

ARTICLES

Patent Quality and Settlement Among Repeat Patent Litigants

JOHN R. ALLISON,* MARK A. LEMLEY** & JOSHUA WALKER***

Repeat patent plaintiffs—those who sue eight or more times on the same patents—have a disproportionate effect on the patent system. They are responsible for a sizeable fraction of all patent lawsuits. Their patents should be among the strongest, according to all economic measures of patent quality. And logic suggests that repeat patent plaintiffs should be risk averse, settling more of their cases and taking only the very best to trial to avoid having their patents invalidated. In this Article, we test those hypotheses. We find that repeat patent plaintiffs are somewhat more likely to settle their cases. But to our surprise, we find that when they do go to trial or judgment, overwhelmingly they lose. This result seems to be driven by two parallel findings: both software patents and patents owned by nonpracticing entities (so-called “patent trolls”) fare extremely poorly in court. We offer some possible explanations for why a group of apparently weak patents nonetheless has so much influence over the patent system and some preliminary thoughts about how these findings should shape the patent reform debate.

TABLE OF CONTENTS

INTRODUCTION	678
I. THE MOST-LITIGATED PATENTS	681
II. RESULTS	686
A. OUTCOMES AND SETTLEMENTS	686

* Spence Centennial Professor, McCombs Graduate School of Business, University of Texas at Austin. © 2011, John R. Allison, Mark A. Lemley & Joshua Walker.

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

*** CEO, Lex Machina, Inc. Thanks to participants at the Harvard Law and Economics Workshop, the Columbia Law and Economics Workshop, the Conference on Empirical Legal Studies, the Northwestern University Patent Litigation Roundtable, and the National Bureau of Economic Research Summer Intellectual Property Institute, and to Ian Ayres, Jonathan Barnett, Jim Bessen, Colleen Chien, Wesley Cohen, Terry Fisher, Michael Frakes, Alberto Galasso, Rose Hagan, Dietmar Harhoff, Scott Hemphill, Brian Jones, Louis Kaplow, Rob Kunstadt, Clarisa Long, Paul Morgan, Craig Nard, Michael Risch, Ben Roin, Dave Schwartz, Steve Shavell, Ted Sichelman, and Scott Stern for helpful comments, and to David Rizk for research assistance. Ying He, Ph.D. candidate, McCombs School of Business, provided invaluable statistical assistance.

B.	ENTITY SIZE AND STATUS	689
1.	Entity Size	689
2.	NPE Status	691
C.	INDUSTRY AND TECHNOLOGY DIFFERENCES	695
D.	LOGISTIC REGRESSION	697
III.	EXPLANATIONS AND IMPLICATIONS	700
A.	OTHER EXPLANATIONS FOR SETTLEMENT DATA	700
1.	Number of Defendants	700
2.	License Agreements and Plaintiff Stakes	704
B.	WEAK LITIGATED PATENTS	705
1.	Multiple Lawsuits As Evidence of Obviousness?	706
2.	Overclaiming?	707
3.	Are Software Patents or Trolls the Problem?	707
C.	OUTLIERS: CRAZY LIKE A FOX?	709
D.	DO WE UNDERSTAND PATENT VALUE?	710
E.	DO WE UNDERSTAND LITIGATION BEHAVIOR?	711
	CONCLUSION	712

INTRODUCTION

Patent owners who file lawsuits put their underlying patents at risk. A significant percentage of litigated patents are held invalid,¹ and a finding of invalidity is the death knell for a patent. Because of the arcane civil procedure doctrine of offensive nonmutual collateral estoppel, the consequences of validity and invalidity holdings are highly asymmetric.² A patentee who wins a suit against defendant *A*, having proven the patent infringed and fought off a validity challenge, gets no credit for the win in a subsequent suit against defendant *B*.³ Because *B* was not a party to the first suit, it is entitled to once again challenge

1. See, e.g., John R. Allison & Mark A. Lemley, *Empirical Evidence on the Validity of Litigated Patents*, 26 AIPLA Q.J. 185, 205 (1998) (finding that 46% of patents litigated to judgment are held invalid); Paul M. Janicke & LiLan Ren, *Who Wins Patent Infringement Cases?*, 34 AIPLA Q.J. 1, 5–6 (2006) (finding that patentees win only 25% of cases litigated to judgment, in part because of invalidity and in part because of noninfringement).

2. E.g., *Blonder-Tongue Labs, Inc. v. Univ. of Ill. Found.*, 402 U.S. 313, 325 (1971) (discussing consequences and fairness of nonmutuality of estoppel in patent litigation).

3. Indeed, the Federal Circuit has even refused to allow the prior judgment to be considered by the jury in a subsequent lawsuit. See *Mendenhall v. Cedarapids, Inc.*, 5 F.3d 1557, 1575 (Fed. Cir. 1993).

the validity of the patent, even on the very same grounds rejected in the first lawsuit.⁴ The same is true in subsequent suits against defendants *C*, *D*, *E*, and so on.⁵ Indeed, for this reason Federal Circuit Judge Rich used to insist that patents were not held valid, but merely held “not invalid.”⁶ By contrast, should *A* succeed in proving the patent invalid, the game is up. The doctrine of collateral estoppel will prevent the patentee from enforcing the patent against *B*, *C*, *D*, or *E*; each of those defendants is entitled to rely on the patentee’s prior loss to defeat the lawsuit.⁷ And even existing licensees will be permitted to stop paying royalties and file their own challenge to the patent.⁸

If you are a patent owner who faces multiple infringers, the deck is stacked against you. How might patentees respond to this asymmetry? One possible option is to sue all the defendants at once. Doing so makes the resulting case more complex, but it insulates the patentee from the risk of having to litigate validity again and again. And even the complexity of the multi-defendant suit can redound to the patentee’s advantage if defendants cannot agree to present a unified front on issues like claim construction.⁹ Alternatively, patentees that are involved (or expect to be involved) in multiple lawsuits might reasonably be more likely to settle cases rather than risk taking them to judgment because an adverse judgment wipes out all the other lawsuits as well.¹⁰

Corollaries of these anticipated patentee responses are that we should expect rational patentees to prefer filing single lawsuits against multiple entities rather than filing multiple suits in parallel,¹¹ that patentees who file multiple lawsuits should settle more often than other patentees, and that those patentees who do file multiple lawsuits and litigate those suits to judgment might do so because

4. *Id.*

5. *Id.*

6. *See, e.g.*, Thomson, S.A. v. Quixote Corp., 166 F.3d 1172 (Fed. Cir. 1999).

7. In patent law, this is true even if the patentee has already litigated and won one or more cases before its first loss. *See, e.g.*, Mendenhall v. Barber-Greene Co., 26 F.3d 1573, 1577 (Fed. Cir. 1994); Miss. Chem. Corp. v. Swift Agric. Chems. Corp., 717 F.2d 1374, 1379 (Fed. Cir. 1983); Stevenson v. Sears, Roebuck & Co., 713 F.2d 705, 709 (Fed. Cir. 1983).

8. MedImmune, Inc. v. Genentech, Inc., 549 U.S. 118, 137 (2007); Lear, Inc. v. Adkins, 395 U.S. 653, 673–74 (1969).

9. The patentee’s costs go up with more defendants because they must prove more infringements. But the increase is not linear because many of the issues—claim construction, validity, unenforceability—will take just as much time to litigate against one defendant as against ten. On the defense side, there may be some cost savings from sharing research, but it has been our experience that trying to coordinate strategy among many different lawyers is often more rather than less costly.

10. For a discussion of the game-theoretic of claim preclusion, see Bruce L. Hay, *Some Settlement Effects of Preclusion*, 1993 U. ILL. L. REV. 21 (1993). Hay suggests that claim preclusion might not drive more settlements because the parties could bargain for their own preclusion rules. But he is focused only on mutual collateral estoppel, which governs subsequent suits between the same parties. His argument does not apply to nonmutual collateral estoppel—the kind we are concerned with here. *Cf.* Michael Risch, *Patent Challenges and Royalty Inflation*, 85 IND. L.J. 1003 (2010) (modeling repeat-play behavior in patent litigation).

11. Some patentees will file “test suits” against small companies early on, in hopes of building a war chest and a record of enforcement before suing a larger array of companies. But that strategy supports serial patent lawsuits, not the suits in parallel that we observe.

they have a stronger than average patent and therefore face less risk of invalidity. In this Article, we test each of these hypotheses using a unique database, the Stanford IP Litigation Clearinghouse. In Part I, we explain the source of our data and our methodology.

Part II provides our results. There has been an increase in the number of defendants named in a particular suit. Nonetheless, many patents are enforced in multiple lawsuits. We focus on the most-litigated patents—all 106 patents that have been the subject of eight or more lawsuits since the year 2000. We emphasize that being involved in eight or more lawsuits does not mean that these patents went to judgment eight or more times. Most cases settle, and many cases in our data set were still pending as of our closing date. In fact, very few cases in our data set had more than one judgment.

Those most-litigated patents exhibit characteristics that economists have associated with value. Indeed, economists looking at these patents would say they are far more valuable than the average patent or even than the average-litigated patent.¹² We also find that serial patent litigants do react to the greater consequences of losing; cases involving the most-litigated patents are indeed more likely to settle than ordinary-litigated patents with a high degree of significance. But to our great surprise, we find that the willingness of these patentees to litigate their cases to judgment is a mistake. Far from being stronger than other litigated patents, the most-litigated patents that go to judgment are far more likely to be held invalid or not infringed. The differences are dramatic. Once-litigated patents win in court almost 50% of the time, while the most-litigated—and putatively most valuable—patents win in court only 10.7% of the time.¹³

The results are equally striking for patents owned by nonpracticing entities (NPEs) and for software patentees. NPEs and software patentees overwhelmingly lose their cases, even with patents that they litigate again and again.

12. See Gwendolyn G. Ball & Jay P. Kesan, *Transaction Costs and Trolls: Strategic Behavior by Individual Inventors, Small Firms and Entrepreneurs in Patent Litigation 2* (Ill. Law & Econ. Papers Series, Research Papers Series No. LE09-005, 2009), available at <http://papers.ssrn.com/abstract=1337166> (finding “evidence that when suing a large alleged infringer, small parties are only enforcing their most ‘valuable’ patents.”); Timo Fischer & Joachim Henkel, *Patent Trolls or Markets for Technology—An Empirical Analysis of Trolls’ Patent Acquisitions 2* (Dec. 14, 2009) (working paper), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1523102 (finding that nonpracticing entities (NPEs) purchased patents that had characteristics associated with patent value).

13. Comparing them shows that patent owner win rates are greater with an exceptional degree of statistical significance—the p-value (the statistical measure of the likelihood that the result arose by chance) is 0. This calculation includes plaintiff wins by default judgment in the once-litigated set of patents; there were no default wins in the most-litigated set. These percentages are obtained when settlements are excluded from the denominator, with only win rates being compared. The percentage win rates are naturally much smaller for both plaintiffs and accused infringers than when the huge numbers of settlements are in the denominator, but the comparison is just as striking—with settlements in the denominator, patent owners in the most-litigated group prevailed only 1% of the time, while those in the once-litigated group prevailed 7.6% of the time. A statistical comparison still yields a p-value of 0.

Software patentees win only 12.9% of their cases, while NPEs win only 9.2%.¹⁴

In Part III, we offer some preliminary thoughts about what might explain these results. We first investigate whether the outcome data are the result of clustering—a few cases in multidistrict litigation that invalidate or hold not infringed multiple patents at once. We find some evidence of clustering but not enough to explain the full differences in the outcomes. One possible explanation is that defendants rather than plaintiffs are driving the decision to take a weak patent to judgment by simply refusing to settle. Another possibility is that the enforcement of a patent against multiple infringers is an indication of widespread simultaneous invention and, hence, of obviousness. A third possibility is that these plaintiffs are by definition outliers, a status that may affect both their rationality and their skill at litigation. For instance, the decision to sue in multiple different suits rather than consolidating the suits may itself provide evidence about the sophistication of the plaintiff and, hence, the likely outcome. Finally, it is possible that the economics of patent litigation make it profitable to enforce even patents that are overwhelmingly likely to lose in court.

