

The Court of Appeals for the Federal Circuit's Impact on Patent Litigation

Matthew D. Henry and John L. Turner[†]

June, 2005

[†]Matthew D. Henry and John L. Turner are at the Department of Economics, Terry College of Business, University of Georgia. We thank Scott Stern, John Allison, Starling Hunter, David Mustard, Paul Heald, Alan Marco, James Bessen, an anonymous referee and seminar participants at the University of Georgia, the 2004 Southern Economic Association Annual Meetings and the 2005 International Industrial Organization Conference for helpful comments. Financial support from the University of Georgia Research Foundation is gratefully acknowledged.

Abstract

More than twenty years after the establishment of the Court of Appeals for the Federal Circuit (CAFC), research has yet to explain accurately the new court's impact on patent litigation, patenting and inventive activity. To address this shortcoming in the literature, we analyze a novel data set that permits us to separately consider the issues of validity and infringement in comparing the tendencies of the CAFC to those of its predecessor appeals courts. Our analysis of district and appellate decisions spanning 1953-2002 yields a recasting of the "pro-patent" nature of the CAFC: while it has been significantly more reluctant than its predecessors to affirm "invalid" decisions, it has not been more reluctant to affirm "not infringed" decisions. Because of the CAFC's tendencies, district courts have decided patents "invalid" significantly less often, patentees have appealed "invalid" decisions significantly more often and infringement has become the more frequently decisive inquiry.

1. INTRODUCTION

In the 1970's, the United States federal judiciary faced an enormous problem. The growth of the number of appeals had far outstripped the growth in the number of appellate judges. In response, the 92nd US Congress established the Commission on Revision of the Federal Court Appellate System, Structure and Internal Procedures, better known as the Hruska Commission, to make recommendations for reform (Public Law 489. 92nd Cong., 2d sess., 13 October 1972). In 1975, the Hruska Commission released its report (67 F.R.D. 195 [1975]). Among other things, the Commission identified the lack of uniformity in US patent law across the geographical circuits and the resulting rampant practice of forum shopping in patent cases.¹ The forum for a patent suit was strategically important, in great part, because the Circuit Courts of Appeal varied widely in the frequency with which they upheld the validity of patents.² To address these problems, the Hruska Commission recommended that a national appeals court be established to handle patent litigation (67 F.R.D. at 371). Congress responded to these recommendations with the Federal Courts Improvement Act of 1982 (Public Law 164. 97th Cong., 2d sess., 2 April 1982), which created the Court of Appeals for the Federal Circuit (CAFC) and established it as the sole US appeals court in patent cases.

At the beginning of its tenure, the CAFC effectively ended any influence of precedents from the geographical Circuit Courts of Appeal. In its very first appeal, *South Corporation et al. v. United States* (690 F.2d 1368, 215 U.S.P.Q. 657 [Fed. Cir. 1982]) the CAFC

¹Rochelle Cooper Dreyfuss (1989, p. 7), who discusses the Hruska Commission, notes “forum shopping was rampant, and . . . a request to transfer a patent infringement action from Texas, in the Fifth Circuit, to Illinois, in the Seventh Circuit, would be bitterly fought in both circuits and, ultimately, in the Supreme Court.”

²Robert Harmon (1992, p. 574), for example, states “When this author broke into the business, and for many years after, it was quite clear that there was no such thing as a valid patent in the Eighth Circuit, and the climate in the Ninth Circuit was not much more hospitable. In the Seventh Circuit, on the other hand, patent infringement could get a client into big trouble. Each of the other circuits occupied its own band in the enforcement spectrum,…”

adopted the holdings of its “parents,” the US Court of Claims and the US Court of Customs and Patent Appeals, as binding precedent.³ Thus, it immediately unified the patent law and eliminated appellate court forum shopping.

The unification of the patent law brought additional consequences, because early CAFC precedents strengthened the statutory presumption of patent validity (35 U.S.C. 282), making an invalidity defense less viable.⁴ For example, the CAFC has required the alleged infringer or patent challenger to produce “clear and convincing evidence” to prove invalidity.⁵ Prior to the creation of this court, the geographical Circuit Courts of Appeal had different standards of validity that were often clearly less favorable to the patentee. The Sixth Circuit endorsed the weaker “preponderance of the evidence” standard in *Dickstein v. Seventy Corporation et al.* (522 F.2d 1294, 187 U.S.P.Q. 138 [6th Cir. 1975]), noting in the opinion that the Second Circuit used this standard as well.⁶ Numerous patent attorneys and legal scholars (Lawrence Kastriner 1991; Harmon 1992; Quillen 1993; Paul Goldstein 1993) note the CAFC’s strengthening of the presumption,⁷ while Donald Dunner, Michael

³*South Corporation* was actually a trademark case, but the opinion included a general opening statement about court precedent. These ancestor courts had heard patent appeals from primarily the US Patent and Trademark Office’s Board of Patent Appeals and Interferences and the US International Trade Commission.

⁴According to Cecil Quillen (1993, pp. 192-95), the CAFC’s three most important changes were the elevation of nonstatutory factors in the determination of nonobviousness, the narrowing of the scope of prior art as it pertains to the obviousness issue, and the enforcement of the “clear and convincing evidence” standard for proving invalidity.

⁵An early example of the use of this standard can be found in *Connell v. Sears Roebuck* (722 F.2d 1542, 1549, 218 U.S.P.Q. 31, 38 [Fed. Cir. 1983]).

⁶“We are therefore of the opinion that in this case, and in the usual patent case in which validity is proved with similar evidence, a preponderance of the evidence is sufficient to establish invalidity. In this regard we are in agreement with the Second Circuit” (187 U.S.P.Q. at 140). As reported in a survey (*United States v. Fatico* (458 F.Supp. 388, 410 [E.D.N.Y. 1978])), judges from the Eastern District of New York, though not unanimous, generally viewed the “preponderance of the evidence” standard as representing a likelihood of 50+%, the “clear and convincing evidence” standard as representing a likelihood of 60-70% and the “beyond a reasonable doubt” standard as representing a likelihood of 80-90%. Christine Jolls, Cass Sunstein and Richard Thaler (1998, pp. 1530-31) offer a broad discussion of the different evidentiary standards.

⁷Kastriner (1991, p. 10) refers to the enforcement of the presumption of validity as “the first step taken by the CAFC which materially strengthened patents.” Harmon (1992, p. 575) writes that “the Federal Circuit’s rigorous observation of the presumption of validity” has made obviousness a more difficult defense in patent litigation. Goldstein (1993, p. 365) states, “The CAFC has not only eliminated intramural conflict and forum shopping. The court has also buttressed the patent grant itself, giving new force to the statutory presumption of validity.” Robert Merges (1997) also discusses the “pro-patent” reputation of the CAFC.

Jakes and Jeffrey Karceski (1995) measure an empirical rate of patent validity (under the CAFC) substantially higher than Gloria Koenig (1980) finds for the geographical Circuit Courts of Appeal prior to 1982.⁸ Not surprisingly, the CAFC has earned a reputation as a “pro-patent” court.

However, the magnitude and scope of the CAFC’s impact, including the impact of its strengthening of the presumption of validity, remain unclear. As we discuss in section 2, research has yet to measure accurately how pro-patent it has been (relative to appeals courts before 1982) and has yet to test, statistically, whether it has been pro-patent with respect to patent validity, patent scope or both. Given that there is some confusion within the economics literature about the way in which the CAFC has been pro-patent,⁹ the latter inquiry takes on greater importance. In addition, research has yet to identify how the CAFC’s changes to patent law have affected the behavior of lower courts and litigants.

As a result, it remains difficult to assess its impact on the incentives to invent, write and litigate patents, and on economic welfare generally. In a related literature, for example, several papers (Kortum and Lerner 1998; Bronwyn Hall and Rosemarie Ham Ziedonis 2001; Hall forthcoming) study the “Friendly Court Hypothesis,” which contends that the establishment of the CAFC is responsible for the simultaneous surge in patenting in the early 1980s. This research has yet to settle on whether this hypothesis is correct.¹⁰ In our view,

⁸Koenig (1980) reports that during 1968-78, Circuit Courts of Appeal found 31% of adjudicated patents valid. Dunner et al. (1995) report that during 1982-94, the CAFC found patents valid in 67% of adjudications. Koenig (1980) gathers data from the United States Patents Quarterly (USPQ). She excludes cases where the patent was found not infringed and validity was not an issue. Dunner et al. (1995) provide more specific data. They report the disposition of the court on patent validity under sections 102, 103 and 112 of Section 35 of the US Code. Disregarding instances where the CAFC vacated a lower court ruling (which occurred in, respectively, 10%, 11% and 17% of cases), it held patents valid under these sections, respectively, 64%, 66% and 77% of the time. The 67% figure combines the data from all three sources of patent validity adjudications. Note that a patent may be adjudicated under more than one section.

⁹Samuel Kortum and Josh Lerner (1998, p. 255), for instance, attribute it to a wider application of the “doctrine of equivalents,” which pertains to the infringement inquiry only.

¹⁰Kortum and Lerner (1998) argue that a simultaneous surge in the productivity of R&D explains the surge in patenting, while Hall and Ziedonis (2001) and Hall (forthcoming) conclude that the CAFC is responsible for some of the surge.

this failure stems, in great part, from the fact that no paper in this literature gathers data on patent litigation to study the CAFC directly.

In this paper, we make three key contributions. First, we identify and estimate a “transition matrix” of conditional probabilities of patent litigation outcomes in the United States for 1953-2002, including (1) probabilities of the three most common types of district court decisions (“invalid,” “not infringed” and “valid and infringed”), conditional on one of the three decisions being issued; (2) probabilities of appeal, conditional on the district court decision; and (3) probabilities of appeals court affirmation, also conditional on the district court decision. By conditioning on the district court decision, we can separately analyze courts’ tendencies with respect to the validity and infringement inquiries. In section 3, we report estimates of these transition matrices for the First through the Tenth Circuits and for the CAFC and compare the tendencies across the different appeals court regimes. To our knowledge, our data form the first comprehensive set gathered from a single source that spans significant portions of both the pre-CAFC and CAFC eras.

