

Understanding the Realities of Modern Patent Litigation¹

John R. Allison,^{*} Mark A. Lemley^{**}
& David L. Schwartz^{***}

Sixteen years ago, two of us published the first detailed empirical look at patent litigation. In this Article, we update and expand the earlier study with a new hand-coded data set. We evaluate all substantive decisions rendered by any court in every patent case filed in 2008 and 2009—decisions made between 2009 and 2013. We consider not just patent validity but also infringement and unenforceability. Moreover, we relate the outcomes of those cases to a host of variables, including variables related to the parties, the patents, and the courts in which those cases were litigated. The result is a comprehensive picture of the outcomes of modern patent litigation, one that confirms conventional wisdom in some respects but upends it in others. In particular, we find a surprising amount of continuity in the basic outcomes of patent lawsuits over the past twenty years, despite rather dramatic changes in who brought patent suits during that time.

Sixteen years ago, two of us published the first detailed empirical look at patent litigation.² That study provided a wealth of valuable information about patent-validity litigation, including the discovery that nearly half of all patents litigated to judgment were held invalid.³ But it was also limited in various respects. The study was based only on patent-validity decisions that finally resolved the case on the merits and only on reported decisions

1. ©2014 John R. Allison, Mark A. Lemley & David L. Schwartz. We thank David Abrams, Jeremy Bock, Shari Diamond, Stuart Graham, Rose Hagan, Jay Kesan, Naomi Lamoreaux, Arti Rai, Michael Risch, Ted Sichelman, and participants at the *Texas Law Review* Symposium on “Steps Toward Evidence-Based IP,” the Works-in-Progress IP Conference at Santa Clara, and the IP2 conference at the Hoover Institution for comments and suggestions. We would like to especially thank Fang Tang for his help with data analysis and statistics and Andrew Thompson for his research assistance.

* Mary John and Ralph Spence Centennial Professor of Business Administration, McCombs Graduate School of Business, The University of Texas at Austin.

** William H. Neukom Professor of Law, Stanford Law School; Partner, Durie Tangri LLP.

*** Associate Professor of Law and Co-Director, Center for Empirical Studies of Intellectual Property, Chicago-Kent College of Law.

2. See John R. Allison & Mark A. Lemley, *Empirical Evidence on the Validity of Litigated Patents*, 26 AIPLA Q.J. 185, 188–93 (1998) (discussing the existing literature and “lack of empirical evidence on the function and impact of the patent system” at the time).

3. *Id.* at 205.

available in a particular legal reporter.⁴ The latter limitation meant that almost half of the decisions were appellate.⁵ Importantly, the cases serving as the study's data sources are now on average more than twenty years old.⁶

In this Article we update and expand the earlier study with a new hand-coded data set. We evaluate all substantive decisions rendered by any court in every patent case filed in 2008 and 2009—decisions made between 2009 and 2013. We consider not just patent validity but also infringement and unenforceability. Moreover, we relate the outcomes of those cases to a host of variables, including variables related to the parties, the patents, and the court in which the case was litigated. The result is a comprehensive picture of the outcomes of modern patent litigation, one that confirms conventional wisdom in some respects but upends it in others.

In Part I, we discuss previous efforts to evaluate patent litigation empirically. In Part II, we discuss our methodology and the choices we made in study design. We present our results in Part III.

I. The Prior Art: Efforts to Understand Patent Litigation So Far

A number of scholars have empirically studied specific patent law doctrines. Claim construction is the most common, with most articles focusing on appellate cases.⁷ Obviousness has also been a point of

4. *See id.* at 194 (“These cases represent all written, final validity decisions by either district courts or the Federal Circuit reported in the *U.S.P.Q.* during an almost eight-year period . . .”).

5. *Id.* at 240.

6. *See id.* at 194 (studying cases decided between 1989 and 1996).

7. *E.g.*, J. Jonas Anderson & Peter S. Menell, *Informal Deference: A Historical, Empirical, and Normative Analysis of Patent Claim Construction*, 108 NW. U. L. REV. 1, 6 (2014) (analyzing claim construction from the Federal Circuit between 2000 and 2011); Gretchen Ann Bender, *Uncertainty and Unpredictability in Patent Litigation: The Time is Ripe for a Consistent Claim Construction Methodology*, 8 J. INTELL. PROP. L. 175, 175 (2001) (noting that the Federal Circuit frequently changes the trial court's claim construction); Christian A. Chu, *Empirical Analysis of the Federal Circuit's Claim Construction Trends*, 16 BERKELEY TECH. L.J. 1075, 1078–79 (2001) (examining the effects of the Federal Circuit's de novo review of claim construction); Shawn P. Miller, *“Fuzzy” Software Patent Boundaries and High Claim Construction Reversal Rates*, 17 STAN. TECH. L. REV. (forthcoming 2014) (manuscript at 2), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2139146 (determining that the Federal Circuit is more likely to find error in district court constructions of software patents); Kimberly A. Moore, *Are District Court Judges Equipped to Resolve Patent Cases?*, 15 HARV. J.L. & TECH. 1, 2 (2001) [hereinafter Moore, *Equipped*] (presenting study results “that show[] that district court judges improperly construe patent claim terms in 33% of the cases appealed to the Federal Circuit”); Kimberly A. Moore, *Markman Eight Years Later: Is Claim Construction More Predictable?*, 9 LEWIS & CLARK L. REV. 231, 245–46 (2005) [hereinafter Moore, *Markman*] (investigating the Federal Circuit's response to *Markman* and finding a higher reversal rate than in her 2001 study); David L. Schwartz, *Courting Specialization: An Empirical Study of Claim Construction Comparing Patent Litigation Before Federal District Courts and the International Trade Commission*, 50 WM. & MARY L. REV. 1699, 1702–04 (2009) [hereinafter Schwartz, *Courting Specialization*] (examining the performance of the U.S. International Trade Commission in patent construction cases); David L. Schwartz, *Practice Makes Perfect? An Empirical Study of*

scholarly interest,⁸ as have inequitable conduct⁹ and the doctrine of equivalents.¹⁰ Although the empirical research into patent litigation just mentioned has largely focused on Federal Circuit decisions,¹¹ with research employing data from district court litigation typically having been limited to opinions available on Westlaw or Lexis,¹² there have been a few notable exceptions.¹³ Unreported decisions, especially denials of summary judgment and the results of jury trials, are lacking from these data sets.¹⁴

Claim Construction Reversal Rates in Patent Cases, 107 MICH. L. REV. 223, 224–25 (2008) [hereinafter Schwartz, *Practice Makes Perfect?*] (questioning whether U.S. district court judges improve their patent claim construction decisions with experience); David L. Schwartz, *Pre-Markman Reversal Rates*, 43 LOY. L.A. L. REV. 1073, 1075–76 (2010) [hereinafter Schwartz, *Pre-Markman Reversal Rates*] (providing data from almost two decades of Federal Circuit opinions to investigate if changes in procedure changed reversal rates).

8. *E.g.*, Christopher A. Cotropia, *Nonobviousness and the Federal Circuit: An Empirical Analysis of Recent Case Law*, 82 NOTRE DAME L. REV. 911, 914 (2007) (studying cases decided under the nonobviousness requirement); Lee Petherbridge & R. Polk Wagner, *The Federal Circuit and Patentability: An Empirical Assessment of the Law of Obviousness*, 85 TEXAS L. REV. 2051, 2054–56 (2007) (arguing on the basis of data that the Federal Circuit’s doctrine of obviousness “appears relatively stable and increasingly flexible”); Jason Rantanen, *The Federal Circuit’s New Obviousness Jurisprudence: An Empirical Study*, 16 STAN. TECH. L. REV. 709, 710–11 (2013) (assessing two prior predictions about obviousness).

9. *E.g.*, Christian E. Mammen, *Controlling the “Plague”: Reforming the Doctrine of Inequitable Conduct*, 24 BERKELEY TECH. L.J. 1329, 1331–33 (2009) (reporting data on inequitable conduct allegations and advocating reforms to the doctrine of inequitable conduct); Lee Petherbridge, Jason Rantanen & Ali Mojibi, *The Federal Circuit and Inequitable Conduct: An Empirical Assessment*, 84 S. CAL. L. REV. 1293, 1318–19 (2011) (observing that the Federal Circuit’s doctrine of inequitable conduct is stricter than the doctrine as applied by lower courts and other judicial panels); Robert D. Swanson, Comment, *The Exergen and Therasense Effects*, 66 STAN. L. REV. 695, 717–18 (2014) (finding that inequitable conduct allegations have dropped dramatically in recent years).

10. *E.g.*, John R. Allison & Mark A. Lemley, *The (Unnoticed) Demise of the Doctrine of Equivalents*, 59 STAN. L. REV. 955, 957–58 (2007) (studying the doctrine of equivalents over three time periods and finding that it rarely matters anymore); Lee Petherbridge, *On the Decline of the Doctrine of Equivalents*, 31 CARDOZO L. REV. 1371, 1379 (2010) (confirming the Allison–Lemley findings); David L. Schwartz, *Explaining the Demise of the Doctrine of Equivalents*, 26 BERKELEY TECH. L.J. 1157, 1159 (2011) (arguing that the decline in the doctrine of equivalents resulted from “doctrinal reallocation” and “doctrinal displacement”).

11. *See, e.g.*, Chu, *supra* note 7, at 1092 (examining Federal Circuit decisions); Petherbridge et al., *supra* note 9, at 1305 (same); Petherbridge & Wagner, *supra* note 8, at 2071 (same).