None of these explanations is entirely satisfactory. The result is a bit of a puzzle. The most-litigated patents—the patents that by all measures should be the most important—seem mostly to fail in court. So too do the patents that have occupied the most public attention—software patents and patents filed by NPEs. That fact has implications for patent policy and, in particular, for patent reform directed at litigation abuse. It appears that, as a society, we are spending a disproportionate amount of time and money dealing with a class of weak patents. Our results may also have implications for our models of patent value and of rational behavior in litigation because it appears we know quite a bit less than we thought about what makes patents valuable.

I. THE MOST-LITIGATED PATENTS

In prior work, two of this Article's authors demonstrated that litigated patents have significantly different characteristics than other patents.¹⁵ They include more claims, cite more prior art, are cited more often by later patents, file more continuation applications, and come from larger "families" of patents.¹⁶ They

14. These percentages also are calculated with settlements excluded from the denominator, and comparisons in the case of both software patentees and NPEs show that they lose when compared with non-software patentees and product-producing companies, respectively, with an extremely high degree of statistical significance—again, with p-values of 0. Likewise, when the calculations are made with settlements left in the denominators, the statistical differences are just as striking, with p-values once again of 0.

15. John R. Allison et al., *Valuable Patents*, 92 GEO. L.J. 435 (2004) [hereinafter ALMT]; see also John R. Allison & Thomas W. Sager, *Valuable Patents Redux: On the Enduring Merit of Using Patent Characteristics to Identify Valuable Patents*, 85 TEX. L. REV. 1769 (2007) (defending the statistical power of the results in the earlier study); James Bessen, *The Value of U.S. Patents By Owner and Patent Characteristics*, 37 RES. POL'Y 932, 932 (2008) (finding that "[l]itigated patents are more valuable, as are highly cited patents").

16. ALMT, *supra* note 15, at 451–58.

are also concentrated in some industries but not others; semiconductor patents are particularly unlikely to be litigated.¹⁷ Many of these characteristics are within the control of the patent applicant, and most are known by the time the patent issues.¹⁸ Allison, Lemley, Moore, and Trunkey (ALMT) suggest that these characteristics are evidence of the private value of patents,¹⁹ following a significant economic literature correlating each of these attributes with value to the patent owner.²⁰

That prior work depended significantly on a randomly selected sample of cases actually litigated, collected by hand from district courts around the country. The development of the Stanford IP Litigation Clearinghouse in December 2008 opened up a second alternative.²¹ The Clearinghouse collects every patent lawsuit filed since January 1, 2000, in searchable format and links those suits to the patents involved.²² Using that database, in a prior paper we identified every patent that has been litigated eight or more times between January 2000 and February 2009 (including cases still pending). There were 106 such patents, which have been litigated in a total of 2,987 different patent-suit pairs in 478 different suits, often against multiple defendants.²³ For

17. *Id.* at 438.

18. *Id.* at 460.

19. *Id.* at 460–65.

20. See, e.g., Bessen, *supra* note 15, at 932 (finding that citation and litigation both correlate with patent value); Jean O. Lanjouw & Mark Schankerman, *Characteristics of Patent Litigation: A Window on Competition*, 32 RAND J. ECON. 129, 140 (2001) (finding that litigation correlates with patent value and that “[t]he number of claims is another . . . indicator of the ‘bits of information’ contained in a patent, and therefore its value”); Stuart J.H. Graham et al., *Post-issue Patent “Quality Control”: A Comparative Study of US Patent Re-examinations and European Patent Oppositions* (Nat’l Bureau of Econ. Research, Working Paper No. 8807, 2002) (comparing USPTO and EPO opposition mechanisms and finding that the most valuable patents were challenged in both systems); Dietmar Harhoff & Markus Reitzig, *Determinants of Opposition Against EPO Patent Grants—The Case of Biotechnology and Pharmaceuticals 1* (Ctr. for Econ. Pol’y Res., Discussion Paper No. 3645, 2002) (confirming that “valuable patents are more likely to be attacked”); Jean O. Lanjouw & Mark Schankerman, *Enforcing Intellectual Property Rights 5* (Nat’l Bureau of Econ. Res., Working Paper No. 8656, 2001) (“[M]ore valuable patents . . . are much more likely to be involved in suits.”); cf. Dietmar Harhoff et al., *Citations, Family Size, Opposition and the Value of Patent Rights*, 32 RES. POL’Y 1343, 1345 (2003) (finding that “[p]atents which are upheld against opposition . . . are particularly valuable”); Jean O. Lanjouw & Josh Lerner, *The Enforcement of Intellectual Property Rights: A Survey of the Empirical Literature*, 49/50 ANNALES D’ECONOMIE ET DE STATISTIQUE 223 (1998) (surveying the literature on the issue).

21. That site is presently operated by a private company, Lex Machina, Inc., and hosted at <http://www.lexmachina.org>.

22. Due to increased availability over time of electronic filings in federal court, the ability to identify patents in suits improves markedly later in time, particularly from 2003 forward. Moreover, electronic access also varies by district, potentially making this patent data set under-inclusive for certain districts despite hand-collection of cases from those districts. Nevertheless, the patents identified represent the best, most representative data set available.

23. John R. Allison, Mark A. Lemley & Joshua Walker, *Extreme Value or Trolls on Top? Evidence From the Most-Litigated Patents*, 158 U. PA. L. REV. 1 (2009). For purposes of this analysis we include declaratory judgment actions as well as actions filed by the patent owner; until 2007 the rules for declaratory judgment required a clear threat of suit by the patent owner. *Teva Pharm. USA, Inc. v. Pfizer, Inc.*, 395 F.3d 1324, 1333 (2005), *abrogated by* *MedImmune, Inc. v. Genentech, Inc.*, 549 U.S.

purposes of that study, we identified a randomly selected control set of 106 patents that have been litigated only once during this time period. This allowed us to extend the work ALMT did in 2003, comparing ordinary litigated patents (already outliers, as we have seen) to the most-litigated patents. To assess the ALMT value hypothesis, we collected data on the number of continuation applications filed leading to issuance of the patent; the raw and adjusted number of “forward citations” (citations to the patent by later patents); the number of “backward citations” to U.S. patents, foreign patents, and non-patent prior art (“prior art references”—the citations the patent makes to prior art); and the number of claims in each patent. Each of these factors has been identified in the economic literature as evidence of the value of a patent.²⁴ As ALMT predicted, those most-litigated patents exhibited even more evidence of private value and even more of an industry skew than did the average litigated patent.²⁵

In this Article, we use the same most-litigated cases, but collect information about the outcome of each case. In addition, we expand our sample of once-litigated cases to 343 (after exclusion of cases that were still pending or that were the subject of purely procedural dispositions) in order to generate a more valid statistical sample of cases that went to judgment or settlement.²⁶

For each litigated patent, we also determine small entity status (in other words, whether the patent owner at issue was an individual, university, or small business, collectively referred to as “small entities”), whether the patent being litigated covers a software invention, whether the patent is assigned before litigation, and—following Lemley and Myhrvold²⁷—the nature of the patent plaintiff, divided into one of twelve different “entity status” categories listed in Table 1.

Of the twelve classes of entity, only one (Class 8) involves enforcement by a patent owner that actually makes products. The rest are different types of NPEs, sometimes called “patent trolls” for the practice of hiding under a bridge they did not build and demanding a toll from surprised passers-by. Rather than take a

118 (2007). We count only separate lawsuits; many patent lawsuits are filed against multiple defendants in a single proceeding. If we count the number of patent assertion-defendant combinations, the number of observations is more than 10,000. There are 478 separate lawsuits involving these patents, but because many of those lawsuits involve several different patents in the data set, the total number of patent-suit pairs of the 106 patents is 2,987.

The choice of the lower bound of eight suits for repeat litigants is arbitrary. It could have been seven, or nine, or ten. We chose a number that gave us enough different patents to run meaningful statistical tests, while avoiding blurring the line between the most- and once-litigated patents.

24. See ALMT, *supra* note 15.

25. *Id.* at 462–63.

26. Every most-litigated patent was at one time a once-litigated patent. To avoid biasing our sample by including patents that simply were not yet multiply litigated, but were headed that way, we checked the status of the once-litigated patents a year after the cut-off date and excluded those that had additional suits filed. This check is not perfect, but makes it less likely that the sample of once-litigated patents includes cases that would eventually qualify as most-litigated patents.

27. Mark A. Lemley and Nathan Myhrvold have developed a twelve-class taxonomy of patent plaintiffs, which we follow here.

Table 1. Entity Status Classes

Entity Class 1 (Acquired patents)
Entity Class 2 (University heritage or tie)
Entity Class 3 (Failed startup)
Entity Class 4 (Former product company no longer producing)
Entity Class 5 (Individual inventor started company)
Entity Class 6 (University/government/NGO)
Entity Class 7 (Startup, preproduct)
Entity Class 8 (Product company)
Entity Class 9 (Individual)
Entity Class 10 (Undetermined)
Entity Class 11 (Industry consortium)
Entity Class 12 (IP subsidiary of product company)

position on what, if any, NPEs should be considered trolls, we classify each patent owner and let the reader decide. We do, however, report the results for practicing versus nonpracticing entities (that is, Class 8 versus all other classes excluding Class 10).²⁸ Notably, virtually all of the NPEs in our data set fall into Classes 1 (companies in the business of acquiring and asserting patents from others) and 5 (inventor-owned companies). So our data do not depend at all on things like whether universities, university spin-offs, or intellectual property (IP) subsidiaries count as patent trolls.

We also categorize each patent into both a technology area and an industry area in order to ascertain whether significant differences existed in the technology and industry areas.²⁹ In our description of technology and industry areas for

28. For a few patent owners, we could not identify their entity status after a diligent search. We classed those entities as Entity Class 10 (Undetermined) and excluded them from our entity status analyses. As a practical matter, however, that a diligent search could not identify what an entity did suggests that it is likely some form of NPE.

29. We did not attempt to create a comprehensive typology of such areas, but for obvious reasons we only identified and defined those technology and industry areas we actually encountered in the population of most-asserted patents and the sample of once-litigated patents. Although the size of our data sets is sufficient for sound statistical analysis, the relatively small number of observations necessarily results in our having encountered fewer technology and industry areas than we would have found in a much larger patent data set. The technology categories are not necessarily mutually exclusive because modern inventions so often involve multiple technologies.

Our industry categories are also not all mutually exclusive, reflecting the reality of modern industry crossovers. For example, a software-implemented telecommunications process or product rightly belongs in both a computer and a communications industry category. There are, however, fewer inventions belonging in more than one *industry* category than there are inventions belonging in more than one *technology* category because mixes of technologies in inventions are more common than industry crossovers.

Table 2. Technology and Industry Areas

<i>Technology Areas</i>
(1) Software
(2) Pure software
(3) Software business method
(4) Mechanical
(5) Electronics
(6) Optics (other than imaging)
(7) Imaging
(8) Biotechnology
(9) Chemistry
<i>Industry Areas</i>
(1) Computer
(2) Semiconductor
(3) Electronics
(4) Medical
(5) Pharmaceutical
(6) Biotechnology
(7) Chemical
(8) Communications
(9) Transportation
(10) Energy and utility services
(11) Financial
(12) Consumer goods and services
(13) Construction

inventions that we actually encounter in our data sets, we attempt to define the areas in a comprehensive way, and our definitions are thus broad enough to include specific inventions not actually found in our data sets. We report the list of technology and industry areas here; they are defined in detail in our prior work.³⁰

We test each of the results we report in this Article for statistical significance. We report such significance along with each result. In addition, we conducted a multivariate regression analysis that evaluates the interaction between the variables.