Second, we analyze annual time series of estimates of conditional probabilities, aggregated across all courts and litigants (section 4). We test for persistence and structural stability in these series. This analysis yields considerable insight into how the impact of the CAFC on patent litigation has differed with respect to the validity and infringement inquiries. In particular, we can explain the impact of the CAFC on the validity inquiry with a simple structural break model, but the same model fails to accurately describe its impact on the infringement inquiry. With respect to infringement, our analysis identifies several new questions, clearly justifying our approach of separately analyzing these inquiries and motivating further research.

Third, by showing how the CAFC has affected litigation outcomes at the district level and probabilities of appeal of those outcomes, our analysis provides a rare glimpse at how

an entire legal system adjusts to a major policy change. The contrasts between the validity inquiry, where the most important CAFC changes were immediate and well understood, and the infringement inquiry, where some changes introduced after 1982 were not well understood, are particularly telling. For instance, it is intuitive that courts and litigants would respond to the CAFC's strong stance on validity, but very surprising that their responses have had significant *external* consequences for the law on infringement. In one (perhaps unintended) consequence, patent cases have come to hinge on the infringement inquiry with increased frequency *because* invalidity became difficult to prove. As a result, the law on infringement has come under substantial stress and has seen several significant modifications.

To our knowledge, this is the first time series analysis of appeals court tendencies and their consequences in patent litigation. Our primary model, which tests for a structural break in the mean of the series, was developed by Donald Andrews (1993) and refined by Timothy Vogelsang (1997). It identifies the most likely break date and tests for the significance of the mean break in that year. The presence of a structural break indicates that an exogenous change to the mean of the series occurs around the time of the break year. In our case, we wish to know if there are such changes to the conditional population probabilities in the aggregate transition matrix that can be linked to the establishment of the CAFC and/or its stronger presumption of patent validity in 1982.¹¹

In our main set of results, we identify significant structural breaks in the part of the aggregate transition matrix pertaining to “invalid” patent decisions and argue that they are linked directly to the establishment of the CAFC. Specifically, we find breaks in time series of each of the three conditional probabilities during 1982-83 and show that each break is

¹¹Structural break models have been used extensively to test for stability in macroeconomic time series, both as specification tests and as a way of identifying the timing of breaks. In the literature on patents, Jon Merz and Nicholas Pace (1994) identify an increase in the mean and trend of new litigation filings in 1982, while Hall (forthcoming) and John Turner (2005) identify a significant increase in the growth rate of patent applications in the early 1980s as well.

consistent with a stronger patent grant under the CAFC. We estimate that, since the breaks occurred, district courts have been roughly *half as likely* to issue an “invalid” decision, patentees have been about 25% more likely to appeal these decisions, and the appeals court has been nearly *three times more likely* to not affirm an “invalid” decision. Together, the timing, synchronicity and intuitive consistency of these breaks are strong evidence that the CAFC’s stronger presumption of validity has had a significant impact on these probabilities. Also strengthening the results, the dynamics of these series are simple and fit the structural break model quite well.

The dynamics of the transition probabilities pertaining to “not infringed” and “valid and infringed” decisions, however, are each quite distinct from those for “invalid” decisions and are far more complicated. Notably, the CAFC’s probability of affirmation of “not infringed” patents is not significantly different from the aggregate pre-CAFC probability, but the conditional probability of “not infringed” decisions (at the district court level) surges after 1982. This indicates that the infringement inquiry has been crucial to the outcome of patent cases more frequently. Bolstering this, we also find that “valid and infringed” decisions have hinged more frequently on the infringement inquiry under the CAFC (section 4.2). Both the conditional probability of “valid and infringed” decisions (at the district court level) and the probability of appeals court affirmation of such decisions surge initially after 1982, but decline sharply in the early 1990s. Given that the CAFC’s major innovations to the law surrounding the infringement inquiry came well after 1982, it is quite likely that such innovations contributed to these later changes in conditional probabilities of “valid and infringed” decisions and probabilities of affirmation.

It is important to note that our estimates of conditional probabilities are not necessarily good estimates of analogous *unconditional* probabilities and are subject to several selection biases. These have been noted by numerous past researchers and clearly affect the

interpretation of the transition matrices for the different appeals court regimes.¹² However, selection effects that are difficult to identify among the individual appeals court regimes may be more easily seen in a time series analysis because they tend to emerge dynamically. In our case, the analysis of structural breaks and persistence actually helps to identify the presence or absence of some selection effects, because we can identify both regime changes (structural breaks) and statistical properties of the transition between regimes (persistence). When subject to selection effects, a series will not typically adjust to a new regime instantaneously; rather, it tends to persist away from its mean. By showing that time series for “invalid” decisions have little persistence, while some of the time series for “not infringed” and “valid and infringed” decisions have significant persistence, we show that the latter two categories are subject to far more selection effects. We explore likely sources of such selection in the discussion in Section 4.2. We conclude with a comprehensive explanation of the CAFC’s impact that incorporates our statistical results as well as some speculation about likely selection effects.

2. PATENT LITIGATION

2.1. The Course of Litigation

Either the patentee or alleged infringer may initiate patent litigation.¹³ To win, the patentee must show that his patents are both valid and infringed. The alleged infringer need

¹²For example, George Priest and Benjamin Klein (1984) argue that selection effects will tend to keep observed litigation success rates near 50%. However, numerous studies (including Kevin Clermont and Theodore Eisenberg [1992] and Daniel Kessler, Thomas Meites and Geoffrey Miller [1996]) find litigation win rates that clearly diverge from 50%.

¹³Frequently the patentee initiates the litigation. If, however, the alleged infringer can demonstrate that he is likely to be sued for infringement in the future, he can file for a declaratory judgment for a finding of non-infringement or a finding of invalid, under 28 U.S.C. 2201, 2202. If the patentee counterclaims for infringement, the resulting suit is similar in nature to a suit initiated by the patentee. Because the patentee and the alleged infringer will typically prefer different lower courts for the trial, he who files first may gain a tactical advantage by naming the venue.

show only invalidity or non-infringement to win. The three most common ultimate outcomes for a patent involved in an infringement suit are: (1) the patent is found “invalid,” (2) the patent is found “not infringed,” or (3) the patent is found “valid and infringed.”

The owners of “valid and infringed” patents are entitled to compensatory damages from the infringers.¹⁴ The owners of “invalid” and “not infringed” patents are not entitled to anything from the alleged infringer. When a court rules a patent “invalid,” it essentially states that at the time the patent was issued, the patented invention did not satisfy the legal requirements of patentability when the inventor applied for the patent and therefore the USPTO was in error in granting the patent.¹⁵ When a patent is decided “invalid,” the question of infringement is mooted because, in the eyes of the court, the inventor’s exclusivity over the exploitation of his invention no longer exists.¹⁶

When a lower court decides a patent “not infringed,” it essentially states that the patent grant does not cover the alleged infringer’s exploitation of the products and/or processes that are under consideration by the court. Here, the patent’s validity remains intact, but the alleged infringer is free to continue using his (non-infringing) technology. The scope of a patent’s protection is within the discretion of the court, and this determination is often crucial in determining whether infringement has occurred. Thus, a finding of “not infringed,” though not invalidating the patent, may leave it with such a narrow scope as to be equally worthless to an inventor.

Each year, numerous “invalid,” “not infringed” and “valid and infringed” decisions are appealed. In each category of appeal, there are several possible outcomes. The most common

¹⁴These damages often take the form of a royalty paid for units of the infringing product sold, but are supposed to reflect “lost profits” stemming from the illegal competition introduced by the presence of the infringing product. However, the court is required only to reimburse the patentee for damage caused by the infringement, and the degree of damage is within the discretion of the court. Damage awards for “valid and infringed” patents vary considerably.

¹⁵For instance, the court may determine that at the time the patent application was filed, either the invention was either not novel or was obvious.

¹⁶In anticipation of potential appeals, courts do sometimes rule on infringement, however.

outcome is that the lower court’s decision is affirmed, while the second is that the appeals court reverses the lower court. Appeals courts vacate some decisions as well. When an appeals court reverses the lower court, it has two primary options. It can replace the lower court’s decision with a new one (e.g. reverse “valid and infringed” to “invalid”) or remand the case back to the lower court for further findings. If reversed or vacated, the lower court is not necessarily precluded from reaching the same ultimate conclusion as earlier, but must use different reasoning.¹⁷ A “reverse and remand” decision may also restrict the court’s attention to validity only or infringement only (this happens most often in cases where the lower court’s original decision was “valid and infringed”).

2.2. Empirical Work

In the past fifty years, several researchers have performed empirical studies of patent litigation. Two of these studies are primarily a written statistical record; they include little rigorous analysis and few substantive conclusions. P.J. Federico (1956), at the request of a US Congressional subcommittee, reports data on patent decisions published in the *United States Patents Quarterly* (USPQ) during 1925-54.¹⁸ Koenig (1980), using the same source, records slightly different statistics on patent decisions from 1953-78.¹⁹

¹⁷For instance, if an appeals court reverses the lower court’s finding of invalidity because of improper specification, the lower court may still ultimately decide the patent invalid for obviousness.

¹⁸At the request of the Subcommittee on Patents, Trademarks, and Copyrights of the Committee on the Judiciary, United States Senate, 84th Congress, 1955. Federico reports district and appeals court decisions from the *United States Patent Quarterly* (USPQ) from January 1948 to September 1955 (volumes 76-106). He mainly considers final adjudications, and separates adjudicated patents into the categories “valid and infringed,” “invalid” and “not infringed.” He includes utility patents, design patents, and plant patents, but excludes reissue patents. He also presents yearly numbers of patent adjudications from 1925-47, which were apparently gathered separately.