12. *See, e.g.*, Moore, *Equipped*, *supra* note 7, at 8 n.36 (indicating that the database of district court opinions was compiled by running Westlaw searches).

13. Several studies have used comprehensive data from U.S. district courts rather than depending on published opinions from Westlaw or Lexis. *See, e.g.*, John R. Allison, Mark A. Lemley & Joshua Walker, *Extreme Value or Trolls on Top? The Characteristics of the Most-Litigated Patents*, 158 U. PA. L. REV. 1, 3 & n.3 (2009) [hereinafter Allison et al., *Most-Litigated Patents*] (using litigation data from Lex Machina); John R. Allison, Emerson H. Tiller, Samantha Zyontz & Tristan Bligh, *Patent Litigation and the Internet*, 2012 STAN. TECH. L. REV., art. 3, ¶¶ 10–11 (2012) [hereinafter Allison et al., *Patent Litigation*] (using litigation data from Derwent LitAlert database and Lex Machina); John R. Allison, Mark A. Lemley & Joshua Walker, *Patent Quality and Settlement Among Repeat Patent Litigants*, 99 GEO. L.J. 677, 682 & n.21 (2011) [hereinafter Allison et al., *Patent Quality*] (using litigation data from Lex Machina); John R.

Recently, the underlying documents, including motions and opinions, from district court litigation became more readily available. Electronic filing requirements meant that the online filing tool, Public Access to Court Electronic Records (PACER), has a nearly complete collection of litigation documents from patent cases.¹⁵ Some scholars have taken advantage of PACER data to analyze district court decisions.¹⁶ But the raw data provided by the Administrative Office of the United States Courts is notoriously error-prone,¹⁷ and it does a poor job of classifying outcomes.¹⁸

II. Our Methodology

In this Part, we explain in detail the techniques we used to locate and collect the data. We describe the data sources and provide information about the coders. And we describe our process of selecting data for inclusion in the data set.

A. Data Collection

We used the Lex Machina database as our data source.¹⁹ Lex Machina provides convenient access to cleaned and verified PACER data for district court patent litigation, which permitted us to evaluate all patent lawsuits. Lex Machina data offer three primary benefits. First, it includes all lawsuits, even those without a decision available on Westlaw or Lexis, so

Allison, Mark A. Lemley, Kimberly A. Moore & R. Derek Trunkey, *Valuable Patents*, 92 GEO. L.J. 435, 443–45 (2004) [hereinafter Allison et al., *Valuable Patents*] (using litigation data from local court records of individual U.S. district courts).

14. See, e.g., Paul M. Janicke & LiLan Ren, *Who Wins Patent Infringement Cases?*, 34 AIPLA Q.J. 1, 10 (2006) (excluding cases that were finally resolved at the trial court level and not appealed). A notable exception is Moore, *Equipped*, *supra* note 7, at 8, which is limited to appellate decisions but does evaluate both unpublished decisions and even one-word Rule 36 affirmances. Schwartz, *Practice Makes Perfect?*, *supra* note 7, at 238, also evaluates both unpublished decisions and Rule 36 affirmances.

15. For a discussion of PACER coding and its shortcomings, see generally Matthew Sag, *Empirical Studies of Copyright Litigation: Nature of Suit Coding* (Loyola Univ. Chi. Sch. of Law, Pub. Law & Legal Theory, Research Paper No. 2013-017), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2330256.

16. See, e.g., Jay P. Kesan & Gwendolyn G. Ball, *How Are Patent Cases Resolved? An Empirical Examination of the Adjudication and Settlement of Patent Disputes*, 84 WASH. U. L. REV. 237, 261 (2006) (examining the online docket reports available through the PACER system).

17. See *id.* at 261 tbl.1 (finding a substantial percentage of cases misclassified as patent cases); Kimberly A. Moore, *Judges, Juries, and Patent Cases—An Empirical Peek Inside the Black Box*, 99 MICH. L. REV. 365, 381 (2000) (eliminating some cases misclassified as patent trials from the data set).

18. See Kesan & Ball, *supra* note 16, at 265 (explaining that the Administrative Office of the District Courts' categories for case disposition are "rather ambiguous").

19. LEX MACHINA, <http://www.lexmachina.com>.

we do not overcount appellate decisions.²⁰ Second, Lex Machina has cleaned and evaluated the PACER data, eliminating many of the errors in the raw data.²¹ Finally, Lex Machina has indexed the cases to identify all summary judgment rulings, trial events, and appeals.²²

Our study covers all patent lawsuits filed in a federal district court between January 1, 2008 and December 31, 2009. We selected 2008 and 2009 for several reasons. First, those years are sufficiently recent to provide a snapshot of current patent litigation. Second, because the cases were initiated several years ago, the overwhelming majority of those cases were finally resolved or settled before our project began.²³ Lex Machina graciously provided us with a list of 2008 and 2009 lawsuits that contained at least one ruling on summary judgment or trial. Lex Machina furnished us a second list of 2008 and 2009 lawsuits, the second list including cases with an appeal but without a summary judgment ruling or trial. The second list allowed us to capture cases in which the parties stipulated to judgment based upon a claim construction decision with the goal of placing the case in condition for appeal. Both lists provided by Lex Machina included basic information about each lawsuit, including the judicial district in which the case was filed, the identity of the district court judge, and the filing date of the lawsuit.

From the cases provided by Lex Machina, we excluded lawsuits that did not include a complaint for infringement of a utility patent, or declaratory relief of noninfringement or invalidity of a utility patent. Thus, we excluded inventorship and licensing disputes, malpractice actions, and allegations of design or plant patent infringement. After removing these lawsuits, we reviewed the docket report in detail, reading all relevant orders, opinions, motions, verdicts, appellate rulings, and other necessary court documents to code the litigation outcomes.

Because many of the dockets were extremely complicated—it was not uncommon for a patent case to have over 500 docket entries—we felt that student coders would be ill-suited to the task. Coding of outcomes,

20. See *Features*, LEX MACHINA, <https://lexmachina.com/features/> (“[V]iew all patent case outcomes for a specific judge or district, displayed in easy-to-read charts and graphs supported by interactive case lists.”).

21. See *How It Works*, LEX MACHINA, <https://lexmachina.com/features/how-it-works/> (“Lex Machina cleans, codes, and tags all data . . .”).

22. See *id.* (“We identify all asserted patents, findings, and outcomes, including any damages awarded. We also build a detailed timeline linking all the briefs, motions, orders, opinions, and other filings for every case.”).

23. We conducted the coding in the late summer and fall of 2013. By February 2014, it appears that only 2%–3% of 2008 and 2009 cases were still open. See Dennis Crouch, *Pendency of Patent Infringement Litigation*, PATENTLY-O (Feb. 17, 2014), <http://patentlyo.com/patent/2014/02/pendency-infringement-litigation.html>; see also Kesan & Ball, *supra* note 16, at 246 (defending the decision to study cases by year filed rather than by year terminated).

especially in patent cases, is notoriously difficult and time consuming, requiring deep knowledge of patent law and litigation, and the motivation to devote long hours to the task. Consequently, Lemley and Schwartz each personally coded the litigation-outcome information for approximately half of the lawsuits. Both Lemley and Schwartz are experienced patent litigators who understand how to read a docket and appreciate complex litigation rulings. The hand coding was extremely time intensive; it took several hundred hours in the aggregate. To permit an evaluation of the reliability and consistency of the coding, Lemley and Schwartz also overlapped in their coding of approximately ten percent of the lawsuits.²⁴

Our study uses a patent–case combination as the unit of analysis. For each case, we coded the outcome separately for each asserted patent. For instance, if the jury returned a verdict on two patents, then we recorded separately what occurred for each patent.²⁵ For each patent, we also obtained various patent demographic information from Thomson Innovation Solutions, including citations received (or “forward citations”),²⁶

24. Lemley and Schwartz both initially coded approximately 5% of the cases. Thereafter, they compared results and fine-tuned the codebook. For coding of the remaining cases, Lemley and Schwartz overlapped in 10% of the initial list of cases provided by Lex Machina. Some of the cases provided by Lex Machina turned out to not have relevant merits decisions. After a manual review of the dockets, the 10% overlap resulted in 30 patent–cases with duplicate coding. To increase the amount of overlap and permit the use of statistical tests to report inter-coder reliability, Schwartz additionally coded another random 15% overlap with Lemley, for an additional 46 patent–cases with duplicate coding. We chose “Cohen’s Kappa” (*kappa*) as the measure of inter-coder reliability. Mark A. Hall & Ronald F. Wright, *Systematic Content Analysis Of Judicial Opinions*, 96 CALIF. L. REV. 63, 113–14 (2008) (stating that the best practice for evaluating coding reliability is to report an agreement coefficient, such as *kappa*). *Kappa* ranges from 0 to 1, with numbers near 1 indicating a higher degree of reliability. *See id.* (explaining that a 0 indicates “agreement entirely by chance” and a 1 indicates “perfect agreement”). For the basic definitive and interim winners in cases, *kappa* was 0.9534, equating to near perfect agreement. For grants of motions for summary judgment of invalidity and noninfringement, *kappa* was 0.9793, which also equates to near perfect agreement for times in which we both identified motions. However, one of us found 1 additional motion for summary judgment of invalidity (40 v. 39). For motions for summary judgment of noninfringement, we each identified motions that the other did not (42 motions were found by both authors; one found 43 motions, while the other identified 44 motions). We revisited the overlapping case dockets again to understand these additional rulings, and we found that the additionally identified rulings should be included. We corrected all known disagreements in the data set. We believe that these differences in coding are due to the complexity of the dockets, and we do not believe that they are biased in one direction or another. We do believe, however, that the reliability information suggests that we slightly undercounted the numbers of merits rulings, although we cannot be sure whether the actual number should have more denials or grants.