30. Allison, Lemley & Walker, *supra* note 23, at 6–11.

Finally, we evaluate the status of each of the cases as of August 2009. We categorize the outcomes of patent cases in the following ways: (1) settlements, including consent judgments; (2) procedural dispositions; (3) pending (including stays and transfers); (4) patentee wins, with and without default judgments; and (5) accused infringer wins. We count a case as a win by one party if that party has prevailed on a case-dispositive issue in the highest court to rule on the merits. Most of our outcomes are district court outcomes, though some have gone to the Federal Circuit. Notably, we count a case as a win for a party if the case against at least one defendant went to judgment, even though other defendants might have settled before judgment.³¹

II. RESULTS

A. OUTCOMES AND SETTLEMENTS

As noted above, our previous work suggests that the most-litigated patents are also the most valuable patents by all the available measures economists have used. Assuming that those value measures are accurate, it is reasonable to expect that the most-litigated patents are more likely to succeed in litigation than the once-litigated patents. At the same time, the doctrine of offensive nonmutual collateral estoppel means that the most-litigated patents are also the most vulnerable. A patentee who files one suit puts the patent at risk, but unless the patent is widely licensed to others, the scope of that risk is not much greater than the risk of losing the one suit. By contrast, a patentee who has multiple lawsuits pending at the same time should be very worried about the possibility of losing even one of those suits because doing so will bring down the whole edifice. These two assumptions can reinforce each other; it is reasonable to assume that where a serial patent plaintiff does not settle, but actually takes the case to judgment, the patents selected for litigation should be even stronger than the average most-litigated patent and more likely still to succeed in court.

As a result, we test two hypotheses: first, that the most-litigated patent owners will be risk averse, and therefore more likely than other litigants to settle their lawsuits before judgment; and second, that where the most-litigated patents *do* get litigated to judgment, the patentee is more likely to prevail than other litigants.

The basic descriptive results of our outcome study are presented in Table 3.

31. We considered and rejected the idea of studying outcomes on a defendant-by-defendant basis. While in some sense doing so would be the most accurate way to report the results, it would bias the statistical analysis by introducing large numbers of observations that are not truly independent. For instance, suppose that there were only two suits in our sample: a plaintiff that defeats a single defendant, and a plaintiff that loses a patent suit against fifty defendants who are in a joint defense agreement. To treat that as fifty patentee losses and one win, and conclude that patentees only win 2% of their lawsuits, would be misleading.

If we were to recalculate our results on a per-defendant basis, the most-litigated patentee win rates would drop even further, largely due to the invalidation of some Katz patents asserted against hundreds of defendants.

Table 3. Outcome Data³²

Outcomes	Most-Litigated Patents	Once-Litigated Patents
settlement/consent judgment	716	288
procedural disposition	15	19
pending (including stays and transfers)	40	83
transfers (including multidistrict litigation)	139	11
default judgment	0	10
plaintiff win	8	16
defendant win	67	29
plaintiff win rate (including defaults)	10.7%	47.3%

It is worth noting the substantial number of most-litigated patent suits that are resolved on procedural grounds, whether involving personal jurisdiction, standing, transfers to another venue, or multidistrict litigation (MDL) consolidation. It should not be surprising that these non-merits resolutions are more common in the more complex environment of multiple litigation. But the widespread use of the MDL process, which is quite new to patent litigation, is notable, as is the sizeable number of transfers outside the MDL context.³³ For purposes of our statistical analyses in the remainder of this Article, we exclude all pending cases (including stays), transfers, and procedural dispositions other than default judgments because none represent a resolution of the case either on the merits or through settlement. The result is a data set of 1,134 results, either on the merits, default judgment, or settlement, 343 of which are once-litigated patent outcomes and the rest are most-litigated patent outcomes.

Several facts stand out in these results. Most notably, our hypothesis that the most-litigated patents will fare well in litigation does not hold up. The most-litigated patent plaintiffs won only 10.7% of their cases, compared with 26% across all lawsuits³⁴ and 47.3% in the once-litigated set.³⁵

Statistical tests bear this out. We compare the proportion of win rates, testing the null hypothesis that there is no difference between the most-litigated and once-litigated patent outcomes.³⁶ We test the proportions in several ways, both

32. The rows that are struck through are those we do not include in subsequent tests.

33. Jamie H. McDole & Aaron D. Charfoos, *Multidistrict Litigation in Patent Infringement Cases 2* (Mar. 2005) (unpublished manuscript), <http://www.ipo.org/AM/Template.cfm?Section=Home&Template=/CM/ContentDisplay.cfm&ContentID=22938>.

34. Janicke & Ren, *supra* note 1, at 5–6.

35. The once-litigated win rate is only 35.6% if we exclude default judgments; this is discussed in more detail below.

36. Because the number of wins on the merits was arguably too small to use the Chi-square test, we used Fisher's exact test as well. Fisher's exact test is commonly used to test differences in proportions, particularly when the number of observations is relatively small and when a Chi-square test may not

Table 4. Most-Litigated vs. Once-Litigated Win Rates

	Most-Litigated (Without Default Judgments)	Once-Litigated (Without Default Judgments)	p-Value (Both Chi-sq. & Fisher's Exact)	Most-Litigated (Including Default Judgments)	Once-Litigated (Including Default Judgments)	p-Value (Both Chi-sq. & Fisher's Exact)
Settlements Included in Denominator	1.00%	4.80%	0	1.00%	7.60%	0
Settlements Excluded from Denominator	10.70%	35.60%	<0.01	10.70%	47.30%	0

including and excluding settlements in the denominator of decided cases, and both including and excluding default judgments as plaintiff wins. No matter which test we use, the differences are highly statistically significant—the most-litigated patentees were more likely to lose.³⁷

Our hypothesis of asymmetric stakes and a consequently greater tendency to settle among repeat patent litigants was confirmed with statistical significance, but the degree of confirmation is not as striking as the disconfirming results for win–loss rates. Settlement rates, with pending cases, procedural dispositions,

produce reliable results. Fisher's exact test is widely used because it requires fewer assumptions than alternative tests and worked well in the comparisons we made in this study. *See generally* Rick Routledge, *Fisher's Exact Test*, in 3 *ENCYCLOPEDIA OF BIostatISTICS 1961–64* (Peter Armitage & Theodore Colton eds., 2d ed. 2005).

Actually, the absolute differences in our comparisons were usually so great that the Chi-square test also found statistical significance with high degrees of confidence (low p-values).

37. There is a potential problem of censoring of results because a ruling that a patent is invalid will, if final, prevent the patentee from enforcing the patent in its other suits. If, for instance, a large percentage of the most-litigated patents were invalidated in the first lawsuit, they would not be enforceable in subsequent lawsuits. There is also a potential selection effect in the most-litigated set because cases that involve eight or more different lawsuits might do so only because they survived the first seven lawsuits (that is the patent was not held invalid in those earlier suits). Finally, findings might not be independent: whether one court holds a patent invalid or not infringed may strongly influence subsequent courts deciding on the same patent.

Censoring and selection do not seem to be a significant problem in our data set, for the simple reason that most of the most-litigated patents involve cases that were filed roughly in parallel, not in series. Most cases in our study settle, and others are still pending, meaning that we have only a small share of the most-litigated patents that are the subject of even one ruling on the merits, and very few that get more than one ruling on the merits. Some of those merits rulings are on infringement, not invalidity, and so will not create estoppel against subsequent lawsuits. Further, a ruling by the district court will not create an estoppel effect as long as that ruling is on appeal. The vast majority of our outcomes are district court rulings, which means the other lawsuits are not automatically precluded by even a finding of invalidity.

The independence concern is potentially more significant. If, for instance, we were to find that a few most-litigated patents were invalidated by a single court, and then multiple courts adopted that reasoning and invalidated the patent, that finding would not necessarily indicate that the most-litigated patents overall were likely invalid. But that is not what we find. There are virtually no instances of multiple rulings on the merits of the same patent in our data set. And when we exclude the few instances that do exist, the results do not change.

Table 5. Most-Litigated vs. Once-Litigated Settlement Rates

	Settlement Rate	p-Value (Both Chi-sq. & Fisher's Exact)
Most-Litigated Patents	90.50%	<0.01
Once-Litigated Patents	84.00%	

and transfers removed from the denominators, are reported in Table 5.³⁸

Settlements account for 90.5% of the outcomes in the most-litigated patent set and 84.0% of the outcomes among the once-litigated patents. This difference is highly statistically significant, confirming our hypothesis of a greater tendency to settle among repeat patent litigants.

Thus, when settlement fails and cases go to judgment, our hypothesis regarding the strength of the most-litigated patents does not hold up. Patent plaintiffs are repeatedly litigating patents that are less successful than the once-litigated patents. And although they are more likely to settle than the once-litigated patents, the disparity is less than theory would predict because the most-litigated patentees regularly lose the cases that do go to judgment. In Part III, we consider possible explanations for this surprising result.

B. ENTITY SIZE AND STATUS

1. Entity Size

We also test the relationship between entity size and entity status and the outcome of cases. The United States Patent and Trademark Office (PTO) distinguishes between “large” and “small” entities. A small entity is an individual, a university, a nonprofit organization, or a company with fewer than 500 employees.³⁹ Small-entity status is determined based on the entity that owned the patent when it issued from the PTO; one complication is that many patents are sold before litigation.⁴⁰ Nonetheless, PTO small-entity status is a useful proxy for the size of the patent plaintiff. We test the effect of the plaintiff’s size on the outcome of patent litigation.⁴¹

Considering only the patents themselves, the proportions of initial ownership by large and small entities are almost equal in the most- and once-litigated data

38. Of necessity, there were no MDL transfers in these cases as there were among the most-litigated patents.

39. 35 U.S.C. § 41(h)(1) (2006) (incorporating by reference the Small Business Act § 3, 15 U.S.C. § 631 (2006)).

40. Kimberly A. Moore, *Populism and Patents*, 82 N.Y.U. L. REV. 69, 96–97 (2007) (documenting a high rate of prelitigation transfer of patents).

41. As we did in our prior article, we removed Ronald S. Katz Technology Licensing LLP from the entity-size analysis because Katz owns some patents classed as small-entity patents and others classed as large-entity patents. In our most-litigated set, Katz did not claim small-entity status (with the resulting lower filing fees) that he was entitled to claim, and thus, these patents were recorded as having been issued to large entities when they should not have been.

Table 6. Large-Entity vs. Small-Entity Settlement Rates Across Both Most- and Once-Asserted Data Sets

	Settlement Rate	p-Value (Both Chi-sq. & Fisher's Exact)
Large Entities	87.10%	0.9
Small Entities	87.50%	

sets: 53.5% of most-litigated patents and 47.8% of once-litigated patents were issued to large entities. The picture is different, however, when one looks at the proportion of actual patent-suit pairs in litigation, where large entities account for a surprisingly small percentage of the most-litigated patents.⁴² Because small entities are disproportionately represented in the actual litigation of most-litigated patents (those patents that were asserted against accused infringers at least eight times), patents that were initially issued to large entities represent only 22.4% of the patent-suit pairs in the most-litigated group, compared to 47.8% of the once-litigated group.

We find no significant difference in the propensity of large and small entities to settle their cases. Indeed, the settlement rates are virtually identical at 87.1% and 87.5%, respectively.