¹⁹Koenig’s statistics exclude plant and design patents, as well as cases from the Court of Claims and cases decided on the basis of issues such as estoppel and misuse. She separates the decisions on individual patents into the following categories: “valid and infringed,” “valid and not infringed,” “not infringed,” “validity not decided” and “invalid.” She presents data on both the number of cases in each category and the corresponding affirmation rates. She also presents data on the number of applications and approval rates of patents. Based on our research, it appears that her yearly figures reflect the publication dates of patent decisions rather than the dates when the decisions were rendered.

Other studies draw more conclusions from their data. Lawrence Baum (1974), who studies the validity of patents involved in litigation from 1921-73, identifies a drop in patent validity rates between the early 1930's and 40's, but does not estimate the year or significance of this drop.²⁰ Dunner (1985) and Dunner et al. (1995), who study decisions by the CAFC during 1982-94, treat these data as a population. Their analysis shows that the Federal Circuit was more likely to affirm judgments in favor of patent owners than accused infringers (Dunner 1985, p. 154). They also find that the court was more likely to uphold findings of "not infringed" than it was to reverse them.²¹ John Allison and Mark Lemley (1998) examine validity rulings during 1989-96 and perform a battery of hypothesis tests on the data as one group.²² Jean Lanjouw and Mark Schankerman (2001), who study patents litigated between 1979-95, focus on differences in characteristics of litigated and non-litigated patents and do not assess the impact of the CAFC. Kimberly Moore, using data gathered from the Administrative Offices of the US Courts on patent litigation since 1983, measures differences in outcomes of judge vs. jury trials (2000), the extent of forum shopping (2001) and differences in outcomes for domestic vs. foreign patentees (2003). Her data do not include information about validity or infringement. Finally, Marco (2005) uses USPQ data

²⁰Baum's data come from the *Official Gazette of the Patent Office* (1921-1960) and the *United States Patent Quarterly* (1961-73). His paper examines only the validity of the patents involved in the suits. He asserts a variety of causes for the "break," including the courts' concern with the anti-competitive uses of patents (and the negative attitude toward monopolies), and the leadership of the Supreme Court. He discounts other explanations, such as that weaker patents were being adjudicated because of the increase in the number of declaratory judgment actions (p. 762, n. 11).

²¹The latter study was an expansion of Dunner's earlier work. This work captures all the published and unpublished opinions from the Federal Circuit in cases appealed from the lower courts, the Court of Federal Claims, the International Trade Commission, and the Patent and Trademark Office. They focus on each patent adjudicated, and separate the cases into the statute section under which they were decided.

²²Allison and Lemley exclude cases on enforceability "due to misuse, inequitable conduct, estoppel, or the like" and infringement decisions. They also exclude cases from the International Trade Commission. They report only the last final decision in each case. The study includes all published and unpublished decisions. They separate the patents on the grounds for the validity decision and by the type of the patent involved (more specifically than simply utility, plant, or design). They find that 54% of patents adjudicated were decided "valid." They also conduct a variety of hypothesis tests concerning the effect of juries on validity and the various reasons for invalidating patents on the different types of patents. They also present trial court affirmation or reversal rates (although it appears that they follow cases that were remanded to the district court and present only final decisions by district courts).

from 1977-97 to estimate probabilities of erroneous findings of validity and invalidity by district and appeals courts.

Unfortunately, none of the data sets in these previous studies can be used to compare courts' treatment of patents across the pre-CAFC and CAFC eras, or across time generally. First, none include large amounts of data from both before and after 1982, and their sets cannot be combined because the authors gather their data from different sources and use different criteria. Second, many of these data sets were not constructed for the purposes of statistical analysis, so tests using the reported statistics may lead to biased conclusions anyway. Since we wish to study time series changes of the treatment of patents by the courts, particularly the CAFC, we must build a new data set.

3. THE DATA

3.1. Criteria

Our data set of patent litigation, which spans 1953-2002, includes virtually all patents found “invalid,” “not infringed” or “valid and infringed” by a lower court in a case where at least one decision (including appellate decisions) is recorded in the USPQ.²³ We include all types of patents (design, plant, reissue and utility).²⁴ We construct it from two separately-gathered data sets, one for district court decisions recorded in the USPQ and another for

²³The vast majority of these cases are decided in district courts, but there are a handful from the US Court of Claims and International Trade Commission. The USPQ contains all written opinions from cases involving patents, copyrights, and trademarks. Because the large majority of appeals court cases generate such opinions, one may be relatively certain that most appellate-level decisions have been captured. Federico (1956) estimated that the USPQ contained about half of the district court decisions. Additionally, it should be noted, however, that at various times, the USPQ has published decisions without certification of the deciding court. These decisions may not be subsequently cited and have no precedential value. We have included such cases in our main data set, but duplicate our analysis on a subset of the data that excludes these decisions to check for robustness. In doing so, we find no substantial differences.

²⁴Utility patents comprise about 90% of the data. Our results are qualitatively the same if only these patents are considered.

appeals court decisions recorded in the USPQ. We combine the two data sets, matching district court decisions with appeals court decisions that are part of the same case and filling in as much missing information as possible. That is, for each decision, we search the USPQ and *Westlaw* to identify subsequent decisions in the same case.

The basic unit in the data set is a particular patent in a particular case (a “patent case”).²⁵ For each patent case, we record the following variables: (1) date of district court decision, (2) USPQ Citation(s), (3) district court, (4) decision, (5) date of appeals court decision (if there is one), (6) appeals court, (7) appeals decision, (8) subsequent action, (9) patent number, (10) litigants, (11) first patent inventor, (12) first inventor’s home country, (13) patent assignee, if different from inventor, (14) assignee’s home country, (15) number of claims and (16) product code. All information reflects primary sources—the information for criteria (1)–(10) is from the published opinions on the cases,²⁶ while the information for criteria (11)–(16) comes from the patent documents themselves.²⁷ In cases where there is no appeal of the first district court ruling, we record the appeals court that would have handled the appeal. Prior to discarding patent cases for reasons to be discussed below, our set includes 5,178 patent cases.

We construct this data set to study how courts have handled the issues of patent validity and infringement. As such, we include only cases that discuss one or both issues. Thus, all “housekeeping” decisions are excluded.²⁸ Furthermore, decisions about preliminary injunctions (which are granted to prevent further damages during the course of the trial) are also excluded, because the plaintiff carries a higher burden of proof in such a decision.²⁹ In

²⁵Thus, for a case with, say, four patents at issue, there are four patent cases. Note that in the analysis presented here, the basic unit is a decision in a patent case.

²⁶We did not rely on the “particular patents” section in the USPQ annotation at the beginning of each opinion. We found numerous inconsistencies between this and the body of the opinion, so we relied on the latter.

²⁷Full text can be viewed at the USPTO’s website, www.uspto.gov.

²⁸Examples are improper venue challenges and lack of standing decisions.

²⁹The statute 35 U.S.C.A. 283 gives discretion to the court to grant preliminary injunctions when it

addition, cases where the patentee’s conduct is the primary issue are also excluded.³⁰

Cases involving licenses, either formal or implied, are somewhat more difficult to handle. Most fall into one of two categories: (1) the parties are operating under an agreement when the licensee decides it can do better in the courts than at the bargaining table; or (2) the court determines that the patent is not infringed because of the way a license, explicit or implicit, is construed. Decisions in the former category are typically based on the validity of the patent or infringement by the actions of the licensee (for instance, by continuing production without the necessary license), while decisions in the latter class are not. Thus, decisions in the latter category are excluded from our data.

Many patents form the basis of multiple patent cases in this data set. When two or more cases are substantially similar and are decided in the same circuit, we discard the later case(s). A relatively common example of this occurs when a court finds a patent “valid and infringed” and, subsequently, the infringing party attempts to invent around the patent, prompting the patentee to sue for infringement again. We do not exclude patents adjudicated in different circuits either simultaneously or sequentially, however, because we wish to measure differences in courts’ tendencies across circuits. We exclude cases decided on the basis of collateral estoppel or *res judicata*,³¹ and exclude all decisions where the patentee sues for contempt by the alleged infringer of a previously issued injunction.³² In total, 386 patent cases are excluded, leaving a total of 4,792.

deems it reasonable. The CAFC advocates a four-factor test to determine when granting an injunction is reasonable: 1) a reasonable likelihood of success on the merits; 2) irreparable harm if the injunction were not granted; 3) the balance of the hardships; and 4) the impact of the injunction on the public interest. See, e.g., *Genentech, Inc. v. Novo Nordisk A/S* (108 F.3d 1361, 42 U.S.P.Q 1001 [Fed. Cir. 1997]).

³⁰Examples of cases in this category are cases decided based on inequitable conduct, patent misuse or laches. We extend this principle to exclude cases where the patent was adjudged “invalid” due to fraud on the part of the patent holder (though most courts would say that the patent is “unenforceable” rather than “invalid”).

³¹The plaintiff is estopped from filing suit if the matter was previously adjudicated and the parties in the suit enjoy mutuality.

³²In such a case, the infringement inquiry is virtually identical to the inquiry in a normal patent case, but the validity of the patent is not open to discussion.

