we wouldn't expect this number to approach 100%, even in an ideal adversarial system. Parties sensibly make alternative arguments in their briefs. Appellate courts, being economical, decide appeals on as few grounds as possible. Thus, an appellant may urge reversal based on three alleged errors. But if the appellate court finds one of the errors reversible it may not analyze the other two. With this in mind, the parties understandably cite more authorities than the courts. Yet, 16% struck us as an unexpectedly low stickiness rate for cases cited in the parties' briefs, and the 38% stickiness rate for cases that had been cited in both parties' briefs struck us as especially surprising. A case that both parties deemed relevant to discuss still had a 62% likelihood of not being mentioned by the court. Maybe this is the product of too many string citations in the briefs or too many alternative arguments, but maybe it is also the product of courts failing to adequately engage with the authorities cited by the parties.[156]

---

153. *See* Figure 22, *supra* page 104. For the four brackets with more than ten results, the stickiness rates based on years since commission were—in ascending order—49%, 48%, 55%, and 49%.

154. *See* Figure 23, *supra* page 104. For the five brackets with more than ten results, the stickiness rates based on total judging years were—in ascending order—48%, 49%, 54%, 47%, and 49%.

155. When sorted by Westlaw topic, the stickiest topics were Government (53%), Labor and Employment (52%), and Sentencing (52%), and the least sticky were Bankruptcy (41%) and Patents (42%). But none of those topics had more than twenty cases in our dataset.

156. Indeed, too many string citations could prevent a court from adequately engaging with the cited authorities. *See* Chew, *supra* note 138, at 854 (explaining that excessive

We observed that courts tended to borrow citations more heavily from winning parties' briefs than from losing parties' briefs.[157] Thus, the key to getting the court to borrow more of a party's citations may simply be for the party to win the case. Or, perhaps, the way for a party to win a case is to have the court adopt many of its case citations. Apart from winning and losing, some pre-decisional attribute or set of attributes may tend to make some attorneys' briefs stickier than others. Stay tuned for further research.

V. CONCLUSION

The big takeaway from this empirical study is that courts—not parties—accounted for more than half of the 7552 cases cited in the 325 judicial opinions we looked at. This result surprised us, as it surprised the small Twitter and conference audiences that we polled. Based only on this study, we don't know why so many cases cited in opinions were endogenous. But our discovery that more than 50% of the cases cited in judicial opinions are endogenous is a novel finding.

Earlier studies that looked at stickiness used much smaller samples of cases decided in only a few courts; our study provides results based on a much more robust dataset.[158] But beyond just filling a gap in the data, we hope that our study of citation stickiness provides an empirical foundation for future discussions on legal persuasion. In theory, the parties' research should have a great deal of influence on judicial decisions. Yet in practice, citation stickiness is both variable and, on average, lower than that theory predicts.

Our study suggests that scholars of legal persuasion and judicial decisionmaking should consider citation stickiness because it's an easy-to-calculate proxy for the influence that parties' research has on judicial decisions: divide the total number of cases cited in both an opinion and the briefs by the number of cases cited in the opinion. This quick calculation reveals to what degree the parties' research "stuck" to the

---

citations "drain[] reader energy" and summarizing various legal writing experts' distaste for string citations).

157. *See supra* notes 98–103 and accompanying text.

158. *See supra* Section I.A.

judicial decision, which in turn can inform any decision about whether additional, more labor-intensive study into the whys and hows is warranted.

APPENDIX A
APPEAL CHARACTERISTICS

	Opinions	Cases Cited in Briefs	Cases Cited in Opinions	Average Cases Cited per Opinion	Sticky Cites in Opinions	Super-Sticky Cites in Opinions
Overall						
All Appeals	325	23479	7552	23.2	3699	1613
Circuit						
1st Circuit	25	1770	626	25.0	354	152
2d Circuit	25	2312	663	26.5	299	137
3d Circuit	25	1867	796	31.8	333	128
4th Circuit	25	2002	436	17.4	240	121
5th Circuit	25	1587	574	23.0	277	109
6th Circuit	25	1571	617	24.7	301	110
7th Circuit	25	1762	397	15.9	196	83
8th Circuit	25	1613	439	17.6	249	110
9th Circuit	25	2110	552	22.1	283	136
10th Circuit	25	1947	777	31.1	352	141
11th Circuit	25	2054	667	26.7	324	129
D.C. Circuit	25	1455	520	20.8	265	159
Fed. Circuit	25	1429	488	19.5	226	98
Civil or Criminal						
Civil	213	15962	4772	22.4	2280	1019
Criminal	112	7385	2780	24.8	1419	594
Westlaw Topic						
Bankruptcy	14	1099	364	26.0	150	57
Civil Rights	27	2003	856	31.7	377	150
Commercial	14	1077	367	26.2	167	83
Government	20	1359	316	15.8	168	90
Labor/Employment	32	2281	699	21.8	360	167
Litigation	19	1591	421	22.2	199	97
Patents	15	830	247	16.5	103	41
Sentencing	33	1740	699	21.2	364	142