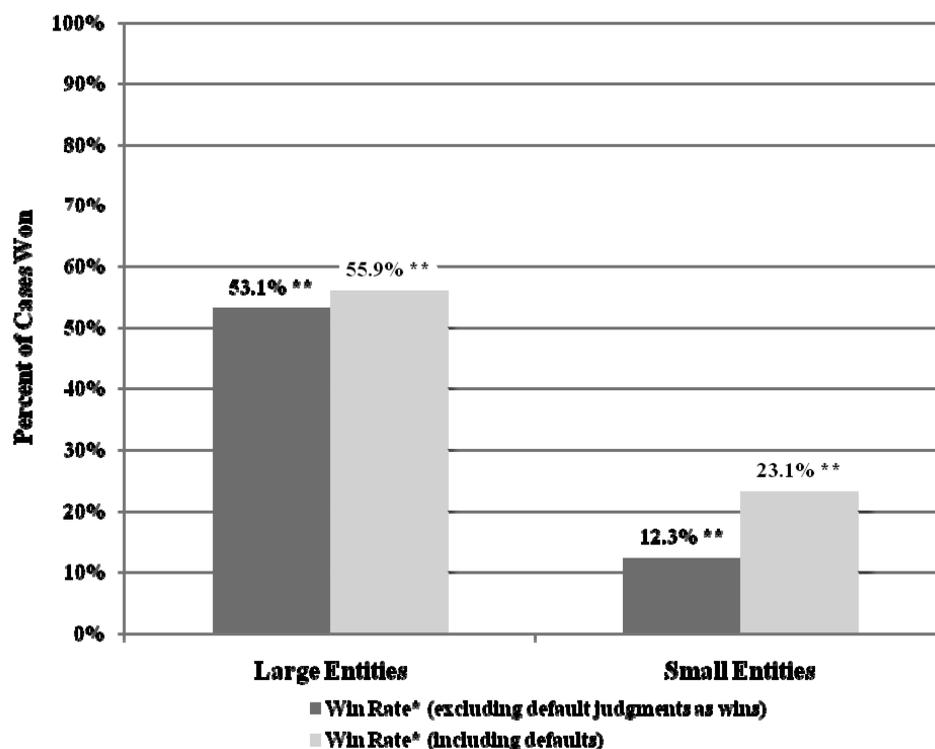
When the cases do not settle, large patent plaintiffs are significantly more likely than small ones to win, regardless of how the data are sliced. When we combine the two data sets, large-entity plaintiffs win 53.1% of the cases decided on the merits (55.9% if default judgments are included), but small-entity plaintiffs win only 12.3% of their cases (23.1% if default judgments are included). These differences are highly statistically significant. Adding settle-

Table 7. Large-Entity vs. Small-Entity Win Rates Across Both Most- and Once-Litigated Data Sets

	Large Entity (Without Default Judgments)	Small Entity (Without Default Judgments)	p-Value (Both Chi-sq. & Fisher's Exact)	Large Entity (Including Default Judgments)	Small Entity (Including Default Judgments)	p-Value (Both Chi-sq. & Fisher's Exact)
Settlements Included in Denominator	6.50%	1.40%	0	7.20%	2.90%	<0.01
Settlements Excluded from Denominator	53.10%	12.30%	0	55.90%	23.10%	<0.01

42. The number of assertions is the number of cases in which a patent is litigated, which in our study varies from a low of eight to a high of ninety-seven suits.

Figure 1. Large & Small Entities' Merits-Based Win Rates Once Settlement Fails (Most- and Once-Litigated Patent Cases Combined)



* Excludes settlements & consent judgments from baseline total of all cases

** p < .05

ments into the denominator naturally reduces the number of patentee wins but does not change the relationship: large entities win judgments in 6.5% of all cases in the combined data sets (7.2% if default judgments are included), compared to 1.4% of small entities (2.9% if default judgments are included). These differences are also highly significant.

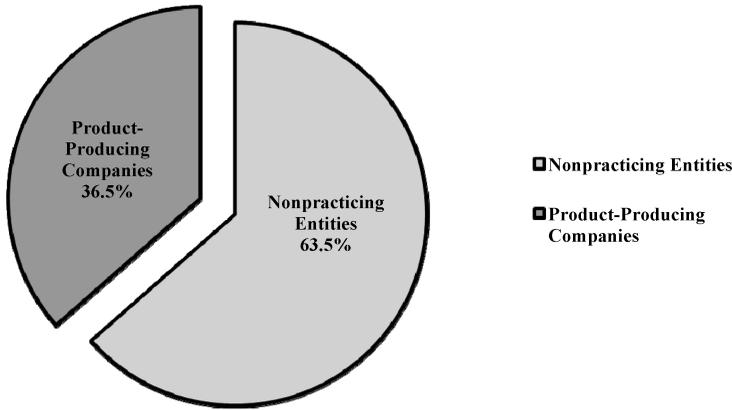
In short, patents originally issued to large entities do substantially better in litigation than those originally issued to small entities.

2. NPE Status

Because many patents are sold before being litigated,⁴³ however, entity size is at best an indirect proxy for the nature of the current plaintiff.⁴⁴ To get more

43. See, e.g., Moore, *supra* note 40 (finding that roughly half of patents enforced by corporations were purchased before suit, generally from an individual).

44. It may, however, reflect differences in the prosecution of patents between large and small entities.

Figure 2. Composition of Plaintiffs in Most-Litigated Patent Cases

directly at the question of whether different types of plaintiffs fare differently in court, we characterized each of the patent owners into one of the twelve entity-status categories described above. For statistical analysis purposes, we aggregate the patents into practicing entities (Class 8) and NPEs (all other classes except Class 10).

As with small entities, NPEs are much more prevalent among the most-litigated patents. Figures 2 and 3 show that 79.2% of the once-litigated patent plaintiffs were companies that made products, while only 36.5% of the most-litigated patents were owned by product-producing companies. And if we look not just at the patents themselves, but at the number of patent-suit pairs of each patent in litigation, the differences are even more dramatic: just 16.7% of the assertions of the most-litigated patents were made by product-producing companies. It is not merely that NPEs are more likely to own the most-litigated

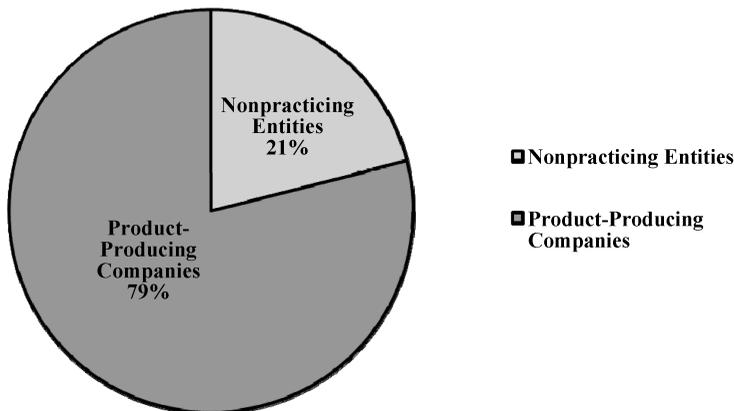
Figure 3. Composition of Plaintiffs in Once-Litigated Patent Cases

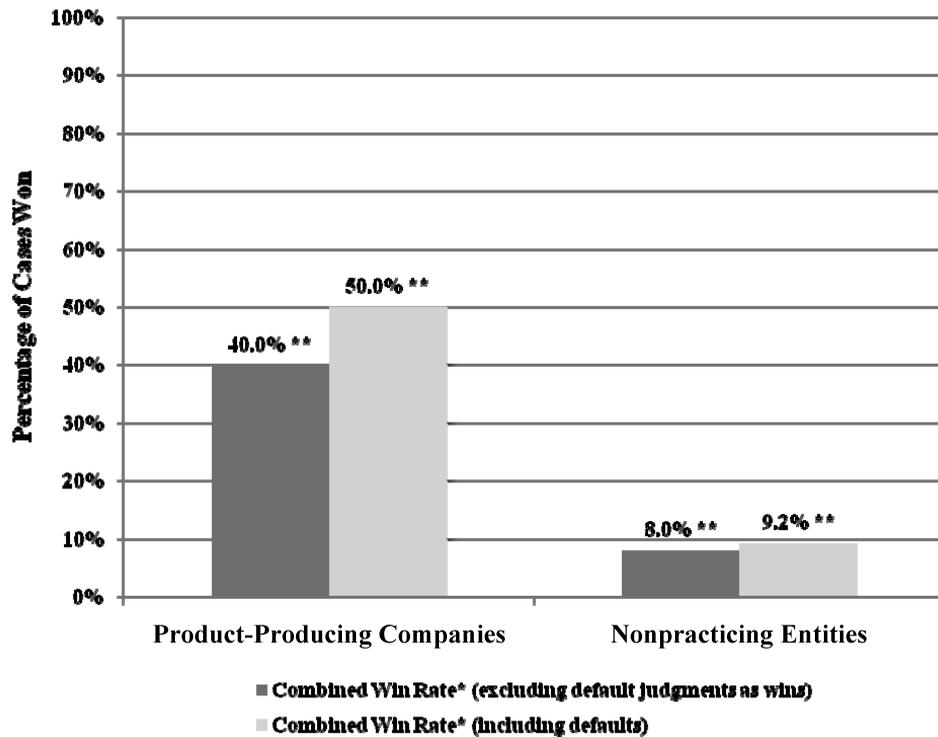
Table 8. Product Company vs. NPE Win Rates Across Both Most- and Once-Litigated Data Sets

	Product Company (Without Default Judgments)	NPE (Without Default Judgments)	p-Value (Both Chi-sq. & Fisher's Exact)	Product Company (Including Default Judgments)	NPE (Including Default Judgments)	p-Value (Both Chi-sq. & Fisher's Exact)
Settlements Included in Denominator	4.50%	0.80%	0	6.70%	1.00%	0
Settlements Excluded from Denominator	40.00%	8.00%	0	50.00%	9.20%	0

patents; they also file more lawsuits on those patents than their product-producing counterparts in the most-litigated set.

As for outcomes, the results, as seen in Table 8 and Figure 4, are dramatic—no matter how the data are sliced, product-producing entities are far

Figure 4. Plaintiffs' Merits-Based Win Rates Once Settlement Fails (Most- and Once-Litigated Patent Cases Combined)



* Excludes settlements & consent judgments from baseline total of all cases

** p < .05

Table 9. Product Company vs. NPE Settlement Rates Across Both Most- and Once-Asserted Data Sets

	Settlement Rate	p-Value (Both Chi-sq. & Fisher's Exact)
Product Companies	86.60%	0.15
NPEs	89.60%	

more likely to win their cases than NPEs.⁴⁵

If we consider just patent-owner wins and defendant wins on the merits, product owners win 40% of their cases across both the most-litigated and once-litigated data sets, while NPEs win only 8%. If we include default judgments, product-producing companies win 50% of their cases, while NPEs win only 9.2%. Each of these results is highly statistically significant.

As with our other results in this Article, including settlements in the denominator reduces the percentage of wins but does not change the relationship. Once settlements are included, product-producing companies win judgments in 4.5% of their suits, while NPEs win judgments in only 0.8% of their suits. Adding default judgments changes these numbers to 6.7% for product-producing companies and 1.0% for NPEs.⁴⁶

Given this dramatic difference in win rates, it is quite striking that NPEs are not significantly more likely to settle their lawsuits than are product companies. As shown in Table 9, product companies settled 86.6% of their cases. While NPEs settled more (89.6%), the difference is not statistically significant.