Courts occasionally judge some of a patent’s claims differently than they do others. Following Federico (1956), we distinguish such patents as follows. If any claim is held “valid and infringed,” we record the patent as “valid and infringed,” no matter how many other claims are found to be “invalid” or “not infringed.” The reason for this determination is that the patentee is winning something from his patent, and damages are seldom increased or decreased based on the number of infringed claims.³³ If there are no “valid and infringed” claims, then if any claim is held “invalid” while others are held “not infringed,” we record the patent as “invalid.” Finally and in similar fashion, “not infringed” claims trump “not adjudicated” claims (i.e., undecided or remanded for other issues).³⁴

3.2. Transition Probabilities

We identify and estimate a “transition matrix” of conditional probabilities of patent litigation outcomes. This matrix includes (1) probabilities of the three most common types of district court decisions (“invalid,” “not infringed” and “valid and infringed”), conditional on one of the three decisions being issued; (2) probabilities of appeal, conditional on the district court decision; and (3) probabilities of appeals court affirmation, also conditional on the district court decision. These three measures, which give transition probabilities in and out of various states within a patent litigation, yield information about the tendencies of litigants, district courts and appeals courts. Our point estimates of each of these transition probabilities are straightforward.

We estimate the first set of probabilities, which we refer to as “shares” of decisions for convenience, with the fractions of each type of decision of the total number of (“invalid,” “not infringed” or “valid and infringed”) district court decisions. When considered alongside

³³Damages are usually calculated by assessing royalties for the number of infringing devices sold, not based on “how much” they infringed the patent.

³⁴We record patents found “valid but not infringed” as “not infringed.”

information about appeals court tendencies, these estimates help explain the decisions of district court agents, such as judges and juries. For a judge, the tendencies of the appeals court relevant to him are important because he will generally prefer to have his decisions affirmed. A reversal suggests that the judge handled the case improperly. When a judge repeatedly has his decisions reversed, he may develop a poor reputation as a jurist, which may jeopardize his chances for career advancement. In addition, a reversal creates additional work if the case is remanded, thereby crowding a court's docket. Thus, we expect that appeals court tendencies should affect the decision-making of judges and, hence, that of lower courts as a whole.³⁵

We estimate probabilities of appeal, which we refer to as “rates” of appeal for convenience, with the fraction of district court decisions for which we record subsequent action in the same case.³⁶ We construct estimates for each type of district court decision. These estimates convey information on the behavior of litigants, who choose whether to file suits, when to settle them (if ever) and whether to appeal decisions. If an appeals court is more (less) likely to overturn a particular decision, it is intuitive that the party losing in the lower court decision is more (less) inclined to appeal. Thus, rates of appeal provide further insight into the importance of appeals court tendencies.

We estimate probabilities of affirmation, which we also refer to as “rates,” with the fraction of district court decisions that are affirmed by the appeals court. These estimates give the cleanest available measures of appeals courts' tendencies. An affirmation represents an endorsement of the district court decision and a clear defeat for the appellant, while any other decision represents some unwillingness of the circuit court to endorse the district court

³⁵The impact on juries is less clear, but appeals court precedent should still affect their decision-making. Juries comprise a very small percentage of patent cases in the pre-CAFC era, but have become more common under the CAFC. Even in a jury case, the judge may still have a great deal of control over the outcome, through tools such as the *judgment non obstante veredicto* (*JNOV*), jury instructions, etc. For further discussion, see Moore (2000).

³⁶We include all unpublished appellate decisions in the calculation of these estimates.

decision.³⁷ In considering appeals of “invalid” decisions, for instance, we can focus on the appeals courts’ tendencies in handling the validity inquiry, because the appeals court nearly always addresses it directly.³⁸ Infringement is more complicated because it is not frequently considered when invalidity is established; for studying it, we consider both “not infringed” and “valid and infringed” decisions.

3.3. Descriptive Statistics

First, consider district court decisions. Our data include many observations where the appeals court decision is recorded in the USPQ but the district court decision is not. Because including these observations in estimating probabilities of appeal would bias these measures upward,³⁹ we omit them here. We also omit “second” decisions in a case—i.e. those coming after at least one appeals court decision.

After these adjustments, we have 3,315 “first” district court decisions. For 2,327 of these, the relevant appeals court is one of the geographical circuits.⁴⁰ We refer to the set of these situations, for convenience, as the pre-CAFC era. For the remaining 988, the CAFC is the relevant appeals court.

Table 1 presents total numbers of decisions broken down by geographical circuit (for the pre-CAFC decisions), estimates of the conditional probabilities of the decisions (henceforth

³⁷Nearly all non-affirmations represent a net victory for the appellant. There are some exceptions—for instance, a reversal of “not infringed” to “invalid” is, ultimately, a net defeat for the patentee. Decisions in this latter category represent a negligible percentage of cases, however, and our results do not change if they are counted differently.

³⁸One important exception to this occurred during 1987-1992, when the CAFC had a practice of vacating district court “invalid” decisions when non-infringement was established. It began this practice with its decisions in *Vieau v. Japax, Inc.* (823 F.2d 1510, 3 U.S.P.Q.2d 1094 [Fed. Cir. 1987]) and *Fonar Corp. v. Johnson & Johnson* (821 F.2d 627, 3 U.S.P.Q.2d 1109 [Fed. Cir. 1987]). The Supreme Court ended this practice with its decision in *Cardinal Chemical v. Morton International Inc.* (508 U.S. 83, 26 U.S.P.Q.2d 1721 [1993]).

³⁹We also expect that including these would bias estimates of the share of decisions, but the direction of the bias is less obvious.

⁴⁰Twenty-two were decisions in lower courts where the relevant appeals court was either the DC Circuit, the 11th Circuit or the CCPA. Because there are so few decisions, we do not report estimates of conditional probabilities for these circuits individually.

“shares” of decisions) and estimates of the conditional probabilities of appeal (henceforth “rates” of appeal). Based on these estimates, it is immediately clear that the CAFC regime is unlike any other. The CAFC-era share of “invalid” decisions, .279, is about half the aggregate pre-CAFC share, .557, and is smaller than the shares for *all* of the geographical circuits. The share of “not infringed” decisions, .352, is nearly three times as large as the aggregate pre-CAFC share, .13, and is larger than the shares for all of the geographical circuits. The shares of “valid and infringed” decisions are similar across eras.

Insert Table 1 here

The rate of appeal has been larger in the CAFC era for all types of district court decisions. The rate for “invalid” decisions, slightly more than 2 in 3, is 14.1 percentage points (slightly more than 25%) larger than the aggregate pre-CAFC rate, and is larger than in all of the geographical circuits except the Fourth and Tenth Circuits. The rate for “not infringed” decisions, slightly less than 3 in 5, is larger than the rate in all circuits except for the Fifth.

Next, consider appeals court decisions. A small percentage of these are not recorded in the USPQ,⁴¹ but we include them in the analysis to come because there are no reasons to expect them to cause bias. We do omit “second” appeals decisions in a case.

We have a total of 3,268 “first” appeals decisions. Of these, 1,927 were issued by the geographical Circuit Courts of Appeal, while 1,341 were issued by the CAFC. Table 2 presents raw numbers of decisions broken down by circuit and by prior district court decision, as well as estimates of affirmation rates.

Insert Table 2 here

⁴¹Among appeals in our data set, most not recorded in the USPQ are recorded in the *Federal Reporter*. A small number are recorded only in *Westlaw*.

The most striking feature of Table 2 is the CAFC's .573 rate of affirmation of "invalid" decisions, which is nearly 28 percentage points less than the .85 aggregate rate in the pre-CAFC era. This implies that the CAFC has been nearly three times more likely to not affirm an "invalid" decision. Its affirmation rate is substantially smaller than the rates for all of the Circuit Courts of Appeal. The CAFC has also affirmed "valid and infringed" patents at a larger rate than all but the Tenth Circuit, and its .723 rate is more than 12 percentage points larger than the aggregate pre-CAFC rate. The CAFC's .685 rate of affirmation of "not infringed" patents is essentially identical to the aggregate pre-CAFC rate.

The aggregate statistics in Tables 1 and 2 identify the most important of the CAFC's tendencies and some of their consequences. As we argue in the rest of the paper, the CAFC's smaller rate of affirmation of "invalid" patents is consistent with the smaller share of district court decisions of "invalid" during the CAFC era and the larger rate of appeal of these decisions. The intuition is that district courts (judges, in particular) have a strong incentive not to have their decisions overturned, so they respond to a lower affirmation rate by issuing fewer "invalid" decisions. Patentees, knowing that the appeals court is less likely to affirm an "invalid" decision, are more inclined to appeal.

4. ANALYSIS

Each of the estimated conditional probabilities is influenced by selection effects. Unobserved factors, such as the types of patent cases that reach trial, the rate of settlement and other sources of heterogeneity, complicate their interpretation. For these reasons, the descriptive statistics presented in the previous subsection are insufficient for analyzing the impact of the CAFC. A time series analysis is warranted.

4.1. Statistical Tests

We now estimate the magnitude and significance of the CAFC’s impact on patent litigation, by comparing its tendencies and those of lower courts to the tendencies of courts in the pre-CAFC era. Since we are interested in learning both the magnitude and timing of changes in the features of patent litigation, we analyze the stationarity of fifty-year annual series of shares of district court decisions, rates of appeal and affirmation rates. This consists of testing for persistence (i.e. serial correlation) and for broken trends. We focus on the magnitude and timing of structural breaks, a particular type of broken trend associated with permanent regime change.

Unit root tests offer evidence of both persistence and broken trends, while structural break tests help to identify or rule out a particular type of broken trend. In testing for unit roots, we use simple augmented Dickey-Fuller (ADF) tests. In testing for structural breaks, we use sup-Wald tests.⁴² The ADF test uses an ordinary least squares (OLS) regression to estimate the following model:

$$\Delta y_t = \alpha + \rho y_{t-1} + \sum_{j=1}^{jmax} \beta_j \Delta y_{t-j} + \epsilon_t, \quad (1)$$

where Δ denotes the first difference of the series. The null hypothesis is

$$H_0 : \rho = 0$$

and the alternative is that $\rho < 0$. The ADF test statistic is a simple t-statistic for ρ .⁴³

The sum of the autoregressive coefficients in the corresponding $AR(jmax + 1)$ model is commonly referred to as the *persistence* of the time series because the larger it is, the less is the tendency of the series to mean-revert. In our specification, the persistence is given by $\rho + 1$. Hence, when ρ is near zero, the persistence is near 1.⁴⁴

⁴²Andrews (1993) pioneered such tests and Vogelsang (1997) extended his work.