25. Occasionally, the court ruled differently on different claims of a patent. For instance, claim 1 may be infringed and not invalid, but claim 2 was not infringed and anticipated. In these cases, we would create a new record for each group of claims that had a different substantive outcome.

26. Because the number of citations received by a patent in later patents (that treat the earlier patent as prior art) is a moving target, the raw numbers of citations received must be adjusted to account for the varying ages of the patents in our data set. We used a commonly employed

each type of prior art reference, maintenance status, number of claims, number of inventors, geographic location of the inventors, and the assignee when there was one. Allison manually coded for whether the patented invention had a U.S. or foreign origin using a decision model that was based on the domicile of a majority of the inventors and resorting to the domicile of the assignee as a tie breaker in the unusual case in which this was required.²⁷ We calculated the age of the patents in our data set as of the filing of the current litigation in 2008–2009. Moreover, we also located the first lawsuit in which each patent had been asserted. From the first lawsuit information, we calculated the age of the patent at first lawsuit. We also determined the age of the patents as of the filing of the 2008–2009 lawsuit.

For each patent in a lawsuit, the coders reviewed and captured all rulings on summary judgment relating to a patent law issue. This includes rulings on motions of summary judgment of noninfringement, infringement, validity, invalidity, inequitable conduct, and no inequitable conduct. We excluded rulings on issues that were not patent-specific, such as laches. We also excluded summary judgment rulings on patent law issues if the court did not reach the merits of the issue—such as denials of summary judgment motions—as being premature. The coders also reviewed and recorded all trial outcomes, whether there was a jury or bench trial, and decisions on post-verdict JMOL motions. Finally, we recorded whether an appeal was lodged and how the appeal was resolved. The resolution data includes whether the ruling on the patent was affirmed or reversed on appeal, or whether an appeal is pending or was dismissed (typically because the case settled). When the underlying trial or appellate court opinion lacked sufficient detail to ascertain the basis for the ruling, we read the underlying briefing by the parties.

We coded merits decisions at a low level of granularity. For invalidity, we coded whether the ruling was based on utility, patentable subject matter, section 102 prior art, obviousness, indefiniteness, written

technique suggested by Bronwyn Hall and her colleagues. Bronwyn H. Hall, Adam B. Jaffe & Manuel Trajtenberg, *The NBER Patent-Citations Data File: Lessons, Insights, and Methodological Tools*, in PATENTS, CITATIONS, & INNOVATIONS 403, 434–41 (Adam B. Jaffe & Manuel Trajtenberg eds., 2002). The method of adjustment to account for the different ages of patents involves placing each patent in the data set into a cohort of other patents in the data set that were issued during the same year. *Id.* at 437. Thus, each cohort is one year, although cohorts of more than one year could be used if necessary even though that would decrease precision somewhat. The number of forward citations received by each patent is divided by the average number of forward citations received by other patents in the same cohort. *Id.* This gives us the adjusted number of forward citations for that patent in the data set. The process is repeated for every other patent in the same cohort and then repeated for each patent in the other year cohorts. To obtain the adjusted number of forward citations for an entire data set, we then averaged the quantity of adjusted number of forward citations received by all patents in the set.

27. Allison also hand-coded the technology and industry categories for each patent. We report those results in a companion paper.

description, enablement, and best mode. We also coded various bases for section 102 invalidity. For infringement, we captured literal infringement, the doctrine of equivalents, and various types of indirect infringement. And we coded unenforceability as well as the basis for the unenforceability argument. In addition to the separate coding of issues for summary judgment and trial, we also recorded the final resolution for each patent on the issues of infringement, validity, and enforceability.

Notably, we coded the issues litigated to decision, whether or not that decision resulted in a trial outcome or a grant of summary judgment. Thus, if an accused infringer argued that the patent was invalid for lack of patentable subject matter, anticipation, and obviousness, and the court denied the first two arguments but granted the third, each of those three rulings shows up in our data set. To understand how the final resolution variables were coded, one should understand that denial of summary judgment does not result in a final resolution. Instead, denial of summary judgment means that there is an unresolved disputed issue of material fact.²⁸ Consequently, denials of summary judgment alone would not result in a final ruling in either direction. If, however, the issue had been resolved at trial, then the final ruling was coded as the trial resolution. If summary judgment had been granted on an issue, then the summary judgment ruling was coded as the final resolution in our coding.²⁹ We coded decisions that finally ruled for a party on an issue as definitive wins, and decisions that ruled for a party but kept the issue alive (largely denial of summary judgment but also remands on appeal) as interim wins.

B. *Potential Limitations*

Our data set and the implications that can be drawn therefrom are subject to several limitations. For brevity, we discuss two important limitations here.

First, our data set is limited to lawsuits filed in 2008 and 2009. It is sufficiently recent, in our opinion, that the results are generally applicable today. However, there have been several legal changes in the interim that may make lawsuits today different from those in our data set. The most salient changes are the passage of the America Invents Act in 2011;³⁰ the Federal Circuit's en banc *Therasense*³¹ decision in 2011; and three Supreme

28. See FED. R. CIV. P. 56.

29. Of course, if the Federal Circuit reversed a ruling relating to a patent on appeal, we updated the final-resolution coding to reflect the appellate decision.

30. Leahy-Smith America Invents Act, Pub. L. No. 112-29, 125 Stat. 284 (2011) (codified in scattered sections of 35 U.S.C.).

31. *Therasense, Inc. v. Becton, Dickinson & Co.*, 649 F.3d 1276 (Fed. Cir. 2011) (en banc).

Court cases involving the doctrine of patentable subject matter in 2010,³² 2012,³³ and 2013.³⁴ The Federal Circuit issued several opinions involving patent damages, which may have affected litigant behavior and settlement.³⁵ These law changes may influence what issues litigants press, and separately, which cases reach the stage of a ruling on the merits. Accordingly, the cases filed today in 2014 may differ from those we studied.

Second and perhaps more importantly, our data set only contains patents which were subject to a ruling on summary judgment, a trial, or an appeal. To be sure, we have the population of cases that resulted in a ruling on a dispositive motion or trial. For these cases, we report statistical results on the outcomes. However, most lawsuits settle,³⁶ and as our data confirms, most lawsuits settle before any ruling on the merits. Cases that settled before any substantive patent ruling are completely absent from our data set, with the exception of some basic descriptive statistics reported in Table 1. Moreover, many disputes do not result in litigation.³⁷ Obviously, our data set lacks unlitigated disputes about patents. The upshot is that our data and results are not generalizable to the cases or disputes that settled without any substantive ruling. Thus, while our data sheds light on who wins and loses patent cases and dispositive motions, it cannot tell us who *would* win cases that were filed but settled without a judgment.

We do not even have a sense of which direction the bias, if any, would point if one were interested in all litigated cases. On the one hand, it may be that the cases that are settled before a merits ruling are mainly strong cases in which the parties overlapped in their expectations on success. If this is true, then the defendant win rates we observe in our data set would be higher than the win rate if all cases were litigated to judgment. On the other hand, it could be that the cases that settled before a merits ruling consist disproportionately of meritless cases that were resolved via cost-of-defense settlements.³⁸ If this alternative hypothesis was true, then our

32. *Bilski v. Kappos*, 130 S. Ct. 3218 (2010).

33. *Mayo Collaborative Servs. v. Prometheus Labs., Inc.*, 132 S. Ct. 1289 (2012).

34. *Ass'n for Molecular Pathology v. Myriad Genetics, Inc.*, 133 S. Ct. 2107 (2013).

35. *See, e.g., Uniloc USA, Inc. v. Microsoft Corp.*, 632 F.3d 1292, 1315 (Fed. Cir. 2011) (prohibiting the use of the 25% rule of thumb for calculating reasonable royalties); *ResQNet.com, Inc. v. Lansa, Inc.*, 594 F.3d 860, 873 (Fed. Cir. 2010) (vacating the district court's damages award because the reasonable royalty determination relied on speculative evidence).

36. Mark A. Lemley, *Rational Ignorance at the Patent Office*, 95 NW. U. L. REV. 1495, 1501 (2001) ("The overwhelming majority of [patent] lawsuits settle or are abandoned before trial.").

37. *See id.* at 1507 (estimating that only 1.5% of patents were litigated).

38. Such claims may be common. *See* Mark A. Lemley & A. Douglas Melamed, *Missing the Forest for the Trolls*, 113 COLUM. L. REV. 2117, 2163 (2013) (noting that patent trolls pursue a large number of cases, many of which a practicing entity would probably not bring, but that these cases are more likely to settle quickly).

estimates of defendant win rates from the cases that reached the merits phase would be lower than the defendant win rate if all filed cases went to judgment. Because almost all of the settlements are confidential,³⁹ we cannot assess the direction of the bias. For these reasons, we urge readers to interpret our results with these limitations in mind.

III. Our Results

In this Part, we present some basic descriptive statistics and then draw some lessons from the data.

A. *Description of the Patents and Cases*

As of the date of our study, there were 949 merits decisions on patents based on infringement lawsuits filed in 2008 and 2009.⁴⁰ Those decisions were made in 462 different cases involving 777 different patents. Most of those cases were concentrated in a relatively small number of judicial districts. Leading the way were the Eastern District of Texas and the District of Delaware, two districts perennially favored by plaintiffs.⁴¹ We present the data in Table 1, along with data on where all of the roughly 5,000 lawsuits filed in 2008 and 2009 were filed.