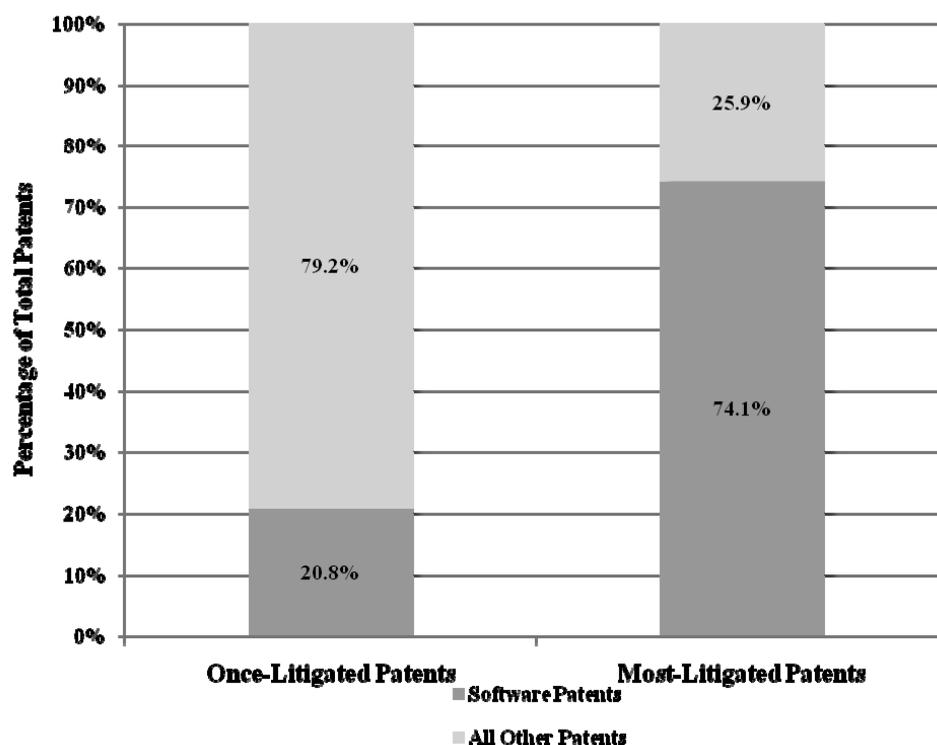
Thus, it appears that NPEs are not as worried about losing as they should be; they take cases to judgment rather than settle them even though they are very unlikely to win those cases. This is particularly surprising given that NPEs are interested primarily in money damages, not excluding a competitor, so they are less likely to have the sort of asymmetric-stakes case that prior literature predicts might not settle.⁴⁷ At the very least, it is surprising that product

45. One scholar has found, to the contrary, that NPEs are no more likely to lose their cases than other types of plaintiffs. See Sannu K. Shrestha, Note, *Trolls or Market-Makers? An Empirical Analysis of Nonpracticing Entities*, 110 COLUM. L. REV. 114, 147–48 (2010). Shrestha also used data from the Stanford IP Litigation Clearinghouse. However, because Shrestha chose a nonrandom sample of NPE cases based on companies reported in the press as NPEs, representing only a small fraction of NPEs, we think the most likely explanation for the difference in our outcomes is selection bias in Shrestha's data. It should not be surprising that companies known to the world as NPEs are known precisely because they fare better than average in litigation. It is also notable that Shrestha excluded Federal Rule of Civil Procedure 12(b)(6) dismissals from his count of results on the merits.

46. Each of these results is highly significant (without default judgments, the p-value is 0.000; with default judgments, the p-value is 0.000).

47. See generally George L. Priest & Benjamin Klein, *The Selection of Disputes for Litigation*, 13 J. LEGAL STUD. 1, 24–29 (1984) (discussing that parties are less likely to settle where there are asymmetric stakes such that one party stands to lose more than the other would gain in a money judgment, for example in an antitrust case).

Figure 5. Software vs. Nonsoftware Inventions Among the Most- and Once-Litigated Patents



companies and NPEs settle at the same rate given their very different win rates in the cases that do not settle. We discuss some possible explanations for that divergence in section III.A.

C. INDUSTRY AND TECHNOLOGY DIFFERENCES

We also divided the data set by industry and technology. Small numbers of litigated cases in many of the industry and technology categories in our data sets required us to make broader comparisons. Moreover, it is the patents covering one type of invention—software—that are one of the most striking features of our study. In our taxonomy, software is a *technology* that is part of the computer *industry*. However, because far more observers are interested in the world of software patenting than of computer-industry patenting more generally, we report here on differences between litigation involving software patents and litigation involving all other types of patents.

The overrepresentation of software patents in the most-litigated set is quite remarkable. Figure 5 shows that software patents constituted 20.8% of the

Table 10. Software Patents vs. Nonsoftware Patents Win Rates Across Both Most- and Once-Litigated Data Sets

	Software Patents (Without Default Judgments)	Non-software Patents (Without Default Judgments)	p-Value (Both Chi-sq. & Fisher's Exact)	Software Patents (Including Default Judgments)	Non-software Patents (Including Default Judgments)	p-Value (Both Chi-sq. & Fisher's Exact)
Settlements Included in Denominator	1.40%	4.00%	0.01	1.40%	7.20%	0
Settlements Excluded from Denominator	12.90%	37.10%	<0.01	12.90%	51.10%	0

once-litigated patents but 74.1% of the most-litigated patents.⁴⁸ And again, these differences are even more dramatic if we measure not patents, but patent-suit pairs. Software patents accounted for 93.7% of the assertions of the most-litigated patents.

The outcomes shown in Table 10 are equally dramatic. No matter how we test it, owners of non-software patents are far more likely to win their cases than are software patent owners. If we consider just patent owner wins and defendant wins on the merits, nonsoftware patent owners win 37.1% of their cases across both the most-litigated and once-litigated data sets, while software patentees win only 12.9%. If we include default judgments, nonsoftware patent owners win 51.1% of their cases, while software patentees win only 12.9%. Each of these results is highly statistically significant.

As with other data, including settlements in the denominator reduces the percentage of wins but does not change the relationship.⁴⁹ Once settlements are included, nonsoftware patent companies win judgments in 4% of their suits, while software patentees win judgments in only 1.4% of their suits. Adding default judgments changes these numbers to 7.2% for nonsoftware patent owners and 1.4% for software patentees. Again, each result is highly statistically significant.

These results appear to be driven by differences in the most-litigated patent set. When we disaggregated the most-litigated and once-litigated data sets, the differences in all four tests remained highly statistically significant for the most-litigated set.⁵⁰ In other words, repeat software patent litigants are particularly unlikely to win their cases.

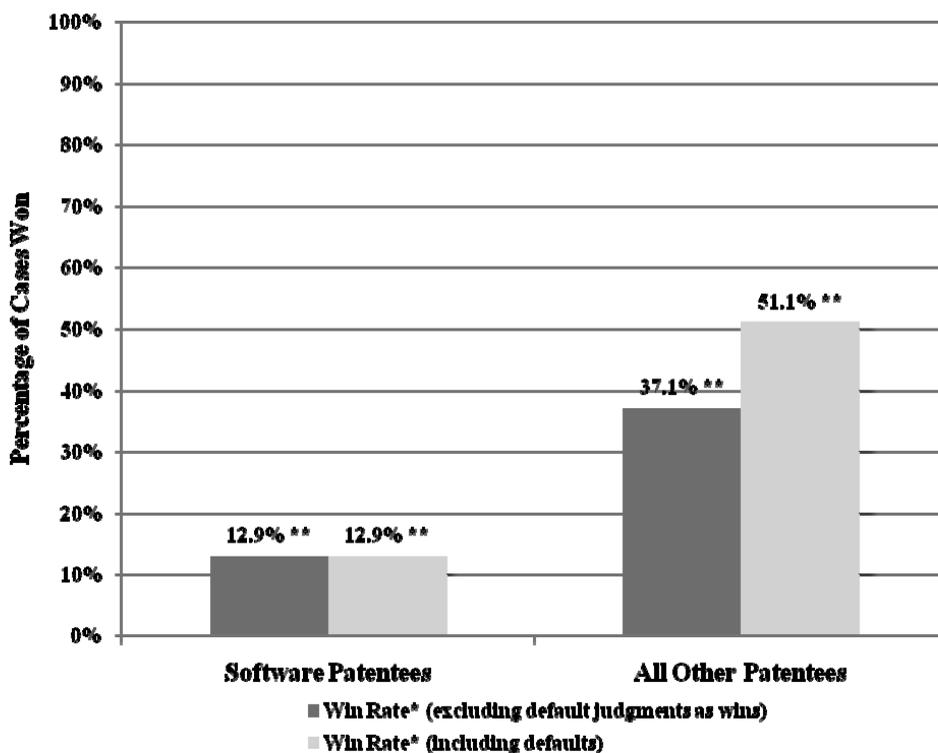
Like NPEs, software patentees are slightly more likely to settle than nonsoft-

48. The descriptive statistics reported in this section differ somewhat from those reported in our prior article, Allison, Lemley & Walker, *supra* note 23, because for this Article we are considering only cases that went to settlement or judgment, and we consider the resolution of each patent a separate event. For this Article, we also increased the sample size for once-litigated patents.

49. See *supra* section II.B.2.

50. Within the most-litigated set, the p-value is 0.035 without settlements, and the p-value is 0.011 with settlements, regardless of whether default judgments were included.

Figure 6. Software vs. Nonsoftware Overall Win Rates (Most- and Once-Litigated Patent Cases Combined)



ware patent owners. Table 11 shows that software patent owners settle 89.5% of their cases across both data sets, compared with 86.0% of nonsoftware patent owners. Unlike the results for NPEs, this difference is (barely) statistically significant at the 10% level. But even so, it does not represent much of an accommodation to the dramatic difference in win rates.

D. LOGISTIC REGRESSION

We also conducted a logistic regression in an effort to measure what influences outcomes when accounting for interactions (correlations) among different independent variables such as entity size, plaintiff type, patent type, and patent characteristics. There were too few patent owner wins in the overall data set to achieve reliable results when regressing these independent variables on patentee wins, but we were able to obtain more robust results by regressing on the more numerous defendant wins. Regressing on settlements proved most satisfactory because these represented the most common outcome and consequently had the largest number of observations.

Another valuable contribution of the logistic regression model was that it permitted us to cluster standard errors on individual patents, avoiding the

Table 11. Software vs. Nonsoftware Patent Owners Settlement Rates Across Both Most- and Once-Litigated Data Sets

	Settlement Rate	p-Value (Both Chi-sq. & Fisher's Exact)
Software Patent Owners	89.50%	0.09
Nonsoftware Patent Owners	86.00%	

potential problem that our results could be affected by some outcomes for most-litigated patents not being totally independent because the same patents with identical characteristics appeared multiple times.⁵¹ While there are such effects, the logistic regression model shows that our results are robust even when taking those effects into account.

When we ran logistic regressions to explain plaintiff wins on the merits (not including default judgments), significant explanatory factors for plaintiff wins were that the plaintiff was a large entity, and curiously, the average number of defendants per case. When we included default judgments, neither entity size nor the number of defendants remained a significant predictor of plaintiff wins. Instead, factors that significantly predicted a plaintiff win were producing a product and not asserting a software patent.

We urge caution in reliance on the logistical regressions here. Doing multivariate analysis on a highly selected group is necessarily fraught with peril. The first problem arises from the relatively small numbers of plaintiff wins—a change that ordinarily would be rather minor, such as adding default judgment wins, produces completely different regression results. Second, the nature of the most-litigated patents likely obscures significant results in this case because of correlations that prevent us from completely disaggregating important variables. Because so many of the most-litigated patents are (1) software patents (2) owned by small entities that (3) do not produce a product and (4) lose their suit, the multivariate model has a difficult time distinguishing those explanatory variables. We can say with confidence that characteristics 1 through 3, taken together, predict a defendant win, but we cannot confidently disaggregate the explanatory variable further.

We hoped to achieve more robust results by regressing on defendant wins because there are far more of them than plaintiff wins. We regressed on defendant wins both with and without settlements in the denominator. In other words, we looked for variables that would predict defendant wins once with settlements included in the denominator and once with them excluded. When settlements are included in the denominator, a predictor of defendant wins is not necessarily a predictor of plaintiff losses because settlement remains a third

51. For example, the Katz patents are asserted in multiple suits; invalidating a Katz patent in one suit may lead to invalidation in other suits.

alternative outcome. When settlements are excluded, however, there are only two possible outcomes and a predictor of defendant wins is a predictor of plaintiff losses. The results of doing this both ways are unfortunately not consistent because the number of settlements overwhelms the numbers of both plaintiff and defendant wins.⁵²

With settlements included in the denominator, we found that the average number of defendants per case was a significant predictor of defendant wins. That is, the more defendants there are, the more likely they are to win. This finding is seemingly inconsistent with what we found when regressing on plaintiff wins without default judgments, where the number of defendants was a significant predictor of plaintiff wins. In fact, however, the two results can be reconciled: an increase in the number of defendants per case reduces the chance of settlement, making both plaintiff and (especially) defendant wins more likely. We say more about this in section III.A.1.