⁴³The critical values for this test are: 10%, -2.57; 5%, -2.86; 1%, -3.43. See James Hamilton (1994, p. 763).

⁴⁴The unit root hypothesis is, essentially, the hypothesis that the persistence of the time series is unity.

We use the lagged endogenous variables at the end of (1) to filter serial correlation out of the errors. We estimate $jmax$ using an iterated test for significance of the $jmax$ th lag. Specifically, we set $jmax = 3$ and run the regression. If the t-statistic on the $jmax$ th coefficient is significant at the 10 % level in a two-sided test, we increase $jmax$ to 6 and run the regression again. Otherwise, we reduce $jmax$ by 1 and run the regression again. The process is repeated until the final lag is significant or no lags are significant. We then use the results from that regression.

The structural break test uses a modified version of the model in (1):

$$\Delta y_t = \alpha + \alpha' 1(t > T_B) + \rho y_{t-1} + \sum_{j=1}^{jmax} \beta_j \Delta y_{t-j} + \epsilon_t \quad (2)$$

where $1(\cdot)$ is the indicator function. After the break date, T_B , the series has a possibly different mean. The null hypothesis, that the mean is the same before and after the break, holds if and only if $\alpha' = 0$.⁴⁵ Thus, we have

$$H_0 : \alpha' = 0.$$

T_B is a free parameter to estimate; we run this regression for each possible T_B , record the Wald statistic on α' , then use the supremum of this set of Wald statistics to identify the break year and test the null hypothesis. Statistical inference of the “sup-Wald” statistic itself depends upon whether the underlying time series includes a unit root, so there are two sets of critical values to consider.⁴⁶ The critical values are larger if the time series includes a unit root, because a highly persistent series is more likely to include a sequence of observations that suggest a structural break is present. The lag length $jmax$ is estimated in precisely the

⁴⁵Note that α is not the mean of the series. Rather, the mean before the break is $\frac{\alpha}{\rho}$, while the mean after is $\frac{\alpha + \alpha' t}{\rho}$. Clearly, these are identical if and only if $\alpha' = 0$. See footnote 52 for details.

⁴⁶We use the critical values published by Vogelsang (1997, pp. 824-25). If the underlying series is stationary, the critical values are: 10%, 9.24; 5%, 10.85; 1%, 14.99. If the underlying series contains a unit root, the critical values are: 10%, 16.14; 5%, 18.20; 1%, 22.64.

same way as in the ADF tests, except that here it must be estimated for each possible break date.

If a structural break is present in an otherwise stationary time series, the estimate of the persistence of the time series will be biased upward if it is estimated using equation (1). In this case, a consistent estimate of the persistence of the series can be obtained, conditional on the break, using equation (2). Moreover, as shown by Pierre Perron (1989), if a structural break is present, an ADF test is biased to not reject the unit root null hypothesis. Fortunately, it is possible, using a test pioneered by Eric Zivot and Donald Andrews (1994), to test for a unit root using equation (2) as well. This “inf-t” test builds the ADF t-statistic on the coefficient estimate of ρ for each possible break date, then chooses the smallest (infimum) of these statistics.⁴⁷

Each year’s statistics aggregate decisions from all courts.⁴⁸ There are nine series in total, built using yearly estimates of the same three conditional probabilities used in section 3: (1) probabilities (shares) of the three most common types of district court decisions, conditional on one of the three decisions being issued (SHARE); (2) rates of appeal, conditional on the district court decision (APPEAL); and (3) appeals court affirmation rates, also conditional on the district court decision (AFFIRM). We build a SHARE, APPEAL and AFFIRM series for each of the three most common types of district court decisions: “invalid,” “not infringed,” and “valid and infringed.” We first test for unit roots, then for structural breaks, so there are 18 tests in total.⁴⁹ The SHARE and AFFIRM run from 1953 to 2002. There are no “valid and infringed” decisions in 2002 in our data, so we cannot estimate a rate of appeal for “valid and infringed” for that year; we use data from 1953 to 2001 to test the

⁴⁷For statistical inference, one can use the critical values published by Zivot and Andrews (1994, p. 256). The critical values are: 10%, -4.58; 5%, -4.80; 1%, -5.34.

⁴⁸None of the appeals decisions through 1982 are due to the CAFC, while nearly all from 1983 on are from the CAFC. After 1982, the geographical circuits handled a small percentage of cases, in which they had built up knowledge of the facts in the case prior to the establishment of the CAFC.

⁴⁹We use MATLAB to estimate the model.

APPEAL series.⁵⁰ The results are presented in Tables 3 and 4.⁵¹

Insert Table 3 here

The most compelling results were found for “invalid” decisions. The unit root test statistics are presented in the bottom row of Table 3. We cannot reject the null hypothesis of a unit root for the AFFIRM and SHARE series, so they are clearly non-stationary. Based on the ADF t-statistic of -2.96, we narrowly reject the null for the APPEAL series at the 5 percent level (the critical value is -2.86).

We identify statistically significant structural breaks in all three series (see Table 4). The sup-Wald statistics for AFFIRM, APPEAL and SHARE each yield rejections of the null of no break at the 1 percent level.⁵² Break years for these series are, respectively, 1983, 1982 and 1982, so the breaks occurred at the onset of the CAFC era.

Based on the construction of the time series, the implied means of the series are:⁵³

$$\begin{aligned} \text{Mean (pre-break)} &= \frac{\alpha}{-\rho} \\ \text{Mean (post-break)} &= \frac{\alpha + \alpha'}{-\rho}. \end{aligned}$$

Estimates of these measures are included in the bottom part of Table 4. Because these time series are constructed using yearly aggregates, the estimated means need not be similar to the aggregate pre-CAFC and CAFC means from Tables 1 and 2. They are, in fact, nearly identical for all “invalid” series (*e.g.*, the estimated “pre-break” affirmation rate of .847 is

⁵⁰At the time these data were gathered, it is probable that many appeals of 2002 decisions were still pending, so we omitted the 2002 year from the “invalid” and “not infringed” series as well.

⁵¹Turner (2005) also reports the results for the AFFIRM series for invalid patents.

⁵²Since we reject a unit root for the APPEAL series, the sup-Wald statistic, 20.55, is compared to the 1% critical value of 14.99. Note that the Zivot-Andrews “inf-t” test for this series also picks 1982 (with *jmax* equal to zero), with a highly significant test statistic of -9.21.

⁵³An AR(*jmax*) centered around a non-zero mean is constructed as follows: $y_t = d_t + x_t$, where $d_t = \mu$ gives the mean and $x_t = \sum_{j=1}^{jmax} \gamma_j x_{t-j} + \epsilon_t$. The series y_t evolves as $\Delta y_t = \alpha + \rho y_{t-1} + \sum_{j=1}^{jmax} \beta_j \Delta y_{t-j} + \epsilon_t$, where $\alpha = -\rho\mu$ and $\rho = -\sum_{j=1}^{jmax} \gamma_j$.

virtually indistinguishable from the “pre-CAFC” affirmation rate of .85). Thus, these tests show that the differences between the courts in the pre-CAFC and CAFC eras, estimated in section 3, are both accurate and highly statistically significant. Since the breaks, district courts have been roughly half as likely to issue an “invalid” decision, patentees have been nearly 25% more likely to appeal these decisions, and the appeals court has been nearly three times more likely to not affirm an “invalid” decision.

Insert Table 4 here

The dynamics of these series are simple and easy to interpret. Consider Figure 1, which plots the AFFIRM and SHARE series along with the fitted lines from the structural break models. The affirmation rate and share of “invalid” decisions fell sharply at the onset of the CAFC era, strongly indicating that this regime had a significant impact on the standard of validity at the appeals level and (as consequence) the district court level. The APPEAL series, though not plotted in the graph, similarly indicates a sharp increase in the rate of appeal of “invalid” decisions at the onset of the CAFC era. This is also intuitively consistent with an increased appeals court standard of validity.

Insert Figure 1 here

Conditional on the structural break, the AFFIRM and SHARE series exhibit no persistence (the estimate of $\rho + 1$ is not significantly different from zero), while the APPEAL series exhibits a small level of negative persistence (-.269). These results indicate that the transition between regimes is, essentially, instantaneous. If it were gradual, by contrast, the time series would converge to its new mean more slowly, and thus *persist* away from its mean. Our estimation would then show significant *positive* persistence. Thus, given our

finding of no positive persistence, we conclude that any selection effects stemming from dynamic responses to the onset of the CAFC regime had no significant impact on time series for “invalid” decisions. Clearly, the structural break model describes the dynamics of these series well.

One additional interesting feature of this analysis is that a unit root is rejected for the APPEAL series, yet we identify a significant structural break. Given the findings of Perron (1989), this is somewhat surprising. The most likely explanation for this is that the structural break, though clear-cut in a statistical sense, is small in magnitude.

The dynamics for the “not infringed” series are dramatically different and more difficult to interpret. A unit root is strongly rejected both for AFFIRM and APPEAL, and neither series has a structural break that is significant at the 10 percent level. In both cases, the series exhibit little if any persistence. Thus, the CAFC’s affirmation rate of “not infringed” decisions has been consistently similar to that of its predecessors, and the rate of appeal of district court decisions of “not infringed” has remained steady as well.