Of the 949 merits decisions, 636 were definitive wins on an issue for one side or the other; the remainder were interim wins (usually the denial of the other side's summary judgment motion). The most common occasions for a merits ruling were summary judgment motions of invalidity (430 observations) and noninfringement (473 observations, increasing to 509 when we added stipulated judgments of noninfringement after claim construction). By contrast, patentees were less likely to seek and obtain a ruling in their favor on summary judgment. Patentees brought and received a ruling on only 125 summary judgment motions on validity⁴² and 128

39. See Scott A. Moss, *Illuminating Secrecy: A New Economic Analysis of Confidential Settlements*, 105 MICH. L. REV. 867, 869 (2007) ("Public settlements are the exception, common in only a few types of cases . . .").

40. We cut off our data collection on June 1, 2013.

41. See Mark A. Lemley, *Where to File Your Patent Case*, 38 AIPLA Q.J. 401, 405 & tbl.2 (2010) (finding both districts to be among the most favored for patent lawsuits). Note, however, that because we count only cases with merits decisions, rather than all cases filed, a district's share of cases in our data set may not match their share of filed cases because cases in some districts are more likely to settle than others.

42. Summary judgment of validity differed from the other summary judgment motions we classified. A motion for summary judgment of validity often encompassed one ground for invalidity. For instance, the patent holder may move for summary judgment of no anticipation. Even if the motion was granted, it would not preclude an accused infringer from contesting the validity on a different basis, such as lack of enablement. Thus, even a successful patent holder on a motion for summary judgment of validity did not necessarily prevail on all invalidity defenses.

summary judgment motions on infringement. They also brought 116 summary judgment motions of no inequitable conduct which resulted in a ruling on the merits. Accused infringers only brought 24 summary judgment motions of inequitable conduct, and none was successful.

Of our 949 merits observations, 290 patents went to trial. Over 70% (206 patents) were heard by juries, with the remainder (84) decided in bench trials. A total of 273 of the 949 merits decisions reached a Federal Circuit decision on appeal, though another 126 merits decisions were appealed and then settled before decision. There are presently 82 merits decisions pending before the Federal Circuit.

Table 1: 2008–2009 Patent Lawsuit Filings

District	% of Merits Decisions in Our Database (#)	% of Lawsuits in Our Database (#)	% of 2008–2009 Lawsuit Filings (#)
TXED	13.5% (128)	13.0% (60)	10.4% (524)
DED	12.9% (122)	10.6% (49)	7.8% (394)
CAND	8.5% (81)	7.4% (34)	6.5% (325)
CACD	5.9% (56)	7.6% (35)	9.0% (454)
CASD	5.3% (51)	5.0% (23)	2.7% (138)
NYSD	5.0% (47)	3.2% (15)	4.3% (216)
ILND	4.2% (40)	4.1% (19)	5.5% (275)
NJD	3.6% (34)	5.0% (23)	6.0% (302)
WIWD	3.2% (30)	3.2% (15)	1.3% (65)
VAED	3.2% (30)	2.6% (12)	2.2% (112)
MAD	2.8% (27)	4.8% (22)	2.2% (108)
TXSD	2.3% (22)	1.5% (7)	1.3% (67)
OHND	1.8% (17)	1.9% (9)	1.8% (89)
All Other Districts	27.8% (264)	30.1% (139)	39.0% (1960)

For summary judgment of invalidity, noninfringement, infringement, inequitable conduct, and no inequitable conduct, the winner of the motion completely resolved the issue in the case.

The columns in Table 1 require some interpretation before being compared with each other. The second column from the left, providing the percentage of merits decisions, is done on a per patent–case basis, as our data is broken down in this manner. A single case may involve multiple patents. The third column from the left collapses our data on merits decisions into a per lawsuit basis, which permits easier comparison with the data on raw lawsuit filings. The far right column utilizes Lex Machina’s raw data on case filings, which is done on a per case basis. While the patent–case and case bases differ, a comparison is useful to see basic trends.


First, less than 10% of the patent lawsuits filed in 2008 and 2009 (462 of 5,029) resulted in any merits decision.⁴³ In other words, more than 90% of lawsuits settle before the court resolves summary judgment or tries the case.⁴⁴

Second, as shown graphically in Figure 1 below, the identity of the districts with the most merits decisions loosely tracks the identity of the districts with the most filings.

43. The percentage is slightly understated because some of the filed lawsuits are still pending and may reach a merits decision after the date of our coding. But there is reason to believe that is true of no more than 2%–3% of cases filed in 2008 and 2009. *See* Crouch, *supra* note 23 (finding that approximately 97% of lawsuits filed in 2008 and 2009 were terminated in district courts by 2013).

44. Kesan & Ball, *supra* note 16, at 271, showed that a relatively large percentage (7%–8%) of summary judgments were being granted in their data set.

Figure 1: 2008–2009 Patent Lawsuit Filings and Merits Decisions by District



However, some districts, such as the Eastern District of Texas, Southern District of California, and the Western District of Wisconsin, appear overrepresented in merits decisions relative to filings. The Western District of Wisconsin, for instance, is known as a “rocket docket,”⁴⁵ which may provide less time for the parties to settle. Other districts, such as the Central District of California, appear underrepresented. The Central District of California has a large number of district court judges—like the Northern District of Illinois, another venue underrepresented in merits decisions—and also has a long average case pendency.⁴⁶ Longer pendency may increase the possibility of settlement before a merits decision.

45. Saurabh Vishnubhakat, *Reconceiving the Patent Rocket Docket: An Empirical Study of Infringement Litigation 1985–2010*, 11 JOHN MARSHALL REV. INTELL. PROP. L. 58, 61 (2011).

46. The median case in the Central District of California went to trial in 955 days. LEX MACHINA, <https://law.lexmachina.com/court/cacd>. By comparison, the median case in the Western District of Wisconsin went to trial in 588 days. LEX MACHINA, <https://law.lexmachina.com/court/wiwd>.

B. *The Realities of Patent Litigation*

In this subpart, we draw a number of lessons from our results—both the descriptive statistics and our multivariate regression analysis.

1. *The Nature of Validity Challenges Is Changing.*—In our 1998 study, we found that decided validity challenges were overwhelmingly based on obviousness—so much so that even though obviousness challenges had one of the lowest win rates, they were also responsible for the largest number of judicial patent invalidations.⁴⁷ Prior art challenges were close behind.⁴⁸ And what section 112 challenges we found were almost entirely enablement- or best-mode-based.⁴⁹

Things have changed. While there are still a sizeable number of adjudicated obviousness challenges (149 summary judgment motions decided), there were fewer decisions on summary judgment motions of obviousness than of anticipation (154). There are a growing number of decisions based on patentable subject matter (26)—a category of minor importance in the 1998 study.⁵⁰ We suspect that if we reviewed lawsuits filed even more recently—such as those filed in 2010 and 2011—the number of summary judgment motions on patentable subject matter would have substantially increased. Recent Supreme Court and Federal Circuit case law likely encouraged more litigation on the doctrine.⁵¹ And the single largest category of adjudicated challenges was for indefiniteness (176), a validity doctrine that barely registered in the 1998 study.⁵²

We attribute the growth of indefiniteness challenges to two factors. First, a major portion of the decisions in our data set involve software patents,⁵³ and the Federal Circuit in the 2000s developed a doctrine that

47. Allison & Lemley, *supra* note 2, at 209 tbl.2.

48. *Id.*

49. *See id.* (finding 29 decisions invalidating patents on enablement, written description, and best mode grounds, while only 8 decisions invalidated patents on claim indefiniteness grounds).

50. *Id.*

51. *See, e.g., Mayo Collaborative Servs. v. Prometheus Labs., Inc.*, 132 S. Ct. 1289, 1295–97 (2012) (holding that a method of administering thiopurine drugs in the treatment of autoimmune diseases was excluded from patentability because the method essentially “set forth laws of nature” and did not do enough to add to natural processes so as to warrant patentability); *Bilski v. Kappos*, 130 S. Ct. 3218, 3231 (2010) (holding that a hedging strategy was based on too abstract a concept to be a patentable process); *CLS Bank Int’l v. Alice Corp. Pty.*, 717 F.3d 1269, 1274 (Fed. Cir. 2013) (en banc), *cert. granted*, 134 S. Ct. 734 (2013) (holding that the asserted method, computer-readable medium, and system claims of the defendant’s patents were invalid for failure to recite patent-eligible subject matter).

52. *See Allison & Lemley, supra* note 2, at 208 tbl.1 (finding that indefiniteness was the grounds for invalidity in only 8 cases).