	Opinions			Briefs
	Percent Sticky	Percent Endogenous	Percent Super-Sticky	Percent Sticky
Overall				
All Appeals	49%	51%	21%	16%
Circuit*				
1st Circuit	57%	43%	24%	20%
2d Circuit	45%	55%	21%	13%
3d Circuit	42%	58%	16%	18%
4th Circuit	55%	45%	28%	12%
5th Circuit	48%	52%	19%	17%
6th Circuit	49%	51%	18%	19%
7th Circuit	49%	51%	21%	11%
8th Circuit	57%	43%	25%	15%
9th Circuit	51%	49%	25%	13%
10th Circuit	45%	55%	18%	18%
11th Circuit	49%	51%	19%	16%
D.C. Circuit	51%	49%	31%	18%
Fed. Circuit	46%	54%	20%	16%
Civil or Criminal				
Civil	48%	52%	21%	14%
Criminal	51%	49%	21%	19%
Westlaw Topic				
Bankruptcy	41%	59%	16%	14%
Civil Rights	44%	56%	18%	19%
Commercial Law	46%	54%	23%	16%
Government	53%	47%	28%	12%
Labor/Employment	52%	48%	24%	16%
Litigation	47%	53%	23%	13%
Patents	42%	58%	17%	12%
Sentencing	52%	48%	20%	21%

*The 95% confidence intervals for each circuit: First Circuit = 52.6%–60.5%; Second Circuit = 41.3%–49.0%; Third Circuit = 38.4%–45.3%; Fourth Circuit = 50.2%–59.8%; Fifth Circuit = 44.1%–52.4%; Sixth Circuit = 44.8%–52.8%; Seventh Circuit = 44.3%–54.4%; Eighth Circuit = 51.9%–61.4%; Ninth Circuit = 47.0%–55.5%; Tenth Circuit = 41.8%–48.9%; Eleventh Circuit = 44.7%–52.4%; D.C. Circuit = 46.6%–55.3%; Federal Circuit = 41.8%–50.8%.

APPENDIX B
JUDGE CHARACTERISTICS

	Opinions	Cases Cited in Briefs	Cases Cited in Opinions	Average Cases Cited per Opinion	Sticky Cites	Super-Sticky Cites
Affiliation of Appointing President						
Democratic	139	10113	3279	23.6	1588	658
Republican	152	10794	3161	20.8	1633	720
Excluded	34	2440	1112	32.7	478	235
Appointing President						
Obama	82	5893	1804	22.0	884	387
GW Bush	73	5186	1625	22.3	798	339
Clinton	48	3404	1253	26.1	612	239
GHW Bush	37	2728	796	21.5	453	220
Reagan	39	2625	615	15.8	322	142
Carter	9	816	222	24.7	92	32
Ford	2	133	85	42.5	31	7
Nixon	1	122	40	40.0	29	12
Excluded	34					
Law School						
BU	10	623	186	18.6	114	48
Georgetown	13	794	290	22.3	115	42
Harvard	43	2739	925	21.5	474	220
Michigan	14	1155	260	18.6	126	59
Texas	11	659	209	19.0	100	47
Tulane	10	747	272	27.2	149	61
UVA	16	1302	268	16.8	148	61
Yale	23	1856	694	30.2	341	121
Sitting by Designation—All						
Designated	20	1663	547	27.4	267	98
Home	271	19244	5893	21.7	2954	1280
Sitting by Designation—District or Appellate						
District	16	1276	402	25.1	206	80
Appellate	275	19631	6038	22.0	3015	1298

	Opinions			Briefs
	Percent Sticky	Percent Endogenous	Percent Super-Sticky	Percent Sticky
Affiliation of Appointing President				
Democratic	48%	52%	20%	16%
Republican	52%	48%	23%	15%
Appointing President				
Obama	49%	51%	21%	15%
GW Bush	49%	51%	21%	15%
Clinton	49%	51%	19%	18%
GHW Bush	57%	43%	28%	17%
Reagan	52%	48%	23%	12%
Carter	41%	59%	14%	11%
Ford	36%	64%	8%	23%
Nixon	73%	28%	30%	24%
Law School				
BU	61%	39%	26%	18%
Georgetown	40%	60%	14%	14%
Harvard	51%	49%	24%	17%
Michigan	48%	52%	23%	11%
Texas	48%	52%	22%	15%
Tulane	55%	45%	22%	20%
UVA	55%	45%	23%	11%
Yale	49%	51%	17%	18%
Sitting by Designation—All				
By Designation	49%	51%	18%	16%
Home	50%	50%	22%	15%
Sitting by Designation—District or Appellate				
District	51%	49%	20%	16%
Appellate	50%	50%	21%	15%