The SHARE series, plotted in Figure 2, has far more complex dynamics. In fact, the structural break model used here is incapable of capturing them. The full series is highly persistent and clearly non-stationary, as evidenced by the ADF t-statistic of $-.12$. There is not a statistically significant structural break in the mean—the sup-Wald statistic of 13.84 (estimated at a break date of 1990) is beneath the 10 percent critical value of 16.14 .

Insert Figure 2 here

From Figure 2, it is clear that the series has some structural instability after 1982. It is not particularly persistent over 1953-82—qualitatively, that part of the series is very similar to the first thirty years of the “invalid” SHARE series plotted in Figure 1. However, the SHARE of “not infringed” grows consistently from the early 1980s on. Thus, the onset of

the CAFC coincides with a surge in “not infringed” decisions despite no significant change in the rate of affirmation of such decisions. Our model, which does not permit trend growth, is not a good description of the dynamics of this series after 1982. We discuss this further later in the paper.

The unit root test rejects the null hypothesis and the structural break test identifies no significant break in the “valid and infringed” APPEAL series. Like the APPEAL series for “not infringed” patents, this series has no significant persistence. Thus, the rate of appeal of “valid and infringed” patents experienced little change under the CAFC.

Insert Figure 3 here

Both the AFFIRM and SHARE series, plotted together in Figure 3, have dynamics too complex for our structural break model. The unit root null cannot be rejected for either series, and neither series contains a significant structural break. Just as with the SHARE series for “not infringed” patents, these two series are not particularly persistent through 1982, but experience dramatic changes afterwards. Most notably, each series experiences a jump that persists through the late 1980s, then experiences a drop. The AFFIRM series falls back to around its pre-CAFC level, while the SHARE series falls beneath its pre-CAFC level.

The dynamics of the “valid and infringed” series clearly change after 1982; there were at least two significant changes in both AFFIRM and SHARE series. Ostensibly, the apparent early rise was favorable to patentees, while the subsequent fall was unfavorable. However, as the structural break model permits neither trend growth nor multiple breaks, it does not provide a good description of the dynamics of these series after 1982.

4.2. Discussion

Litigated patents, particularly those that reach appeal, are endogenously determined, so the estimated breaks are potentially subject to equilibrium effects. Namely, when the CAFC imposes a new patent law, we expect to see changes in its affirmation rates but also expect changes in the number of and distribution of litigated patents. When the “new” distribution of cases works its way through, this may dramatically affect the share of “invalid” decisions, the rate of appeal and the rate of affirmation. It may also indirectly affect the patent law, which must adjust to this new distribution of cases. These effects will tend to add persistence to the time series and potentially introduce multiple structural breaks.

Based on the results of the previous section’s analysis of the three “invalid” series, we conclude that the CAFC’s stronger presumption of validity had an immediate and permanent effect on the ultimate likelihood of patent validity. The CAFC has affirmed “invalid” decisions significantly less often, patentees have appealed “invalid” decisions significantly more often, and district courts have ruled patents “invalid” significantly less often. Thus, all three measures indicate a more favorable environment for patentees. Since the three “invalid” series include very little persistence and the structural breaks are highly significant, this indicates that for the validity inquiry, the adjustment period was nearly immediate.

Clearly, this is not the case for the infringement inquiry. Despite no significant change in the rate of affirmation of “not infringed” decisions, there has been a surge in the share of such decisions, most of which occurred after 1990. Among “valid and infringed” decisions, there was an initial surge in the affirmation rate, which is intuitively consistent with the drop in the affirmation rate of “invalid” decisions. However, this surge was short-lived—the affirmation rate in the 1990s has been similar to that in the pre-CAFC era.

The data reveal two trends in infringement outcomes that are attributable to the CAFC’s stronger presumption of validity. First, “not infringed” decisions are clearly more likely after 1982 (recall Figure 2). However, unlike the share of “invalid” decisions, the change in the

share of “not infringed” occurred over a number of years. Second, the affirmation of “valid and infringed” decisions has been more likely to hinge on the infringement inquiry under the CAFC. In the pre-CAFC era, 71% of appellate non-affirmations of “valid and infringed” decisions were reversals to “invalid,” while only 16% were reversals to “not infringed.” Under the CAFC, by contrast, more non-affirmations were reversals to “not infringed” (40% to 29%).⁵⁴ Interestingly, this disparity is most notable from 1993 on. From 1983-92, more non-affirmations were reversals to “invalid” (41% to 32%). From 1993-2002, more were reversals to “not infringed” (44% to 22%).

Taken together, these statistical features suggest the following explanation. The CAFC enforced a stronger presumption of validity in 1982, reducing the likelihood of an affirmation of an “invalid” decision. The major changes to the validity inquiry were immediate and well understood. As a result, district courts quickly and sharply reduced the frequency of “invalid” decisions and increased the frequency of “valid and infringed” decisions.

This yielded a significant secondary effect, however—with increased frequency, infringement has become the pivotal inquiry in patent cases. When a patent is decided “invalid,” as was more frequently the case prior to the CAFC, infringement need not be considered. Thus, under the CAFC, courts have addressed infringement far more often. This increased attention has exerted additional stress on the law on infringement, producing significant adjustments since 1982. These adjustments, in turn, contributed to the drop in the AFFIRM and SHARE series for “valid and infringed” patents in the early 1990s.

One prominent example of innovation to the law on infringement is the CAFC’s handling of the “doctrine of equivalents,” under which patent scope may be extended beyond the literal language of the claims. When the CAFC was established, the precedent for the application

⁵⁴Clearly, the CAFC has vacated and/or remanded district court decisions more frequently than its predecessors. Among remands, the court has instructed the district court to consider only infringement with greater frequency than it has instructed them to consider only validity.

of the doctrine was *Graver Tank v. Linde Air Products Inc.* (339 U.S. 605 [1950]), the most widely cited patent case: “If two devices do the same work in substantially the same way, and accomplish substantially the same result, they are the same, even though they differ in name, form or shape” (339 U.S. at 608).⁵⁵ According to Dunner and Jakes (1993), the CAFC’s early application of the doctrine was similar to that of courts prior to 1982 but, eventually, there emerged “a division among the judges...and arguably inconsistent decisions between panels” (Dunner and Jakes 1993, p. 857).

In *Pennwalt Corp. v. Durand-Wayland* (833 F.2d 931, 4 U.S.P.Q.2d 1737 [Fed. Cir. 1987]), the CAFC restricted the scope of the doctrine of equivalents by requiring that whenever a claim includes multiple “elements” (e.g. a function and a result), each element must be infringed, either literally or by an equivalent, for the claim to be infringed under the doctrine. Thus, a court cannot find infringement “on the whole.” This, along with several other decisions in the late 1980s and early 1990s, led Harmon (1992, p. 576) to conclude “the court’s approach to the doctrine of equivalents, taken as a whole, is somewhat inconsistent with its reputation as a pro-patent tribunal.”

In *Festo Corp. v. Shoketsu Kinzoku Kogyo Kabushiki Corp.* (234 F.3d 558, 56 U.S.P.Q.2d 1865 [Fed. Cir. 2000]) the CAFC further restricted the application of the doctrine of equivalents. Specifically, the CAFC opinion forbids its application to any claim amended during patent prosecution. The Supreme Court tempered this restriction but did not remove it entirely.⁵⁶

5. CONCLUSION

⁵⁵This theory of the doctrine, restated in *Graver Tank*, comes originally from the opinion in *Union Paper-Bag Machine Co. v. Murphy* (97 U.S. 120, 125 [1877]).

⁵⁶The Supreme Court opinion (535 U.S. 722, 62 U.S.P.Q.2d 1705 [2002]) states that an amendment during prosecution may give rise to a presumption against not applying the doctrine of equivalents, but may not serve as an absolute bar to a finding of infringement under the doctrine.

In sum, we find that the CAFC has been pro-patent, but only with respect to validity. Their stronger presumption of validity has decreased both the likelihood of invalidity and the frequency with which validity is the pivotal inquiry. We estimate that patentees have been nearly three times more likely to achieve some success in overcoming a district court decision of “invalid” in an appeal in the CAFC era. In rationally responding to the CAFC’s stronger presumption of validity, district courts have found patents invalid about half as often as before and patentees appealed about 25% more “invalid” decisions.

As a result, the infringement inquiry has taken on greater importance and most of the CAFC’s innovations to patent law have affected patent scope. Although these changes have contributed to persistence and possibly multiple structural breaks in the “not infringed” and “valid and infringed” time series, it is not clear whether they have, on balance, favored patentees. These results are particularly intriguing in light of research finding that inventors began to seek more patents after the establishment of the CAFC (Hall and Ziedonis, 2001) and began to litigate their patents more as well (Merz and Pace, 1994).⁵⁷ Identifying how these effects fit together, as part of the transition to a new equilibrium under the CAFC, forms an important and interesting topic for future research.

The CAFC’s changes have clearly enhanced the value of patents and have increased the incentives of patentees to sue for infringement. Patent litigation is more likely to succeed, *ceteris paribus*, when a patent is less likely to be invalidated. Additionally, the patentee incurs less risk in filing an infringement suit. The patentee knows that a lost suit will

⁵⁷Specifically, Merz and Pace identify significant structural increases in both the mean and the trend of new patent litigation filings in September 1982. The instantaneous mean break is robust to *trend* growth in the population of patents, for which Merz and Pace do not control. However, the increased trend in new patent litigation filings, post-1982, may be due to this force. Lanjouw and Schankerman (2004), for instance, show that the unconditional likelihood of patent litigation did not increase, with time, across cohorts of patents with application years from 1978-95. However, given their specification of cohort by application year, that it routinely takes multiple years for a patent to be issued on an application, and that it may take longer still for litigation to be initiated, their sample includes a relatively small number of cases filed prior to the establishment of the CAFC. Thus, further research may be necessary to identify whether there was a change in trend growth of *per capita* litigation resulting from this court.

more likely be based on a finding of “not infringed,” in which case he may keep some of the patent’s future protection (if found “invalid,” its future protection would be lost with certainty). Finally, with the court focusing more on infringement, inventors have a greater incentive to build portfolios of patents to minimize gaps in their protection of products and processes. This is consistent with the finding of Hall and Ziedonis (2001) that R&D managers in the US semiconductor industry believe that the CAFC has spurred patent portfolio races. We look forward to further progress in this area.