53. Over one-third of the merits decisions in our study concerned software patents (339 of 949). We discuss technology- and industry-specific results in a subsequent paper.

applied indefiniteness to software means-plus-function claims with more force than elsewhere.⁵⁴ While that likely led to more indefiniteness challenges in software cases, those challenges apparently were not all that successful, perhaps because fewer and fewer claims are written in means plus function format. Second, indefiniteness is a pure question of law that is normally decided in connection with claim construction⁵⁵ because the defendant's argument is that the claim term is not capable of being construed.⁵⁶ Claim construction itself was rare in our 1998 paper.⁵⁷ That study only included data through 1996,⁵⁸ the same year *Markman*⁵⁹ was decided. Today, however, claim construction is the most likely form of substantive ruling in a patent case because it is a prerequisite to virtually any type of summary judgment motion on validity or infringement.⁶⁰

54. See *Function Media, L.L.C. v. Google Inc.*, 708 F.3d 1310, 1318 (Fed. Cir. 2013) (noting that “[w]hen dealing with a ‘special purpose computer-implemented means-plus-function limitation,’” disclosure of the algorithm for performing the function is required); *ePlus, Inc. v. Lawson Software, Inc.*, 700 F.3d 509, 518 (Fed. Cir. 2012) (concluding that the specification at issue did “not disclose sufficient structure for the ‘means for processing’ limitation”); *Noah Sys., Inc. v. Intuit Inc.*, 675 F.3d 1302, 1313 (Fed. Cir. 2012) (distinguishing between cases where a software patent specification “discloses no algorithm” and those where an algorithm is disclosed but still may be “inadequate”); *Ergo Licensing, LLC v. CareFusion 303, Inc.*, 673 F.3d 1361, 1365 (Fed. Cir. 2012) (explaining that “[i]t is only in the rare circumstances where any general-purpose computer without any special programming can perform the function that an algorithm need not be disclosed,” and requiring that such disclosure demonstrate the “step-by-step process” for arriving at a given result); *Typhoon Touch Techs., Inc. v. Dell, Inc.*, 659 F.3d 1376, 1384–86 (Fed. Cir. 2011) (holding that while means-plus-function software claims required disclosure of corresponding structure performing that function in the specification, that structure did not need to be described in the form of software code); *In re Aoyama*, 656 F.3d 1293, 1294, 1297–98 (Fed. Cir. 2011) (holding a means-plus-function software patent claim invalid as indefinite for failure to disclose the corresponding algorithm performing that function); *Aristocrat Techs. Austl. Pty Ltd. v. Int’l Game Tech.*, 521 F.3d 1328, 1337–38 (Fed. Cir. 2008) (requiring algorithm disclosure and indicating that the standard is whether “a person of ordinary skill in the art would not recognize the patent as disclosing any algorithm at all”); *WMS Gaming Inc. v. Int’l Game Tech.*, 184 F.3d 1339, 1349 (Fed. Cir. 1999) (“[T]he disclosed structure is not the general purpose computer, but rather the special purpose computer programmed to perform the disclosed algorithm.”). For further discussion on functional claiming, see generally Mark A. Lemley, *Software Patents and the Return of Functional Claiming*, 2013 WIS. L. REV. 905.

55. See, e.g., *Teva Pharms. USA, Inc. v. Sandoz, Inc.*, 723 F.3d 1363, 1368, 1373 (Fed. Cir. 2013), cert. granted, 134 S. Ct. 1761 (2014) (addressing questions of definiteness and claim construction and noting that both are questions of law).

56. See Peter S. Menell, Matthew D. Powers & Steven C. Carlson, *Patent Claim Construction: A Modern Synthesis and Structured Framework*, 25 BERKELEY TECH. L.J. 711, 772 (2010) (“When a claim cannot be construed, it is indefinite, and therefore invalid. Some authority suggests that all indefiniteness issues boil down to an issue of claim construction.”).

57. See Allison & Lemley, *supra* note 2, at 208 tbl.1.

58. *Id.* at 194.

59. *Markman v. Westview Instruments, Inc.*, 517 U.S. 370 (1996).

60. *Markman* requires courts to decide claim construction as a matter of law. *Id.* at 372. “This process is usually conducted during the pretrial stage in a ‘*Markman* hearing,’ where the judge determines the scope of the patent at issue.” Jay P. Kesan & Gwendolyn G. Ball, *Judicial*

Because courts often decide indefiniteness issues while construing claims, they are likely to see more indefiniteness motions than other forms of invalidity issues. Cases that settle after claim construction, for instance, never reach the merits of other arguments but will decide indefiniteness.⁶¹ Notably, however, software patents are not statistically significantly more likely to be found indefinite than others.⁶²

It appears that the indefiniteness doctrine plays a larger role than previously recognized in patent law. Remarkably, the rise of indefiniteness motions occurred despite Federal Circuit hostility to the doctrine. The Federal Circuit has made it very difficult to prevail on indefiniteness outside of software means-plus-function claims; a claim is indefinite under current law only if it is “insolubly ambiguous.”⁶³ The Supreme Court granted certiorari in January 2014 in an indefiniteness case⁶⁴ and seems poised to broaden the doctrine considerably. If it does, indefiniteness may play an even larger role in patent litigation in the near future.

2. *Individual Validity Challenges Lose.*—The courts ruled on validity in a large number of cases, mostly on summary judgment. Most of those motions failed. Table 2 reports the success rates of summary judgment motions of invalidity, both overall and by specific issue.

Overall, accused infringers won only 31% of the invalidity challenges brought on summary judgment.⁶⁵ For many of the most common sorts of challenges, the win rate was even lower. Patentees defeated summary judgment motions based on prior art, obviousness, and section 112 more

Experience and the Efficiency and Accuracy of Patent Adjudication: An Empirical Analysis of the Case for a Specialized Patent Trial Court, 24 HARV. J.L. & TECH. 393, 415 (2011). “If there are no remaining issues of material fact . . . a case can be resolved on summary judgment or quickly settled as the possible outcomes become more predictable.” *Id.* Thus, claim construction may well be the only substantive ruling in a given case.

61. Notably, the fact that indefiniteness is decided during claim construction means that we may actually undercount the number of indefiniteness motions. Not all indefiniteness motions or rulings are styled “summary judgment;” some rulings on indefiniteness may evade our view because they are buried inside an order that purports to be only about claim construction. While we have done our best to identify all such cases, we cannot guarantee that we have them all. So, if anything, our numbers understate the growth in the importance of indefiniteness.

62. We will discuss industry- and technology-specific results in a separate, forthcoming paper.

63. *Teva Pharms. USA, Inc. v. Sandoz, Inc.*, 723 F.3d 1363, 1368 (Fed. Cir. 2013), *cert. granted*, 134 S. Ct. 1761 (2014) (quoting *Biosig Instruments, Inc. v. Nautilus, Inc.*, 715 F.3d 891, 898 (Fed. Cir. 2013), *cert. granted*, 134 S. Ct. 896 (2014)).

64. *Nautilus, Inc. v. Biosig Instruments, Inc.*, 134 S. Ct. 896 (2014).

65. Unless otherwise noted in the paper, we treated split rulings as separate observations. *See supra* note 25. As a robustness check, we also calculated the summary judgment success rate on invalidity by reweighting these split rulings to normalize all observations on a patent in a case to one. Using this metric, the invalidity rate was still 31%.


than four times in five. Notably, patentable subject matter motions were the only ones to prevail a majority of the time (14 of 26, or 54%, were successful).

Table 2: Success Rates of Invalidity Summary Judgment Motions⁶⁶

Grounds for Summary Judgment	Number of Successful Motions/Total	Percentage of Successful Motions
No Patentable Subject Matter	14/26	54%
Section 102: Prior Art	31/154	20%
Section 103: Obviousness	31/149	20%
Section 112: Indefiniteness	30/176	17%
Section 112: Lack of Enablement	8/63	13%
Section 112: Inadequate Written Description	11/73	15%
Overall	131/430	30%

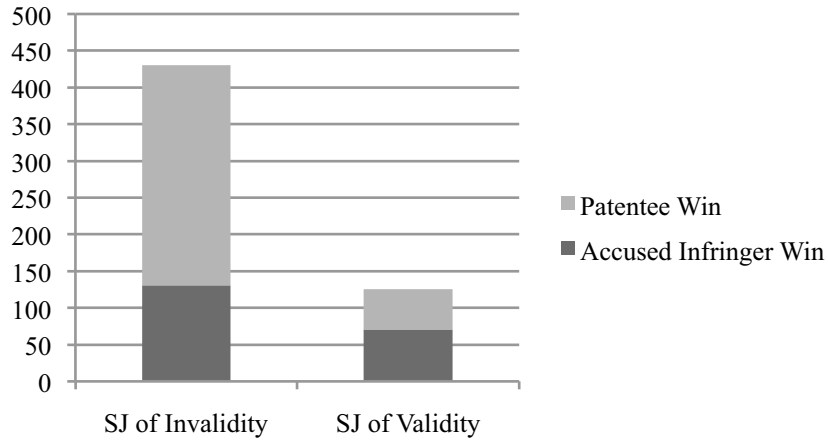
66. The numbers of individual challenges do not add to the total because some motions were brought on multiple grounds. The numbers of successful challenges do not add to the total because a few successful motions were brought on grounds not listed here, like utility or inventorship.

Figure 2: Success Rates on Invalidation Summary Judgment Motions



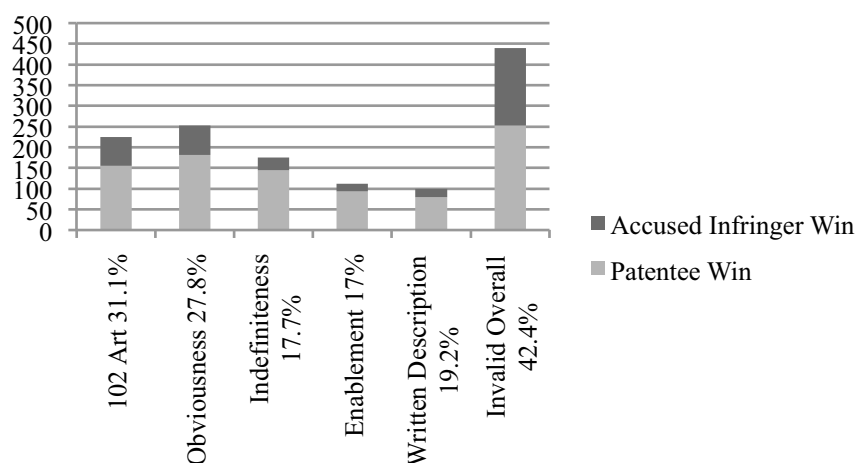
Patentees were much less likely to obtain summary judgment of validity, as Figure 3 shows.