	Opinions	Cases Cited in Briefs	Cases Cited in Opinions	Sticky Cites	Super-Sticky Cites	Average Cases Cited per Opinion
Years from Birth						
41 to 50	17	1222	322	155	60	18.9
51 to 60	63	4756	1352	660	278	21.5
61 to 70	97	6475	2269	1111	474	23.4
71 to 80	78	5812	1634	852	437	20.9
81 to 90	33	2394	797	413	177	24.2
91 to 100	3	248	66	30	11	22.0
Years since Law Degree						
11 to 20	4	384	91	39	20	22.8
21 to 30	55	3819	1116	564	238	20.3
31 to 40	60	4392	1334	668	312	22.2
41 to 50	109	7417	2464	1217	475	22.6
51 to 60	54	4015	1140	578	264	21.1
61 to 70	9	783	284	157	71	31.6
Years from Commission						
1 to 10	91	6706	2035	1003	440	22.4
11 to 20	83	5863	2050	976	395	24.7
21 to 30	71	5080	1476	815	367	20.8
31 to 40	43	3003	754	367	157	17.5
41 to 50	3	255	125	60	19	41.7
Total Judging Years						
1 to 10	46	3245	953	457	204	20.7
11 to 20	71	5728	1773	863	354	25.0
21 to 30	91	6147	1980	1077	490	21.8
31 to 40	58	4002	1212	571	219	20.9
41 to 50	23	1608	470	230	104	20.4
51 to 60	2	177	52	23	7	26.0
Active or Senior						
Active	224	16111	4684	2384	1040	20.9
Senior	67	4796	1756	837	338	26.2

	Opinions			Briefs
	Percent Sticky	Percent Endogenous	Percent Super-Sticky	Percent Sticky
Years from Birth				
41 to 50	48%	52%	19%	13%
51 to 60	49%	51%	21%	14%
61 to 70	49%	51%	21%	17%
71 to 80	52%	48%	27%	15%
81 to 90	52%	48%	22%	17%
91 to 100	45%	55%	17%	12%
Years since Law Degree				
11 to 20	43%	57%	22%	10%
21 to 30	51%	49%	21%	15%
31 to 40	50%	50%	23%	15%
41 to 50	49%	51%	19%	16%
51 to 60	51%	49%	23%	14%
61 to 70	55%	45%	25%	20%
Years from Commission				
1 to 10	49%	51%	22%	15%
11 to 20	48%	52%	19%	17%
21 to 30	55%	45%	25%	16%
31 to 40	49%	51%	21%	12%
41 to 50	48%	52%	15%	24%
Total Judging Years				
1 to 10	48%	52%	21%	14%
11 to 20	49%	51%	20%	15%
21 to 30	54%	46%	25%	18%
31 to 40	47%	53%	18%	14%
41 to 50	49%	51%	22%	14%
51 to 60	44%	56%	13%	13%
Active or Senior				
Active	51%	49%	22%	15%
Senior	48%	52%	19%	17%

APPENDIX C
BRIEF CHARACTERISTICS

	Cases Cited in Briefs	Sticky Cites per Brief	Percent Sticky	Average Cases Cited per Brief
Opening Briefs (325)	11168	2156	19%	34.4
Response Briefs (325)	13166	2759	21%	40.5
Reply Briefs (295)	5827	1245	21%	19.8
Opening and Reply Briefs	14659	2563	17%	——
Cited in Both Parties' Briefs	4276	1613	38%	——
From Response Brief to Reply Brief	1105	283	26%	___
From One Party's Brief to Opinion	23536	3699	16%	——
From Both Parties' Briefs to Opinion	23536	1613	7%	——

APPENDIX D
WINNINGNESS

	Cases Cited in Briefs	Cases Cited in Opinions	Sticky Cites	Super-Sticky Cites	Average Cites per Opinion
Appellant Won (86)	5807	1805	861	421	21.0
Appellee Won (201)	14021	4431	2287	977	22.0

	Citations in Opinions		Citations in Briefs
	Percent Sticky	**Percent Super Sticky**	**Percent Sticky**
Appellant Won	48%	23%	15%
Appellee Won	52%	22%	16%

	Cases Cited in Briefs	Sticky Cites (Out of 6236 Cases Cited in All Opinions)	Percent Sticky Cites in Opinions	Percent Sticky Cites in Briefs
Winning Party (297)	11937	2530	41%	21%
Losing Party (297)	11786	2026	32%	17%