When we then regress on defendant wins with settlements excluded from the denominator, we get cleaner results because there are only two possible outcomes. Here, we find that the assertion of a software patent is a very significant predictor of defendant wins, and thus, of plaintiff losses—yet further confirmation of how poorly software patents fared when going to judgment across both data sets. We also find that when the plaintiff asserts a patent originally issued to a large entity, defendants are more likely to lose—large-entity status is a significant predictor of plaintiff wins.

Our final set of regressions seeks predictors of settlement. The number of defendants per case is a *negative* predictor of settlement. That is, the more defendants there are per case, the less likely the case is to settle. Recall, as well, that we have some evidence from the win-rate regressions that, if the case does not settle but instead goes to judgment, the more defendants there are per case the more likely those defendants are to win. And, because of the much larger number of observations, the result from this regression is almost certainly more trustworthy than the inconsistent finding when plaintiff wins (without default judgments) was the dependent variable. Such evidence may speak against the strategy of suing many defendants in the same case because it is more difficult to settle with those defendants, and if the patent owner does not settle, those defendants are more likely to win.⁵³

52. We did consider employing multinomial regression with three dependent variables—plaintiff win, defendant win, and settlement—instead of a logistic regression. Because of the nature of our data and the small numbers of plaintiff wins, however, we would not have achieved results that were any more satisfactory.

53. Remember, though, that because settlements are in the denominator, the fact that defendants are more likely to win when there are more of them per case does not necessarily mean that plaintiffs are more likely to lose. When we regressed directly on plaintiff wins, the number of defendants in a case was also a significant predictor of plaintiff wins.

III. EXPLANATIONS AND IMPLICATIONS

What explains these surprising results? In this Part, we explore some possible reasons why the most-litigated patents turn out to be weaker than other-litigated patents, as well as some reasons why the owners of those patents nonetheless seem willing to let the cases go to judgment more often than they should—they do settle more, just not nearly often enough.

A. OTHER EXPLANATIONS FOR SETTLEMENT DATA

As we found in the previous Part, our hypothesis that the patentees who file the most suits will be more likely to settle is borne out, but to a surprisingly modest extent. In this section, we explore some other factors that might complicate the hypothesized relationship between repeat play and settlement.

1. Number of Defendants

One possible explanation is that the settlement results are being driven by defendants rather than plaintiffs in the most-litigated cases. The most-litigated patents are generally asserted against multiple defendants in each suit; perhaps the presence of multiple defendants gives the patentee less control over settlement because even one obstinate defendant who refuses to settle (possibly in the hopes of building a reputation that will deter future plaintiffs) can force the case to judgment. On this theory, the more defendants in a case, the more likely it should be to go to judgment with respect to at least one of those defendants.

In the abstract, there are reasons to be skeptical of this “defendant control” theory. To begin, the economic literature suggests the opposite—that invalidation of a patent is a public good because the defendant that takes a case all the way to judgment shoulders all the risk of losing the case, but must share the benefit of invalidating the patent with all of its competitors.⁵⁴ Thus, if anything, we should expect that defendants in multiparty patent cases should be more likely to settle out and leave their competitors holding the bag, particularly because while defendants can share information, they cannot act jointly in deciding to settle.⁵⁵ Similarly, because the patentee faces the risk of losing the patent itself if the case does not settle, it is reasonable to expect that the patentee will often be willing to accept a lower payment in settlement from each individual defendant. This is particularly true if the defendant has discovered a strong argument against the patent, such as a killer piece of prior art. When that

54. Joseph Farrell & Robert P. Merges, *Incentives to Challenge and Defend Patents: Why Litigation Won't Reliably Fix Patent Office Errors and Why Administrative Patent Review Might Help*, 19 BERKELEY TECH. L.J. 943, 952 (2004); Joseph Scott Miller, *Building a Better Bounty: Litigation-Stage Rewards for Defeating Patents*, 19 BERKELEY TECH. L.J. 667, 679–80 (2004); Joseph Scott Miller, *Joint Defense or Research Joint Venture: Reassessing the Patent-Challenge Bloc's Antitrust Status*, 2011 STAN. TECH. L. REV. (forthcoming), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1652188.

55. *Jones Knitting Corp. v. Morgan*, 361 F.2d 451, 459 (3d Cir. 1966) (sharing settlement authority violates the antitrust laws).

Table 12. Number of Defendants for Most-Litigated Set

N	791
Mean	5.19
Median	2
Mode	1
Std. Deviation	7.96
Variance	63.4
Range	68
Minimum	1
Maximum	69

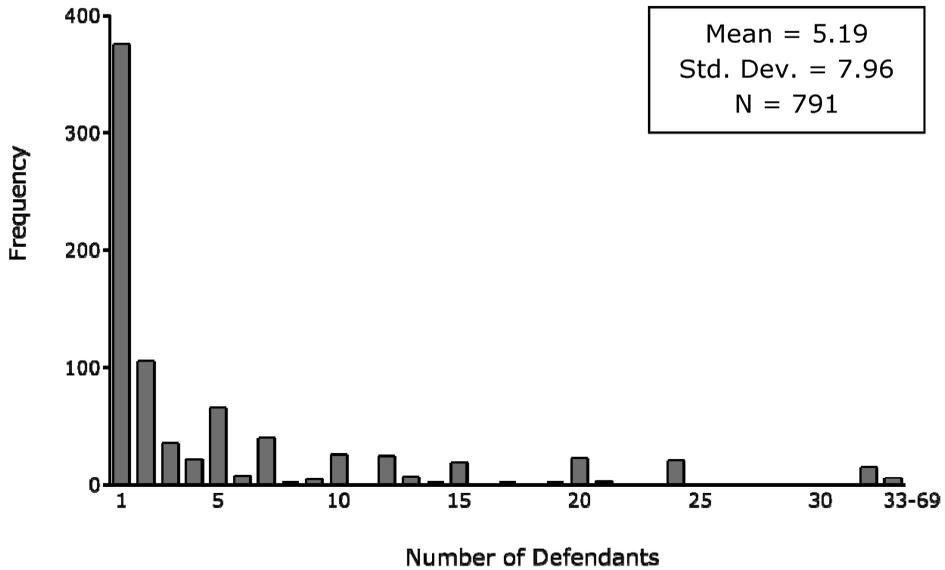
Table 13. Number of Defendants for Once-Litigated Set

N	343
Mean	2.13
Median	1
Mode	1
Std. Deviation	3.46
Variance	12
Range	51
Minimum	1
Maximum	52

happens, it is in the interest of both the plaintiff and the defendant to give that particular defendant a confidential “sweetheart” deal that keeps the argument out of court and keeps the patent alive.⁵⁶ Indeed, there is no reason the patentee in such a case could not simply drop one determined defendant from the suit, before or after settling with others.

We collected data on the number of defendants per suit in the most-litigated and once-litigated sets. We report the results in Tables 12–13 and Figures 7–8. Overall, both the most-litigated and the once-litigated sets showed a strongly skewed distribution, with the modal number of defendants in both sets being 1 and the median number of defendants being 2 in the most-litigated set and 1 in

56. If prior settlement agreements have “most-favored nation” clauses, sweetheart deals may not be feasible. We have no way to assess how common those clauses are, though for the same reason patentees treat settlement agreements as confidential we would expect serial patent litigants to object to the inclusion of such clauses. Our anecdotal sense is that patent settlements are not uniform, but vary depending on the defendant.

Figure 7. Histogram of Most-Litigated Set

the once-litigated set. But the most-litigated patents had a much higher average number of defendants (5.2 versus 2.1), driven by a few outlier suits against as many as 69 different defendants.

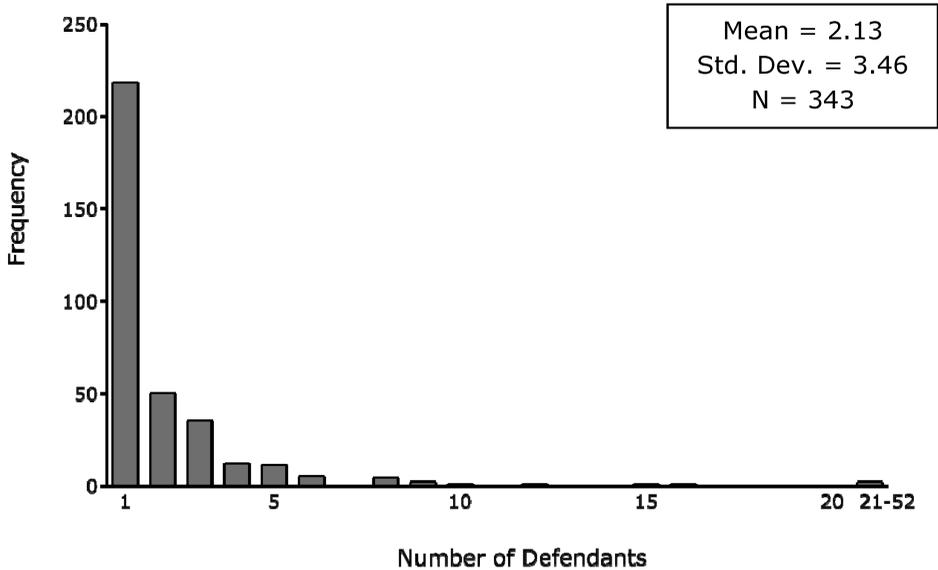
The Lex Machina data allow us to test the effect of multiple defendants on the likelihood of settlement. The results are mixed. When we test the effect of multiple defendants on the likelihood of settlement separately in the most-litigated and once-litigated sets, we find no statistically significant effect. When we combine the two data sets, however, the result is statistically significant: cases with more defendants are less likely to settle.⁵⁷ We report these results in Table 14, where we see that the number of defendants per case is a significant *negative* predictor of settlement.

This data could be consistent with the multiple-defendants explanation because, on average, the most-litigated patents have more defendants per case than the once-litigated patents. But the results are statistically significant only when we combine the data sets, meaning that we cannot draw that conclusion as a matter of statistics.⁵⁸

Even if the larger number of defendants causes cases to be less likely to settle, we are still left with a puzzle. For we also find that the more defendants

57. That is consistent with the results we found in the logistic regression. One can also see in Table 15 that the adjusted number of forward citations is also a significant positive predictor of settlement; for example, the more forward citations the plaintiffs' patents receive, the more likely the case is to settle.

58. This is likely a function of the size of our data set, though we cannot prove that.

Figure 8. Histogram of Once-Litigated Set**Table 14. Effect of Number of Defendants on Likelihood of Settlement Across Both Data Sets**

	Coef.	Robust Std. Err.	Z	P> z	[95% Conf. Interval]	
Number of Defendants	-.0660679	.0155329	-4.25	0.000	-.0965119	-.0356239
Adjusted Number of Forward Citations	1.108155	.4473482	2.48	0.013	.2313682	1.984941
Constant	1.864647	.5993802	3.11	0.002	.6898835	3.039411

are named, the more likely the defendants are to win the case.⁵⁹ We present those results in Table 15, where we see that the number of defendants per case is a significant *positive* predictor of defendant wins.

Thus, a plausible explanation for at least some of the results is that when a patentee sues multiple defendants, the patentee loses control over the case, being forced to trial or judgment even in cases it would prefer to settle, and which it is more likely to lose.