Figure 3: Overall Results on Summary Judgment of Validity



The fact that most individual validity challenges fail is true not just of summary judgment rulings, but also of overall final decisions on validity. Figure 4 shows the overall win rate for validity across all procedural postures.


Figure 4: Overall Invalidity Win Rates



3. *Overall, Challengers Win.*—Notwithstanding our finding that most individual validity challenges fail, the overall picture for patentees is considerably darker. Patentees won only 164 of the 636 definitive merits rulings, or 26%.⁶⁷ Notably, that number is essentially unchanged from Paul Janicke and LiLan Ren’s study nearly a decade ago,⁶⁸ despite substantial changes in the nature of patent plaintiffs in that decade.


67. As a robustness check, we also calculated the definitive-merits-ruling win rate by reweighting the split patents. Using this alternative metric, patentees still only won 26% of the rulings.

68. See Janicke & Ren, *supra* note 14, at 5 (finding that patentees won 25% of cases).

Figure 5: Overall Patentee Win Rate

Why do patentees lose nearly three-quarters of the time when the court definitively resolves the merits? The answer is twofold. First, while courts turn away most validity challenges, patentees do not fare as well when it comes to infringement. Accused infringers won 54% (256 of 473) of their summary judgment motions alleging noninfringement of individual patents.⁶⁹ That number rises to 57% (292 of 509) when we include stipulated judgments of noninfringement after claim construction, which are functionally equivalent to summary judgments of noninfringement; the patentee concedes that it cannot win under a particular claim construction in order to tee the case up for appeal.

69. As a robustness check, we also calculated the summary judgment win rate on noninfringement by reweighting the split patents. Accused infringers won 53% of their summary judgment rulings using that metric.

Figure 6: Summary Judgments of Infringement

Second, the nature of patent litigation requires patentees to win every issue before the court. A patentee who defeats five of six invalidity challenges, only to lose the sixth, loses the case.⁷⁰ So does a patentee who wins on validity and inequitable conduct but loses on infringement.⁷¹ One of us has referred to this as the “fractioning” of patent law.⁷² Our data suggest that it has a significant effect on patent cases overall because many of our cases had motions on multiple issues, and those motions were not always decided in favor of the same party. In patent law, a split decision is almost always a decision for the accused infringer, not the patentee.

The summary judgment process exacerbates the fractioning. Summary judgment in most areas of law is predominantly used by defendants,⁷³ and patent law is no exception. Patentees brought fewer decided motions for

70. See Mark A. Lemley, *The Fractioning of Patent Law*, in *INTELLECTUAL PROPERTY AND THE COMMON LAW* 504, 509 (Shyamkrishna Balganeshe ed., 2013) (describing validity doctrines as having a “patentee-must-win-everything characteristic”).

71. See *id.* at 508 (noting that a patentee must win both invalidity and infringement).

72. *Id.* at 504.

73. Joe S. Cecil et al., *A Quarter-Century of Summary Judgment Practice in Six Federal District Courts*, 4 *J. EMPIRICAL LEGAL STUD.* 861, 886 (2007) (“Defendants’ motions for summary judgment are far more common than plaintiffs’ motions.”).

summary judgment of infringement (128) than accused infringers did seeking noninfringement (473). Accused infringers prevailed on 257 of the decided summary judgment motions of noninfringement, a success rate of 54%. When stipulated judgments of noninfringement are included, accused infringers received favorable pretrial judgments of noninfringement in 316 of 509 instances, a success rate of 62%.⁷⁴ By contrast, patentees won less than a third of their motions for summary judgment of infringement (41 of 128, or 32%). The patentee's burden to be entitled to summary judgment of infringement is higher than the burden on accused infringers for noninfringement. A patentee must show a lack of disputed issues of material fact for *all* elements of the claimed invention, while the accused infringers merely need to show a lack of disputed issues of material fact for *any* element of the claimed invention.⁷⁵

Furthermore, because the defendant only needs to prevail on one defense, it can move on one or more bases for summary judgment.⁷⁶ Even if unsuccessful, the accused infringer has another chance to win the case at trial.⁷⁷ In contrast, the patentee must both survive summary judgment and prevail at trial.⁷⁸ Thus, the accused infringers have several bites at the proverbial apple.

If a case reached the trial stage, patentees fared much better. Overall, patentees won 60.7% of the trials, which included prevailing on 59.4% of patents decided by juries and 63.9% of patents decided by the bench. And it bears repeating that many cases are settled after a denial of summary judgment and before trial. These patents are not included in our statistics on definitive rulings, and many presumably involve a monetary payment to the patentee. These selection issues should be taken into account when considering the statistic that accused infringers win approximately three quarters of the patents that end with a definitive ruling. Patentees often get paid even without a definitive ruling.

4. *Where You Stand Depends on Where You Sit.*—Both patentees and accused infringers engage in forum shopping, filing suit in the district court

74. A patent owner may sometimes stipulate that, if the court construes a disputed claim term in a particular way, there will be no infringement. If the court construes the claim term unfavorably to the patent owner, the result is a judgment of noninfringement in the same manner as though the decision had been in response to a motion for summary judgment of noninfringement.

75. Lemley, *supra* note 70, at 506.

76. *See id.*

77. *See* Morton Denlow, *Summary Judgment: Boon or Burden?*, 37 JUDGES' J., Summer 1998, at 26, 27 (“[A] defendant who brings and loses a summary judgment motion lives to fight another day. The losing plaintiff, however, loses not only the battle, but also the war.”).

78. *See id.* (explaining that a plaintiff must attempt to win the summary judgment decision “just to ensure his case goes forward”).

they believe is likely to be most favorable to their claim.⁷⁹ Our multivariate regression analysis of the merits decisions indicates that several districts are correlated with higher win rates for one side or the other—either overall or on various issues—even after we control for the characteristics of the patents, the patentees, the technology, and the industry.⁸⁰ We report the effects of district on overall win rates in Tables 3A and 3B.⁸¹

The two districts with the most patent cases—the Eastern District of

79. See Kimberly A. Moore, *Forum Shopping in Patent Cases: Does Geographic Choice Affect Innovation?*, 79 N.C. L. REV. 889, 920–23 (2001) (comparing forum shopping by infringers in declaratory judgment actions with that of patentees in infringement cases).

80. We used logistic regression (or logit) models, because each of our dependent variables (specific outcomes) is binary (or “dummy”—“yes” or “no”). Although multivariate regression usually assumes that all variables are independent of one another, this assumption does not hold when applied to studies of patent-infringement litigation. There are several reasons for this: (1) many cases involve the assertion of multiple patents, and decisions about these patents are made by the same judge and jury; (2) it is common to find in a data set that the same patent has been litigated in more than one separate lawsuit against different defendants, and even though the decision makers may be different, the same patent has the same attributes in each case; and (3) some cases will be consolidated, with the same decision maker deciding certain issues—usually only pretrial summary judgments, but sometimes trial decisions as well. Allison & Lemley, *supra* note 2, at 245; Allison et al., *Patent Quality*, *supra* note 13, at 678–79; Kesan & Ball, *supra* note 16, at 261. To account for the lack of complete independence among observations, we clustered on the standard errors of the unique patent numbers.

81. In addition to addressing the problem caused by lack of complete independence among our observations, we also had to contend with the fact that when running multiple tests from the same data set, there is the problem that we might obtain one or more findings of statistical significance by pure chance. Of the various techniques that have been proposed for correcting this problem, we decided that the use of bootstrapping would best serve our needs. To correct for any possible false significance findings (false discovery rate) resulting from doing multiple tests from the same data set, we used a bootstrapping procedure when running the logistic regressions on the various merits decisions. This procedure consisted of first resampling the original data to construct fifty samples with the original size. Thus, we had 949 observations, and from that we took a random sample of 949 fifty different times. Each random sample from the original 949 observations is clearly *not* identical to the original 949 observation sample because of the randomness of the samples—randomness will miss some of the observations and duplicate others. We then ran the logistic regression on the first random sample and generated a coefficient, standard error, and *p*-value. Random sample 1 was then added to the original data set of 949 observations. Then, random sample 2 was taken, another logistic regression was run on this second sample, and a second coefficient was generated, along with a standard error and *p*-value. Random sample 2 was then added back into the set consisting of the original 949 observations plus the first random sample. This process was repeated a total of fifty times. Finally, we averaged the fifty coefficients and derived a final standard error and *p*-value. Note that we clustered on the standard errors of the unique patent numbers when running each of the fifty logistic regressions. Also, the combination of bootstrapping and standard-error clustering was employed for each regression model—there was a separate regression model for each of the merits outcomes. We were required to do separate logits on each merits outcome, and could not combine all of these outcomes into a single multinomial regression model because the different outcomes possible for each patent were not independent of one another. See generally Joseph P. Romano et al., *Control of the False Discovery Rate Under Dependence Using the Bootstrap and Sub-sampling*, 17 TEST 417 (2008) (discussing the merits of the bootstrap method to control for a false discovery rate while testing *s* null hypotheses simultaneously).