REFERENCES

- Allison, John R. and Mark Lemley. 1998. Empirical Evidence on the Validity of Litigated Patents. *AIPLA Quarterly Journal* 26:185-275.
- Andrews, Donald W. K. 1993. Tests for Parameter Instability and Structural Change with Unknown Change Point. *Econometrica* 61:821-56.
- Baum, Lawrence. 1974. The Federal Courts and Patent Validity: An Analysis of the Record. *Journal of the Patent Office Society* 56:758-87.
- Clermont, Kevin M. and Theodore Eisenberg. 1992. Trial by Jury or Judge: Transcending Empiricism. *Cornell Law Review* 77:1124-77.
- Dreyfuss, Rochelle Cooper. 1989. The Federal Circuit: A Case Study in Specialized Courts. *New York University Law Review* 64:1-78.
- Dunner, Donald R. 1985. Introduction to The United States Court of Appeals for the Federal Circuit—The First Three Years. *AIPLA Quarterly Journal* 13:185-94.
- Dunner, Donald R. and J. Michael Jakes. 1993. The Equitable Doctrine of Equivalents. *Journal of the Patent and Trademark Office Society* 74:857-74.

- Dunner, Donald R., J. Michael Jakes, and Jeffrey D. Karceski. 1995. A Statistical Look at the Federal Circuit's Patent Decisions: 1982-1994. *The Federal Circuit Bar Journal* 5:151-79.
- Federico, P.J. 1956. Adjudicated Patents, 1948-54. *Journal of the Patent Office Society* 38:233-49.
- Goldstein, Paul. 1993. *Copyright, Patent, Trademark and Related State Doctrines: Cases and Materials on the Law of Intellectual Property*. Westbury, New York: The Foundation Press, Inc.
- Hall, Bronwyn H. Forthcoming. Exploring the Patent Explosion. *Journal of Technology Transfer*.
- Hall, Bronwyn H. and Rosemarie Ham Ziedonis. 2001. The Patent Paradox Revisited: An Empirical Study of Patenting in the US Semiconductor Industry, 1979-1995. *RAND Journal of Economics* 32:101-28.
- Harmon, Robert L. 1992. Seven New Rules of Thumb: How the Federal Circuit Has Changed the Way Patent Lawyers Advise Clients. *George Mason University Law Review* 14:573-84.
- Jolls, Christine, Cass R. Sunstein and Richard Thaler. 1998. A Behavioral Approach to Law and Economics. *Stanford Law Review* 50:1471-1550.
- Kastriner, Lawrence G. 1991. The Revival of Confidence in the Patent System. *Journal of the Patent and Trademark Society* 73:5-23.
- Kessler, Daniel, Thomas Meites and Geoffrey Miller. 1996. Explaining Deviations from the Fifty Percent Rule: A Multimodal Approach to the Selection of Cases for Litigation. *Journal of Legal Studies* 25:233-59.

- Koenig, Gloria. 1980. *Patent Invalidity: A Statistical and Substantive Analysis*. New York, NY: Clark Boardman Company.
- Kortum, Samuel and Josh Lerner. 1998. Stronger Protection or Technological Revolution: What Is Behind the Recent Surge in Patenting? *Carnegie-Rochester Series on Public Policy* 48:247-304.
- Lanjouw, Jean O. and Mark Schankerman. 2001. Characteristics of Patent Litigation: A Window on Competition. *RAND Journal of Economics* 32:129-51.
- Lanjouw, Jean O. and Mark Schankerman. 2004. Protecting Intellectual Property Rights: Are Small Firms Handicapped? *Journal of Law and Economics* 47:45-74.
- Marco, Alan C. 2005. Learning by Suing: Structural Estimates of Court Errors in Patent Litigation. Working Paper No. 68, Vassar College Working Paper series, February.
- Merges, Robert P. 1997. *Patent Law and Policy*. Charlottesville, VA: MICHIE Law Publishers.
- Merz, Jon F. and Nicholas M. Pace. 1994. Trends in Patent Litigation: The Apparent Influence of Strengthened Patents Attributable to the Court of Appeals for the Federal Circuit. *Journal of the Patent and Trademark Office Society* 76:579-90.
- Moore, Kimberly. 2000. Judges, Juries and Patent Cases: An Empirical Peek Inside the Black Box. *Michigan Law Review* 98:365-406.
- Moore, Kimberly. 2001. Forum Shopping in Patent Cases: Does Geographic Choice Affect Innovation? *North Carolina Law Review* 79:889-935.
- Moore, Kimberly. 2003. Xenophobia in American Courts. *Northwestern University Law Review* 97:1497-1550.
- Perron, Pierre. 1989. The Great Crash, the Oil Price Shock and the Unit Root Hypothesis.

- Econometrica* 57:1361-1401.
- Priest, George L. and Benjamin Klein. 1984. The Selection of Disputes for Litigation. *Journal of Legal Studies* 13:1-55.
- Quillen, Jr., Cecil D. 1993. Proposal for the Simplification and Reform of the United States Patent System. *AIPLA Quarterly Journal* 21:189-212.
- Turner, John L. 2005. In Defense of the Patent Friendly Court Hypothesis: Theory and Evidence. Unpublished manuscript. University of Georgia, Department of Economics, April.
- Vogelsang, Timothy J. 1997. Wald-Type Tests for Detecting Shifts in the Trend Function of a Dynamic Time Series. *Econometric Theory* 13:818-49.
- Zivot, Eric and Donald W.K. Andrews. 1992. Further Evidence on the Great Crash, the Oil-Price Shock, and the Unit Root Hypothesis. *Journal of Business and Economic Statistics* 10:251-70.

Table 1. District Court Decisions 1953-2002

Circuit	N	Invalid (Pr)	Appealed (Pr)	Not Infringed (Pr)	Appealed (Pr)	Valid and Infringed (Pr)	Appealed (Pr)
CAFC	988	.279	.688	.352	.592	.368	.618
Pre-CAFC	2,327	.557	.547	.130	.482	.313	.562
First	145	.641	.634	.117	.471	.241	.800
Second	328	.634	.457	.140	.326	.226	.568
Third	251	.717	.528	.116	.483	.167	.619
Fourth	193	.461	.708	.166	.563	.373	.681
Fifth	232	.401	.613	.159	.649	.440	.480
Sixth	237	.616	.521	.122	.517	.262	.597
Seventh	438	.495	.530	.094	.512	.411	.544
Eighth	120	.600	.528	.192	.435	.208	.680
Ninth	293	.563	.545	.126	.432	.311	.407
Tenth	68	.324	.727	.103	.286	.574	.538
Other	22	.455	.400	.227	.600	.318	.857

Note. These statistics reflect a complete set of “invalid,” “not infringed” and “valid and infringed” district court patent decisions that were rendered during 1953-2002 and are published in the *United States Patents Quarterly*. Decisions are grouped by the circuit of the relevant appeals court. Each decision is for one patent in one case - *e.g.*, when there are two patents in a case, we record two decisions. When neither validity nor infringement of the patent factors in the decision, it is excluded from the present analysis (our criteria for exclusion are described in section 3.1). The numbers in columns 2, 4 and 6 give the estimated probabilities of “invalid,” “not infringed” and “valid and infringed” decisions, conditional on one of these three decisions being rendered. The numbers in columns 3, 5 and 7 give estimated conditional probabilities of appeal. For example, in districts located in the First circuit from 1953 until the establishment of the CAFC, 93 of 145 decisions (or 64.1%) were “invalid,” and we record an appeals decision for 59 of these 93 (63.4%).

Table 2. Appeals Court Decisions 1953-2002

Circuit	Invalid (N)	Affirmed (Pr)	Not Infringed (N)	Affirmed (Pr)	Valid and Infringed (N)	Affirmed (Pr)
CAFC	415	.573	460	.685	466	.723
Pre-CAFC	1,068	.850	232	.694	627	.603
First	66	.909	8	.750	29	.552
Second	109	.899	21	.571	53	.434
Third	112	.866	14	.714	28	.607
Fourth	78	.833	19	.789	54	.574
Fifth	108	.778	51	.549	84	.690
Sixth	113	.912	30	.900	69	.652
Seventh	194	.851	34	.735	147	.619
Eighth	51	.902	8	.750	20	.450
Ninth	190	.816	37	.676	95	.558
Tenth	40	.725	4	.250	35	.829
Other	7	.857	6	1.00	13	.462

Note. These statistics reflect appellate decisions, rendered during 1953-2002, that address “invalid,” “not infringed” and “valid and infringed” district court decisions. All appellate decisions published in the *United States Patents Quarterly* are included, as well as (a small number of) unpublished appellate decisions in cases where the district court decision is published. Each decision is for one patent in one case - *e.g.*, when there are two patents in a case, we record two decisions. When neither validity nor infringement of the patent factors in the decision, it is excluded from the present analysis (our criteria for exclusion are described in section 3.1). Columns 1, 3 and 5 give the number of appeals decisions in each category in our data. Columns 2, 4 and 6 give estimated probabilities that the appeals court affirms a district court decision, conditional on whether that decision was “invalid,” “not infringed” or “valid and infringed,” respectively.