59. The multivariate regression here produced inconsistent results, suggesting that both plaintiff wins and defendant wins are more likely as the number of defendants increases. This seemingly contradictory result likely results from the fact that cases with more defendants are less likely to settle, meaning that there are both more plaintiff wins and more defendant wins in those cases. One can see in Table 15 that the adjusted number of forward citations is also a significant negative predictor of defendant wins, for example, the more forward citations the plaintiffs' patents receive the less likely defendants are to win.

Table 15. Effect of Number of Defendants on Likelihood of Defendant Win Across Both Data Sets

	Coef.	Robust Std. Err.	Z	P> Z	[95% Conf. Interval]	
Number of Defendants	.0752602	.0187087	4.02	0.000	.0385919	.1119285
Adjusted Number of Forward Citations	-1.478648	.5050783	-2.93	0.003	-2.468584	-.488713
Constant	-1.883341	.6104311	-3.09	0.002	3.079764	-.6869182

But this just pushes the puzzle back a level. If patentees lose control over their cases when they sue multiple defendants, why do they choose to sue multiple defendants in the same suit?⁶⁰ One explanation may be cost. Litigation is expensive, and the plaintiff may minimize its costs by filing multiple lawsuits. Having made this choice, plaintiffs could still choose to drop their cases against particular defendants, or settle those cases for a nominal sum, rather than go to judgment and risk invalidation. And according to the data, that is precisely what they should do. But they do not seem to.

2. License Agreements and Plaintiff Stakes

Key to our hypothesis that repeat patent litigants should be more inclined to settle is the assumption that the number of different lawsuits is evidence of how much the patentee has to lose if the patent is invalidated or construed narrowly. This assumption strikes us as a reasonable first-order approximation because we can identify concrete things the plaintiff has to lose if the patent is invalidated. But there may be other evidence we cannot see that changes the risk-aversion calculus. Most notably, if a patentee in the once-litigated patent set has an established set of license agreements with an ongoing royalty, and has happened to sue the only user who did not take a license, that patentee has an unobservable, but quite significant, stake in not having the patent invalidated that may make it risk averse. Similarly, we do not know whether serial patent litigants have an established licensing practice outside of the lawsuits. If there is selection bias in the unobservable settlement of disputes that never make it to litigation, it could affect our results.

To affect our findings, though, it must be true not only that some once-litigated patents have hidden licenses—and therefore unseen stakes beyond the outcome of the instant case—but also, that once-litigated patents are systematically more likely than the most-litigated patents to have such licenses, to such

60. If anything, the law encourages suits against unrelated defendants for infringing the same patent to be separated. *See, e.g.*, Order Severing Parties Due to Misjoinder and Dismissing All But the First Named Defendant, *Finisar Corp. v. Source Photonics, Inc.*, No. C 10-00032 WHA (N.D. Cal. May 5, 2010) (*sua sponte* order severing unrelated defendants in a patent case due to misjoinder).

an extent that those hidden licenses swamp the effect a loss would have on the pending suits of the most-litigious patentees. We can think of no way to test this with available information; licensing agreements outside of litigation are almost always confidential.⁶¹ But one plausible reason that it might be true lies in the reality that the patentee filed so many suits. Although the difference between filing one suit and multiple suits may reflect the value of the patent, the necessity of using the patent, or the number of competitors, it might also reflect the willingness of companies to license the patent from the plaintiff. If two companies, each with twenty infringers, send license demands to all twenty companies, the company that gets paid without litigation will file relatively few lawsuits, whereas the company whose license efforts are rebuffed will have to file more lawsuits if it hopes to collect revenue.⁶²

A second possible explanation has to do with the dependence of plaintiff stakes on the structure of the licenses entered into. If the parties enter into a royalty-bearing license, the obligation of the licensee to pay royalties ends when the patent is held invalid.⁶³ By contrast, if a company pays a lump sum to license a patent, it is far from clear that it can get that money back merely because the patent is later held invalid.⁶⁴ And while the confidentiality of the data prevents a rigorous analysis, our experience has been that settlements of patent lawsuits more commonly involve lump-sum payments rather than ongoing royalties. As a result, a serial patent plaintiff that has settled a number of its cases may feel more comfortable litigating the rest because it probably will not be forced to give the money back even if it loses.

We emphasize that we cannot test any of this empirically. It is all speculation. But it does serve to point out that the relationship between patterns of litigation and party stakes is more complex than we might at first assume.

B. WEAK LITIGATED PATENTS

More surprising to us is the finding that the most-litigated patents that do make it to judgment fare so poorly. Are there structural reasons to think the most-litigated patents are likely to be weaker and, therefore, most likely to be rejected? And if so, what does this say about the evidence economists use to assess the value of patents?

61. *But see* Mark A. Lemley & Nathan Myhrvold, *How To Make a Patent Market*, 36 HOFSTRA L. REV. 257 (2008) (arguing that licensing agreements outside of litigation should not be confidential).

62. Strictly speaking, that company will merely have to sue more defendants; it could sue them all in the same case. But our data suggest that those who sue multiple defendants also tend to file multiple lawsuits. The plaintiffs who filed the most lawsuits also sued more defendants per case on average—5.2 defendants compared with 2.1 in the once-litigated set.

63. *See Lear, Inc. v. Adkins*, 395 U.S. 653, 673–74 (1969).

64. *See, e.g., Studiengesellschaft Kohle v. Shell Oil Co.*, 112 F.3d 1561, 1568 (Fed. Cir. 1997) (pre-*MedImmune*, concluding that a licensee could not recover royalties paid before a patent challenge); *MedImmune Inc. v. Genentech Inc.*, 535 F. Supp. 2d 1000, 1016 (C.D. Cal. 2008) (noting, but not resolving, the issue). As a result, lawyers and scholars have argued that patent owners should license for lump sums rather than ongoing royalties whenever possible. *See* Michael Risch, *supra* note 10, at 1017–18.

1. Multiple Lawsuits As Evidence of Obviousness?

The most-litigated patents are enforced against a wide array of companies—an average of 5.2 defendants per *suit* per patent, with some patents involved in as many as ninety-seven different suits. Although in theory that could result from a pattern of widespread copying of a particular technology, other work suggests that patent plaintiffs are overwhelmingly suing independent developers, not those they accuse of copying their ideas.⁶⁵ Thus, one reasonable inference from the filing of multiple lawsuits against multiple defendants is that the technology in question is one that was widely adopted by an industry as a result of widespread, near-simultaneous invention.

Near-simultaneous invention is sometimes used as evidence that an invention is obvious.⁶⁶ After all, if most of the companies interested in a particular field independently develop the same technology at about the same time, that is pretty good evidence that the technology was within the capability of one of ordinary skill in the art. So one hypothesis is that the enforcement of a patent against an entire industry is evidence of simultaneous invention, and therefore, of obviousness.

We do not believe this can explain our results, however. First, we discovered in our prior article that the most-litigated patents result from long chains of continuation applications—more than four applications on average⁶⁷—and, therefore, were based on applications filed well before the lawsuits.⁶⁸ As a result, the patents in our study were more likely than average to have been invented long before the rest of the industry adopted the technology. Second, and more important, the patentee losses in the most-litigated patents data set were overwhelmingly not based on findings of obviousness. They are split between findings of invalidity, noninfringement, and unenforceability; those based on invalidity are more likely to be findings of on-sale bar or inadequate written description than findings of obviousness. Although there are more findings of invalidity than noninfringement (sixty of the eighty-six cases in the most-litigated set involved findings of invalidity, forty-four involved findings of noninfringement, and three involved findings of unenforceability), the most common grounds for invalidation of the most-litigated patents were written

65. Christopher A. Cotropia & Mark A. Lemley, *Copying in Patent Law*, 87 N.C. L. REV. 1421, 1459 (2009). This is particularly true in the information technology industries in which the vast majority of the most-litigated patents exist.

66. *Hybritech, Inc. v. Monoclonal Antibodies, Inc.*, 802 F.2d 1367, 1380 n.4 (Fed. Cir. 1986) (noting the use of this factor, but not applying it in the case at hand). For discussion of secondary considerations of non-obviousness, see Rochelle Cooper Dreyfuss, *The Federal Circuit: A Case Study in Specialized Courts*, 64 N.Y.U. L. REV. 1, 9–10 (1989), and Robert P. Merges, *Commercial Success and Patent Standards: Economic Perspectives on Innovation*, 76 CALIF. L. REV. 803, 816 (1988). For an argument that simultaneous invention should be used more often, see Mark A. Lemley, *Should Patent Infringement Require Proof of Copying?*, 105 MICH. L. REV. 1525, 1534–35 (2007).

67. Allison, Lemley & Walker, *supra* note 23, at 28.

68. *Cf.* Allison & Lemley, *supra* note 1, at 237 (finding that even the average-litigated patent spent 12.3 years from filing of the application to resolution of the case).

description, anticipation, and on-sale bar.⁶⁹

2. Overclaiming?

Because serial patent litigants have sued so many defendants, it might suggest that they are overclaiming. Perhaps what distinguishes these cases from the once-litigated cases is that the patentees have either obtained a patent that is too broad or that they are overreading the claims they received. If so, the patent in question would be more likely to be found either invalid (in the first case) or not infringed (in the second). The fact that the patentees in this data set filed so many continuation applications may also support the overclaiming hypothesis because one use of continuation applications is to change patent claims over time to cover technologies developed after the original application was filed.⁷⁰ Doing so does not necessarily violate the patent laws, but it is far more likely to involve overclaiming because the patentee is attempting to extend the claims to something it did not have in mind when it filed its original patent application.

It is hard to know whether the patentees in these cases are more likely than others to involve overclaiming, though the explanation is plausible.⁷¹ To measure overclaiming, we would need to know what the “right” scope of the patent was in a Platonic sense, as well as whether the court was right to find the patent invalid or not infringed.⁷²

3. Are Software Patents or Trolls the Problem?

One of the most striking findings is the weakness of software and NPE-owned patents in the overall data set. Given the enormous percentage of the cases we review that involve software patents, especially in the most-litigated set, it seems likely that software patents are dragging down the averages. After all, 134 of the 421 patents we evaluate in both data sets, or 31.8%, involve software. There have been numerous complaints about the quality of software

69. These numbers add up to more than eighty-six cases because in many cases the court ruled on more than one ground.

70. See, e.g., Mark A. Lemley & Kimberly A. Moore, *Ending Abuse of Patent Continuations*, 84 B.U. L. REV. 63, 107 (2004) (“One of the most egregious abuses of continuation applications described above is the use of the process to change patent claims to track inventions first made by one of the applicant’s competitors.”).

71. See Ronald J. Mann & Marian Underweiser, *A New Look at Patent Quality: Relating Patent Prosecution to Validity* 8 (Columbia Law & Econ. Working Paper No. 381, 2010), available at <http://ssrn.com/abstract=1671784> (finding that patents with more claims are more likely to be invalidated, and suggesting that the owners of those patents may be more likely to overclaim them). Mann and Underweiser also find that filing continuation applications is negatively correlated with validity. *Id.* at 18.

72. We cannot distinguish overclaiming from other explanations based on the reason the patent was rejected because an overbroad claim may be invalid either under 35 U.S.C. § 112 for lack of description or enablement or under 35 U.S.C. §§ 102 and 103 because it is written so broadly as to encompass the prior art.

patents; our data may give some empirical support to those assertions.⁷³ If we consider just patent-owner wins and defendant wins on the merits, nonsoftware patent owners win 37.1% of their cases across both the most-litigated and once-litigated data sets, whereas software patentees win only 12.9%. If we include default judgments, nonsoftware patent owners win 51.1% of their cases, whereas software patentees win only 12.9%.⁷⁴

Something similar can be said about suits brought by NPEs. NPE suits, like software suits, constitute a large percentage of the most-litigated cases; they represent 730 of the 1134 cases in our combined data set, or 64.3%. If we consider just patent-owner wins and defendant wins on the merits, product owners win 40% of their cases across both the most-litigated and once-litigated data sets, while NPEs win only 8%. If we include default judgments, product-producing companies win 50% of their cases, whereas NPEs win only 9.2%.⁷⁵

The authors have elsewhere expressed skepticism over efforts to eliminate particular types of patents,⁷⁶ and one has argued that we should not single out patent trolls for special treatment.⁷⁷ But it is important to recognize that software patents and patents asserted by NPEs are both taking disproportionate resources in patent litigation and that the social benefit from those cases appears to be slight.