Texas and the District of Delaware—were both significantly more likely to rule for the patentee in the cases we studied than were the “non-busy” patent districts.⁸² So too was the Southern District of New York. By contrast, only one district was significantly less likely to rule for patentees—the Central District of California.⁸³

Our results are largely but not completely consistent with prior work on district-specific variation in outcomes.⁸⁴ While prior papers have found some differences in district outcomes, some of them were testing different questions. Mark Lemley, Jamie Kendall, and Clint Martin, for instance, tested only trial outcomes, not all case outcomes.⁸⁵ And while those studies used multivariate regressions, each included different variables.⁸⁶ Our findings represent results from a number of multivariate regressions that account for all of the other independent variables in our study. We show results for ten of the major outcomes across some of the top districts in Tables 3A and 3B. All of these outcomes have large enough observations for the percentage rates to be meaningful. Moreover, a test comparing the percentage rates across districts showed that the differences were highly significant among districts for all of the ten outcomes—the significance level for nine of the ten outcomes across districts was <0.01 .⁸⁷ These differences are striking. Forum shopping, it seems, can pay dividends for cases that reach merits decisions.

It is important to keep in mind, however, that the regression results we report show correlations and are not proof of causation. The success of patentees in any particular district may be a function of the quality of cases brought in that district rather than any particular pro- or anti-patent sentiment.⁸⁸ For instance, it is possible that the weaker the patent

82. The omitted districts in this analysis are all districts other than the top thirteen. These other districts were combined into one category for measurement.

83. It is also notable that the Central District of California had fewer merits decisions than most other busy districts, but we have not tested for a relationship between the two.

84. See, e.g., Mark A. Lemley, Su Li & Jennifer M. Urban, *Does Familiarity Breed Contempt Among Judges Deciding Patent Cases?*, 66 STAN. L. REV. (forthcoming 2014) (manuscript at 23) (determining that Delaware courts are more likely to rule for patentees); Mark A. Lemley, Jamie Kendall & Clint Martin, *Rush to Judgment? Trial Length and Outcomes in Patent Cases*, 41 AIPLA Q.J. 169, 185 (2013) (finding no significant differences by district in trial results).

85. See Lemley, Kendall & Martin, *supra* note 84, at 172.

86. See Lemley, Li & Urban, *supra* note 84 (manuscript at 15–16) (selecting dependent and independent variables); Lemley, Kendall & Martin, *supra* note 84, at 176 n.19 (describing the authors’ decision not to include other variables of possible interest).

87. For the other outcome—Invalidity at any stage based on Section 102 prior art—the differences in rates among districts was significant at 0.017, very close to the <0.01 of the other nine outcomes.

88. See Lemley, Kendall & Martin, *supra* note 84, at 184–85 (finding that the evidence “does not support the conclusion that the district in which a case is litigated significantly affects the likelihood that the jury will find for the patentee”).

infringement claim, the more likely an accused infringer is to seek declaratory relief. Separately, it is possible that stronger cases are brought in certain districts. We cannot rule out these possibilities.

Table 3A: Ten Major Outcomes by District

	TX ED	DE D	CA ND	CA CD	CA SD	NY SD	IL ND	WI WD
% Win Rates by District; X = No Observations								
Patentee Definitive Winner	45	33	15	5	20	54	5	32
SJ Invalid Any	18	22	44	59	18	31	56	17
SJ No infr + stip. jdg no infr	45	64	64	64	54	64	41	75
Patentee Trial Win	72	49	50	83	55	100	20	88
Invalidity-All— Any Stage	23	40	68	63	60	25	50	18
Invalidity-102 Prior Art—Any Stage	17	23	44	63	55	20	67	0
Invalidity-103 Obvious—Any Stage	25	23	50	43	73	0	25	20
Invalidity-112 In- definiteness—Any Stage	15	20	40	67	0	0	25	0
Invalidity-112 En- ablement & Writ- ten Descr. —Any Stage	0	32	20	100	57	0	67	0
Total Direct In- fringement—All Stages	48	42	10	12	48	77	29	29

Table 3B: Ten Major Outcomes by District (cont.)

	NJ D	MA D	VA ED	OH ND	TX SD	All Other Dist.	Overall
% Win Rates by District; X = No Observations							
Patentee Definitive Winner	17	0	19	25	33	19	26
SJ Invalid Any	39	13	19	0	11	40	31
SJ No infr + stip. jdg no infr	56	53	65	0	91	58	57
Patentee Trial Win	27	0	43	50	60	71	61
Invalidity-All—Any Stage	60	64	60	0	24	53	42
Invalidity-102 Prior Art—Any Stage	14	43	22	0	27	41	31
Invalidity-103 Obvious—Any Stage	29	71	13	0	0	42	30
Invalidity-112 Indefiniteness—Any Stage	X	0	38	0	8	31	18
Invalidity-112 Enablement & Written Descr.—Any Stage	80	57	100	0	0	6	22
Total Direct Infringement—All Stages	27	23	29	25	50	34	36

Table 4: Definitive Win Rate by District—Multivariate (Logit) Regression Results⁸⁹

District	Patent Owner Definitive Winner
TX ED	1.252*** (0.331)
DE D	0.745** (0.337)
CA ND	-0.316 (0.426)
CA CD	-1.532** (0.607)
CA SD	0.0522 (0.551)
NY SD	1.593*** (0.493)
IL ND	-1.557*** (0.599)
WI WD	0.685 (0.481)
NJ D	-0.12 (0.617)
MA D	X
VA ED	-0.00844 (0.59)
OH ND	0.34 (0.781)
TX SD	0.745* (0.386)
N	620
Standard Errors in Parentheses	
* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$	
X = District omitted because of too few observations, excessive collinearity, or lack of randomness among observations	

89. Because we performed quite a few separate logistic regression tests using the same data set, there is a chance of deriving a finding of statistical significance by pure chance, which is often referred to as the false discovery rate problem (false positive finding of significance). As noted above, we used a bootstrapping methodology to minimize this risk. *See supra* note 81. In the table above, the districts are the independent variables and the specific outcome—Definitive Winner—is the dependent variable.

5. *Diversification Works*.—Modern patent litigation is often about more than enforcing a single patent.⁹⁰ A significant fraction of the cases in our study involved decisions on more than one patent. Notably, we find that cases in our study that evaluated more than one patent were significantly more likely to rule for the patentee, both in final outcome and in interim decisions. Notably, our finding is not merely that patentees who litigate multiple cases are more likely to win on at least one of them, but that the fact that a court rules on multiple patents is associated with an increased patentee win rate on each patent.

In addition to prevailing more overall, patentees also fared better on validity issues in multi-patent decisions. Specifically, patentees were significantly more likely to be granted summary judgment on a validity issue on a particular patent when the court ruled on multiple patents. It is possible that redundancy or diversification works, increasing the chances that the patentee will prevail on each patent. Here, the fractioning of patent law may work in favor of patentees. If a patentee prevails on a single patent in a lawsuit involving multiple patents, the patentee is entitled to damages and possibly an injunction.⁹¹ In fact, the damages may be the same for infringement on a single patent and infringement of multiple, related patents.⁹² Alternatively, it is possible that causation works the other way, and that patentees with stronger inventions are more likely to obtain and assert multiple patents and take the case to judgment. That said, there may be multiple selection effects that contribute to these results, including the fact that the number of patents asserted may affect how parties decide to move for judgment and how courts evaluate those motions, so we urge caution in interpreting this result.

6. *Foreign Inventors Do Just Fine*.—A number of studies have sought to evaluate whether the U.S. patent system is biased against foreigners,⁹³ as a number of foreign companies suspect.⁹⁴ In this study, we look not at the location of the litigants, but at the domicile of the inventors themselves.

90. See Gideon Parchomovsky & R. Polk Wagner, *Patent Portfolios*, 154 U. PA. L. REV. 1, 27 (2005) (proposing that the real value of patents lies in their aggregation into portfolios).

91. See John M. Golden, *Principles for Patent Remedies*, 88 TEXAS L. REV. 505, 514 (2010) (noting that injunctions are a “standard remedy for patent infringement,” and even when such injunctions “are unavailable, a patent owner may recover money damages”).

92. This Article utilizes each patent in a lawsuit as the unit of observation. Future work includes transforming the unit of observation to each lawsuit and performing similar empirical analysis. Analyzing the data using the lawsuit as the unit of observation may shed more light on litigation involving multiple patents.

We define a patent as being of foreign origin if a majority of its inventors were domiciled outside the United States; that is, we emphasize the geographic origin of the underlying *invention*, not its owner.⁹⁵ There were 146 foreign-origin inventions out of 777 patents litigated to a merits decision in our study. 98 of those 146 patents were filed first in a foreign country, but not always in the country where the invention originated. The correlation between foreign-origin invention and foreign-priority-filing country was 0.72, which is high but not extremely high. Patents on only five foreign-origin inventions were filed first in the European Patent Office (EPO).

How did those foreign-origin patents fare in litigation? We find a very strong result: patents of foreign origin in our study were much more likely to prevail in court in a merits decision than those issued to domestic inventors. In addition to being more likely to prevail overall, foreign-inventor patents were less likely to be held invalid, less likely to be held invalid on summary judgment, and less likely to be found obvious.

This result was frankly surprising to us. It may suggest that there is no bias against foreign inventors, though it may be driven in full or in part by selection effects. Kimberly Moore found that foreign litigants were much less likely to enforce their patents in the U.S. courts, suggesting that foreign litigants might be selecting only their best patents for suit.⁹⁶ While we investigate foreign inventors and not necessarily foreign owners, the two are likely to be correlated, and a similar effect might be at work here. It is also possible that the entity size or status of plaintiffs that assert foreign-invented patents differs systematically from domestic ones. If non-practicing entities primarily assert U.S. patents, for instance, and if those entities are more likely to lose, those facts may explain our results. We intend to test this hypothesis in subsequent work.

7. *It's Good to Go First.*—Plaintiffs traditionally go first in litigation and get the last word as well. There is some reason to think that confers an

93. See Kimberly A. Moore, *Xenophobia in American Courts*, 97 NW. U. L. REV. 1497, 1504 (2003) (finding that juries in patent trials are biased against foreign parties); cf. Kevin M. Clermont & Theodore Eisenberg, Commentary, *Xenophilia in American Courts*, 109 HARV. L. REV. 1120, 1122–23 (1996) (finding that foreign parties are more likely to prevail in federal civil actions).

94. Moore, *supra* note 93, at 1497–98.

95. In the unusual case in which there was an equal split between the number of U.S. and non-U.S. inventors, the domicile of the assignee was used as a tiebreaker. There were no cases in which there were an equal number of U.S. and foreign inventors without there also being an assignee to break the tie.

96. Moore, *supra* note 93, at 1505.

advantage in general in litigation.⁹⁷

That seems to be true in patent law as well. Consistent with prior work,⁹⁸ we find that accused infringers who sue for declaratory judgment fare substantially better than other accused infringers in cases that reach a merits decision. They are more likely to win overall, more likely to establish that the patent is invalid, and more likely to win their invalidity argument on summary judgment. Notably, while declaratory judgment allows accused infringers rather than patentees to pick the forum, and we found above that some fora are more favorable to patentees than others, this result is independent of the district-specific effects. That is, the benefit that declaratory judgment plaintiffs get is not simply a function of their ability to have their case heard in a more favorable forum.

Again, however, we encourage the reader not to read too much into this result. Selection effects may be at work. It is possible, for instance, that accused infringers (or their counsel) who file declaratory judgments are more sophisticated than those who just wait to be sued. That greater sophistication may translate into greater win rates. Correspondingly, patentees who actually send threat letters that can trigger declaratory relief may be less sophisticated than others; experienced patent lawyers can generally avoid creating declaratory judgment jurisdiction.⁹⁹ We cannot test the quality of counsel on either side, but it is a possible explanation for these results.

8. *Patent Characteristics Don't Seem to Matter Much.*—Our final finding is quite surprising—the observable characteristics of the patents don't seem to have much, if any, bearing on the outcome of the cases involving those patents. Neither the number of adjusted citations

97. For general arguments that going first is an advantage in litigation, see, for example, Shari Seidman Diamond et al., *Juror Reactions to Attorneys at Trial*, 87 J. CRIM. L. & CRIMINOLOGY 17, 27 (1996). Bernard Chao is studying this effect experimentally. E-mail from Bernard Chao, Assistant Professor of Law, Univ. of Denver Sturm Coll. of Law (Apr. 22, 2014, 1:17 PM) (acknowledging that Chao, along with John Campbell, Chris Robertson, and David Yokum, is conducting a study tentatively titled *Assessing the Substantive Effects of Declaratory Judgment Actions in Patent Litigation*).

98. See Moore, *supra* note 79, at 920–93 (finding that when accused infringers choose the forum, such as through a declaratory judgment action, the infringer “is much more likely to win”); Kimberly A. Moore, *Jury Demands: Who’s Asking?*, 17 BERKELEY TECH. L.J. 847, 859–61 (2002) (“Accused infringers generally bring declaratory judgment actions when they believe they have a strong case on the merits.”).

99. See Kristin Johnson Doyle, *Patent Demand Letters: Avoiding Declaratory Judgment Jurisdiction—Part 2 of 2*, INTELL. PROP. TODAY (Feb. 2010), <http://www.iptoday.com/issues/2010/02/patent-demand-letters-avoiding-declaratory-judgment-jurisdiction-part-2-2.asp> (“[U]se of smart strategies when dealing with alleged infringers may serve to shield the patent owner from declaratory judgment jurisdiction.”).

received¹⁰⁰ nor the number of prior art references have any significant correlation to overall win rates, validity, or infringement outcomes. Citations seem to tell us nothing about whether patents are valid or whether they are likely to be infringed. That is remarkable given how much effort economists have spent measuring the value of innovation by patent citation counts.¹⁰¹

More generally, it is notable how little explanatory power the group of independent variables in our model has. The pseudo R^2 is a measure in logit regression of how much power the independent variables together have in explaining a dependent variable. Stated somewhat differently, it estimates how well the model (group of independent variables) fits the data. The pseudo R^2 s in our regressions reported in Table 5 for ten major outcomes are very low, revealing that most of the variation in patent litigation outcomes is not predictable, at least based upon the extensive variables we captured.¹⁰² In other work we consider some variables not present here, including industry and technology area.¹⁰³ While there are significant differences in patent-litigation outcomes by industry and technology, even including those variables does not explain most of the differences in patent-litigation outcomes. The characteristics of individual lawyers, clients, and judges seem to matter quite a bit. We think that is as it should be.

100. For an explanation of the adjustment process, see *supra* note 26.

101. See, e.g., David S. Abrams, *Did TRIPS Spur Innovation? An Analysis of Patent Duration and Incentives to Innovate*, 157 U. PA. L. REV. 1613, 1616 n.9 (2009) (citing numerous articles assessing patent value by citation count).

102. This contrasts with Michael J. Mazzeo, Jonathan Hillel & Samantha Zyontz, *Explaining the “Unpredictable”: An Empirical Analysis of U.S. Patent Infringement Awards*, 35 INT’L REV. L. & ECON. 58, 67 (2013), which finds that damages (as opposed to liability rulings) are predictable based on some simple variables.

103. See John R. Allison, Mark A. Lemley & David L. Schwartz, PowerPoint: Differences in Patent Litigation Outcomes by Technology and Industry (2014) (on file with authors).

Table 5: Explanatory Power of Patent Characteristics¹⁰⁴

Pseudo R^2		
Outcome (each is a dependent variable in a logistic regression model)	Pseudo R -Squareds for each model (outcome)	N for each model (outcome)
Patentee Definitive Winner	0.145	636
SJ Invalid Any	0.0569	426
SJ No infr + stip. jdg no infr	0.0123	509
Patentee Trial Win	0.0510	290
Invalidity-All—Any Stage	0.0876	439
Invalidity-102 Prior Art—Any Stage	0.0624	231
Invalidity-103 Obvious—Any Stage	0.133	258
Invalidity-112 Indefiniteness—Any Stage	0.171	175
Invalidity-112 Enablement & Written Descr.—Any Stage	0.171	137
Total Direct Infringement—All Stages	0.152	530

Conclusion

The overall picture painted by our data is complex. In many ways, patent litigation is rather different than it was when we conducted our original study. The top districts for patent litigation—the Eastern District of Texas and the District of Delaware—were not nearly as important twenty years ago. The *Markman* hearing did not exist in our original study.¹⁰⁵ Patent assertion entities (referred to by some as “patent trolls”) were a

104. This particular set of regressions also included six technology areas—mechanical, electronics, chemistry, biotechnology, software, and optics, but the pseudo R^2 's were at the same very low levels for regressions run with only the eight patent characteristics, and run in several other ways.

105. *Markman* was not decided until the final year of our study, so its effect on our data was insignificant. *Markman v. Westview Instruments, Inc.*, 517 U.S. 370 (1996); see Allison & Lemley, *supra* note 2, at 194 (stating that the authors used data “from early 1989 through 1996”).

minor feature of patent litigation in the 1990s.¹⁰⁶ And the most successful validity challenges today—patentable subject matter and indefiniteness—were virtually unknown twenty years ago.¹⁰⁷

At the same time, many of our results will sound familiar to experienced students of the patent system. Ten years ago, Janicke and Ren found that patentees won only 25% of decided cases;¹⁰⁸ we find that number virtually unchanged today.¹⁰⁹ Forty-six percent of patents whose validity was decided in the 1990s were held invalid;¹¹⁰ today the invalidation rate is 43%. Much has changed about patent law, but the overall dynamics of patent litigation—in which patentees win at trial but not on summary judgment, and in which patentees win each individual issue but lose overall—remain remarkably similar to the patent litigation we studied twenty years ago.

106. See Sara Jeruss, Robin Feldman & Joshua Walker, *The America Invents Act 500: Effects of Patent Monetization Entities on US Litigation*, 11 DUKE L. & TECH. REV. 357, 358–62 (2012) (reporting an increase in patent-monetization-entity suits in the last decade). *But cf.* Christopher A. Cotropia, Jay P. Kesan & David L. Schwartz, *Unpacking Patent Assertion Entities*, 99 MINN. L. REV. (forthcoming 2014) (finding that nearly all of the supposed increase in patent assertion entity litigation from 2010 to 2012 is explained by the joinder provisions of the America Invents Act; Cotropia et al. hypothesize that large increases occurred prior to 2010).

107. See Allison & Lemley, *supra* note 2, at 208 (finding that, of the 138 patents held invalid in the study population, only 1 was held invalid on patentable subject matter grounds, and only 8 were held invalid on indefiniteness grounds).

108. Janicke & Ren, *supra* note 14, at 5.

109. The continuity may be even greater than that. Matthew Henry and John Turner study patent litigation going back to 2009, and find with two exceptions the patentee's overall odds of winning hover between 27% and 29%. Matthew D. Henry & John L. Turner, *Across Five Eras: Patent Enforcement in the United States 1929-2006*, at 4 (June 2013) (unpublished manuscript), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2274383.

110. Allison & Lemley, *supra* note 2, at 205.