It is less clear what conclusions we should draw from this fact. Perhaps judges (and juries, though most of the outcomes were the result of judicial decisions) simply do not like either NPEs or software patents. After all, both have been subject to substantial criticism by both legal scholars and the popular press.⁷⁸ Our results could accordingly be evidence of antisoftware or anti-NPE

73. See, e.g., Julie E. Cohen, *Reverse Engineering and the Rise of Electronic Vigilantism: Intellectual Property Implications of "Lock-Out" Programs*, 68 S. CAL. L. REV. 1091, 1177–80 (1995); Peter S. Menell, *A Method for Reforming the Patent System*, 13 MICH. TELECOMM. & TECH. L. REV. 487, 506 (2007); Pamela Samuelson et al., *A Manifesto Concerning the Legal Protection of Computer Programs*, 94 COLUM. L. REV. 2308, 2425 (1994).

74. Each of these results is highly significant.

75. Each of these results is highly significant.

76. See, e.g., John R. Allison & Starling D. Hunter, *On the Feasibility of Improving Patent Quality One Technology at a Time: The Case of Business Methods*, 21 BERKELEY TECH. L.J. 729, 789 (2006); John R. Allison & Ronald J. Mann, *The Disputed Quality of Software Patents*, 85 WASH. U. L. REV. 297, 334 (2007).

77. Mark A. Lemley, *Are Universities Patent Trolls?*, 18 FORDHAM INTELL. PROP. MEDIA & ENT. L.J. 611, 630 (2008).

78. On patent trolls, see, for example, Robin M. Davis, Note, *Failed Attempts To Dwarf the Patent Trolls: Permanent Injunctions in Patent Infringement Cases Under the Proposed Patent Reform Act of 2005 and eBay v. MercExchange*, 17 CORNELL J.L. & PUB. POL'Y 431, 438 (2008) ("Most patent trolling behavior thrives on the inequities of enforcing patent rights without contributing anything to either the invention or production of new technologies."), Jason Kirby, *Patent Troll or Producer?*, NAT'L POST (Toronto), Jan. 14, 2006, available at <http://www.financialpost.com/story.html?id=1509d361-0144-4432-b6dc-2c14026c98d6> ("Companies who do the costly grunt work of actually developing and marketing new technologies are being held ransom by tiny outfits whose only assets are 'kooky and vague' patents . . ."), and Joe Beyers, *Rise of the Patent Trolls*, CNET NEWS, (Oct. 12, 2005, 4:00 AM), http://news.cnet.com/Rise-of-the-patent-trolls/2010-1071_3-5892996.html.

On software patents, see, for example, JAMES BESSEN & MICHAEL MEURER, PATENT FAILURE: HOW

“bias.” There is no way to tell whether the rejection of particular NPE and software patents is right or wrong without having some Platonic insight into the “true” validity and scope of each patent. But even if we think that courts are wrong to invalidate these patents or limit their scope, they are nevertheless doing so. Society is spending a large chunk of its patent law resources dealing with what are—for whatever reason—the weakest cases. And patent plaintiffs are pursuing those cases despite the overwhelming odds against them.⁷⁹

What might this mean for patent reform? On the one hand, it should give substantial ammunition to those who argue against software patents and who want to restrain patent trolls. If software and NPE patents are overwhelmingly bad—either invalid or overclaimed—the social benefit of allowing them may well be outweighed by the harm they cause. At the same time, however, one could read this evidence as proof that the system is working—the bad patents are being weeded out of the system and are not stifling innovation.

The truth probably lies somewhere in between. The latter claim—that the widespread invalidation of software and NPE-owned patents shows that the system is working—seems altogether too facile. After all, roughly 90% of those cases settled without judgment. Although those settlements are confidential, we expect that the vast majority involved some sort of payment to the patent plaintiff—a payment that the outcomes data suggests might represent not the acquisition of real legal rights but a nuisance settlement over a likely invalid patent. At the same time, that these patents are so weak should—at least once exposed—limit the value of those settlements and quiet concerns that software or troll patentees will actually shut down many innovative products.

C. OUTLIERS: CRAZY LIKE A FOX?

Finally, it is worth noting that, by definition, our study is a study of outliers. Ronald Katz is obviously an outlier, having filed a sizeable percentage of the patent lawsuits in the last decade against hundreds of defendants.⁸⁰ But in some sense, anyone in the most-litigated patent set is an outlier simply by virtue of being willing to sue multiple times in different courts over the same patent. It may be that outliers are either irrational or simply have motivations that are not shared by the majority of patent owners. That is a possible explanation both for their unwillingness to settle and perhaps for their willingness to pursue losing

JUDGES, BUREAUCRATS, AND LAWYERS PUT INNOVATORS AT RISK 194–98 (2008), and Jacqueline Lipton, *IP's Problem Child: Shifting the Paradigms for Software Protection*, 58 HASTINGS L.J. 205, 208 (2006). *But cf.* Stuart J.H. Graham & David C. Mowery, *Software Patents: Good News or Bad News?*, in INTELLECTUAL PROPERTY RIGHTS IN FRONTIER INDUSTRIES: SOFTWARE AND BIOTECHNOLOGY 45, 73 (Robert W. Hahn ed., 2005) (evaluating claims that software patents are bad for innovation).

79. The same may be true of other types of lawsuits. Antitrust suits, for instance, continue to be brought despite their low win rate. A more general explanation for that phenomenon is outside the scope of this Article.

80. Although Katz has an outsized effect on the most-litigated patents, it is worth noting that even when we exclude Katz patents from the results, as we do with entity size, we still find dramatic differences between the most-litigated and the once-litigated patents. *See supra* note 41 and accompanying text.

cases to judgment. And it may be particularly true of the NPEs that dominate the most-litigated data set, because their motivations are likely to differ systematically from those of product-producing companies.

Alternatively, the patents (rather than the patentees) may be outliers in another sense—they might have so much potential value if enforceable that it is worth suing to enforce them even if the suit is unlikely to succeed. If the payoff for victory at the end of the day is \$1 billion, a 10% chance of success does not sound so bad. On this theory, the most-litigated patents are lottery tickets that rational plaintiffs use to take a chance at a big win. This theory is consistent with most-litigated patents also being litigated against far more defendants; if an entire industry potentially infringes a patent, there may simply be more money at stake. And higher stakes may also explain why the settlement finding is not as strong as we expected—the greater the stakes in a case, the less important attorney's fees are as an inducement to settle the case.⁸¹

We think there is something to both explanations, though, without data on how much was paid to settle these cases, we cannot test them.⁸² But even if the most-litigated patents are outliers, that does not mean that they are irrelevant. They represent a substantial percentage of patent litigation, and—precisely because they appear willing to take to trial weak cases others might settle—they may have an even larger influence on the law. Put another way, the Katzes of the world exist because we have structured our patent system in a way that permits them to exist. Katz may be an outlier, but he is not unique. Patent owners like this are a feature of the system. We should not ignore that reality in evaluating that system.⁸³

D. DO WE UNDERSTAND PATENT VALUE?

Whatever the explanation for the poor performance of the most-litigated patents, such poor performance calls into question the evidence economists have long relied upon to demonstrate patent value. The connection between patent claims, forward citations, backward citations, and application family size on the one hand, and the value of patents on the other, is well-established in the economic literature and forms the basis for a great deal of economic analysis not only of patents, but of innovation and growth more generally.⁸⁴ Allison and Mann have found software patents to be of above-average quality by traditional

81. Future research could investigate how the nature of the fee arrangement with the lawyer affects settlement behavior.

82. We hope in the future to test whether parties enforcing their own patents fared better or worse than parties enforcing patents they bought. The findings may shed some light on the irrational-behavior story.

83. We contemplated excluding Katz from our results. But we think doing so would paint an artificially rosy picture of the patent system.

84. See *supra* notes 15–20 (citing this literature). *But cf.* Bessen, *supra* note 15, at 932 (finding that patent citations are correlated with value, but explain little variance in value); Alfonso Gambardella et al., *The Value of European Patents*, 5 EUR. MGMT. REV. 69, 82 (2008) (finding a value linkage, but noting that the data is noisy).

economic measures.⁸⁵ ALMT found in 2004 a very strong connection between these measures of patent value and whether patents were litigated.⁸⁶ And in our previous work, we found that that connection was even stronger for the most-litigated patents.⁸⁷ Shrestha found that NPE patents had characteristics that indicated that they were more valuable than other patents.⁸⁸ And although we have always emphasized that it is private, not social, value that is measured by these statistics, we would have expected even that private value to be reflected at the end of the day in litigation outcomes. That it is not—that the very patents that by every economic measure are the most valuable ones turn out to be much weaker than other litigated patents—should give economists and other social scientists significant pause in using those measures of value for other purposes as well. Perhaps there is a definition of value that is independent of whether or not the patent is in fact valid and what it covers, but once we understand that the existing measures of value do not correspond to “good” patents, we suspect that definition will have to be very specialized. And in any event, the voluminous literature that measures the value of patents based on these characteristics, and the value of innovation based on the number of patents that have those characteristics, stands on shakier ground than previously thought.

E. DO WE UNDERSTAND LITIGATION BEHAVIOR?

Finally, our results are a bit of a puzzle for the most common law and economics models of litigation. The Priest–Klein model suggests that litigation is about divergent expectations between the parties, and in the absence of a systemic asymmetry, plaintiff win rates in the cases that are selected for trial should approach 50%.⁸⁹ An asymmetric stakes theory would suggest that plaintiffs have more to lose from going to trial than defendants, and so the plaintiff win rate should exceed 50% in the most-litigated cases.⁹⁰ Our data do not support either theory; perhaps they represent yet another nail in the Priest–Klein coffin.⁹¹ But they do beg the question of what is motivating the parties in these cases. Further work could investigate how the repeat-player status of both patentees and defendants affects settlements and outcomes.

85. Allison & Mann, *supra* note 76, at 333–34.

86. ALMT, *supra* note 15, at 437.

87. Allison, Lemley & Walker, *supra* note 23, at 3.

88. Shrestha, *supra* note 45, at 145–46.

89. Priest & Klein, *supra* note 47, at 17–24.

90. See, e.g., Robert Cooter, Stephen Marks & Robert Mnookin, *Bargaining in the Shadow of the Law: A Testable Model of Strategic Behavior*, 11 J. LEGAL STUD. 225 (1982) (discussing how asymmetric stakes might affect win rates). It is possible to read asymmetry the other way in certain respects. For example, before 2006 defendants faced an injunction if they lost, and the risk of that injunction might have been so catastrophic as to raise their stakes. But that should be equally true for the once- and most-litigated patents; any defendant cares only about the threat a patent poses to him, and he will only be sued once on any given patent.

91. For other criticism of Priest–Klein, see, for example, Steven Shavell, *Any Frequency of Plaintiff Victory at Trial Is Possible*, 25 J. LEGAL STUD. 493, 493–94 (1996).

CONCLUSION

We designed this study to explore the effects of repeat play on litigation behavior, contributing to a literature on the economics of civil procedure as well as the substance of patent law. But what we found was dramatic and unexpected: the patents and patentees that occupy the most time and attention in court and in public policy debates—the very patents that economists consider the most valuable—are astonishingly weak. Nonpracticing entities and software patentees almost never win their cases. That may be a good thing, if you believe that most software patents are bad or that NPEs are bad for society. But it certainly means that the patent system is wasting more of its time than expected dealing with weak patents. And it also suggests that both our measures of patent value and our theories of litigation behavior need some serious reconsideration.