Table 3. Unit Root Tests

	<u>Invalid</u>			<u>Not Infringed</u>			<u>Valid and Infringed</u>		
	AFFIRM	APPEAL	SHARE	AFFIRM	APPEAL	SHARE	AFFIRM	APPEAL	SHARE
α	.101 (.098)	.347 (.117)	.040 (.039)	.777 (.107)	.543 (.111)	.021 (.023)	.267 (.125)	.695 (.090)	.082 (.053)
ρ	-.163 (.129)	-.569 (.192)	-.106 (.085)	-1.117 (.146)	-1.011 (.206)	-.010 (.089)	-.409 (.191)	-1.210 (.152)	-.266 (.154)
β_1	-.722 (.160)	-.397 (.136)	-.522 (.145)		.280 (.180)	-.815 (.157)	-.703 (.183)		-.357 (.152)
β_2	-.331 (.141)		-.305 (.133)		.398 (.144)	-.351 (.153)	-.344 (.138)		
R^2	.472	.558	.312	.550	.446	.428	.604	.573	.261
T	47	47	47	49	46	47	47	48	48
ADF-t	-1.27	-2.96*	-1.25	-7.67**	-4.91**	-.12	-2.14	-7.94**	-1.73

Note. These results reflect regressions using the equation $\Delta y_t = \alpha + \rho y_{t-1} + \sum_{j=1}^{j^{max}} \beta_j \Delta y_{t-j} + \epsilon_t$. All time series $\{y_t\}$ are annual: the AFFIRM and SHARE series run from 1953-2002, while the APPEAL series run from 1953-2001. The null hypothesis of the augmented Dickey-Fuller (ADF) test is that there is a unit root in the time series, or $\rho = 0$. Standard errors are in parentheses. Statistical significance is represented for the unit root tests only and with the following convention: * Significant at the 5% level, ** Significant at the 1% level.

Table 4. Structural Break Tests

	<u>Invalid</u>			<u>Not Infringed</u>			<u>Valid and Infringed</u>		
	AFFIRM	APPEAL	SHARE	AFFIRM	APPEAL	SHARE	AFFIRM	APPEAL	SHARE
α	.961 (.118)	.697 (.078)	.500 (.068)	.794 (.103)	.584 (.105)	.107 (.030)	.250 (.122)	.812 (.104)	.193 (.058)
α'	-.319 (.048)	.172 (.038)	-.239 (.037)	-.262 (.135)	.216 (.082)	.219 (.059)	-.073 (.045)	-.108 (.053)	-.114 (.034)
ρ	-1.134 (.136)	-1.269 (.138)	-.910 (.121)	-1.126 (.140)	-1.135 (.197)	-.664 (.181)	-.355 (.189)	-1.290 (.151)	-.538 (.160)
β_1					.359 (.170)	-.314 (.139)	-.759 (.181)		-.264 (.139)
β_2					.434 (.135)		-.384 (.137)		
R^2	.591	.644	.540	.583	.520	.509	.625	.607	.405
T	49	48	49	49	46	48	47	48	48
sup- Wald	43.81**	20.55**	42.23**	3.77	6.93	13.84	2.65	4.06	11.32
Year	1983	1982	1982	2000	1996	1990	1991	1969	1993
$\frac{\alpha}{-\rho}$.847	.549	.549	.705	.515	.161	.703	.630	.358
$\frac{\alpha+\alpha'}{-\rho}$.566	.685	.287	.473	.705	.491	.496	.547	.147

Note. These results reflect regressions using the equation $\Delta y_t = \alpha + \alpha' 1(t > T_B) + \rho y_{t-1} + \sum_{j=1}^{j^{max}} \beta_j \Delta y_{t-j} + \epsilon_t$. All time series $\{y_t\}$ are annual: the AFFIRM and SHARE series run from 1953-2002, while the APPEAL series run from 1953-2001. The null hypothesis of the sup-Wald test is that there is no structural break in the time series, or $\alpha' = 0$. Standard errors are in parentheses. Statistical significance is represented for the sup-Wald tests only and with the following convention: * Significant at the 5% level, ** Significant at the 1% level.

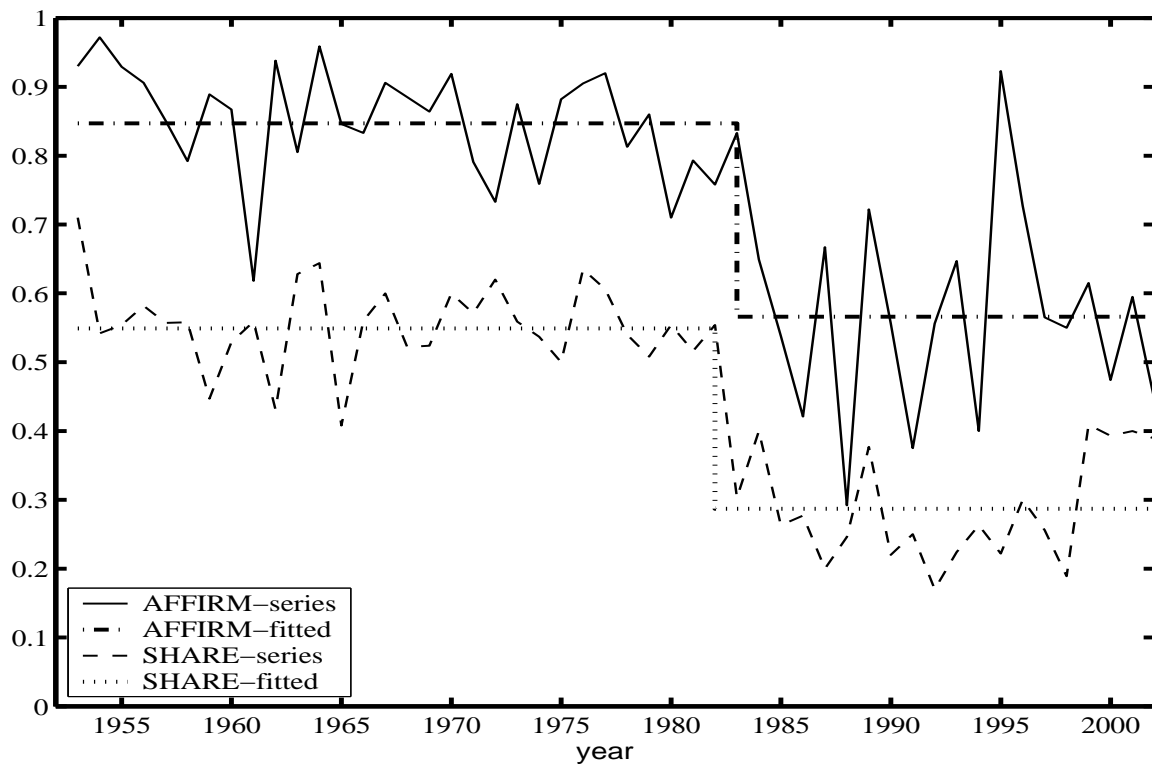


Figure 1: *“Invalid” Series*

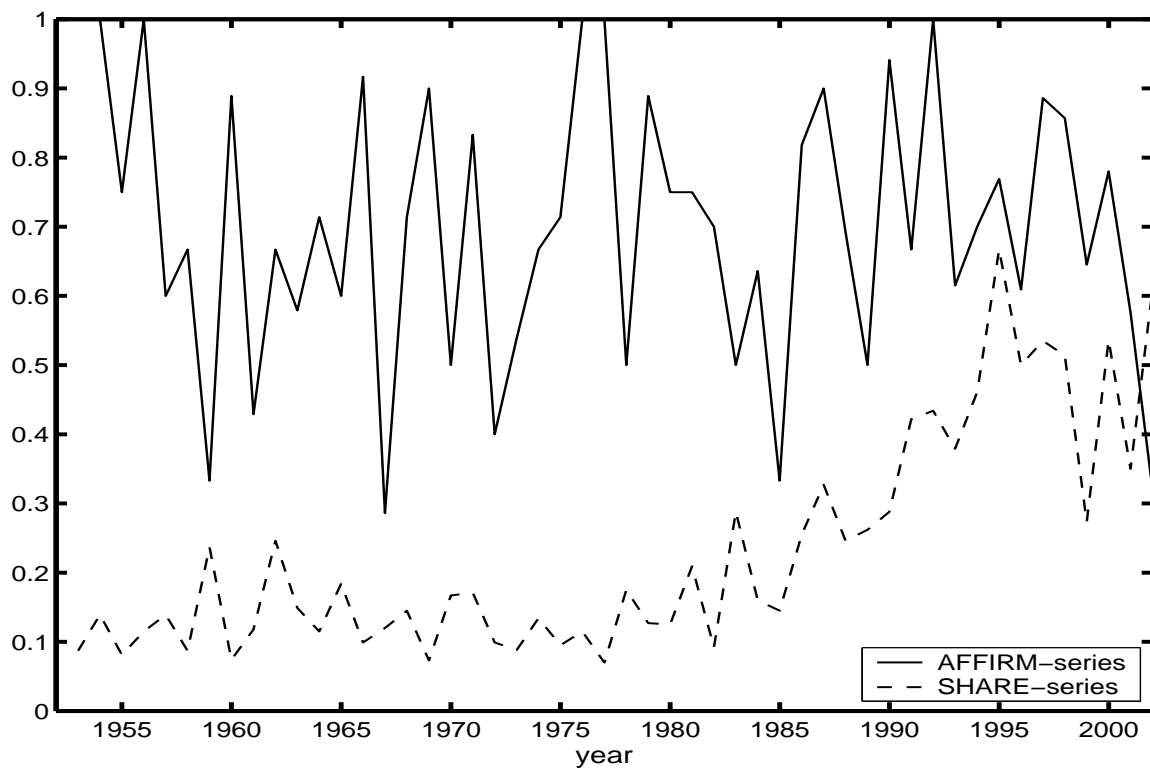


Figure 2: "Not Infringed" Series

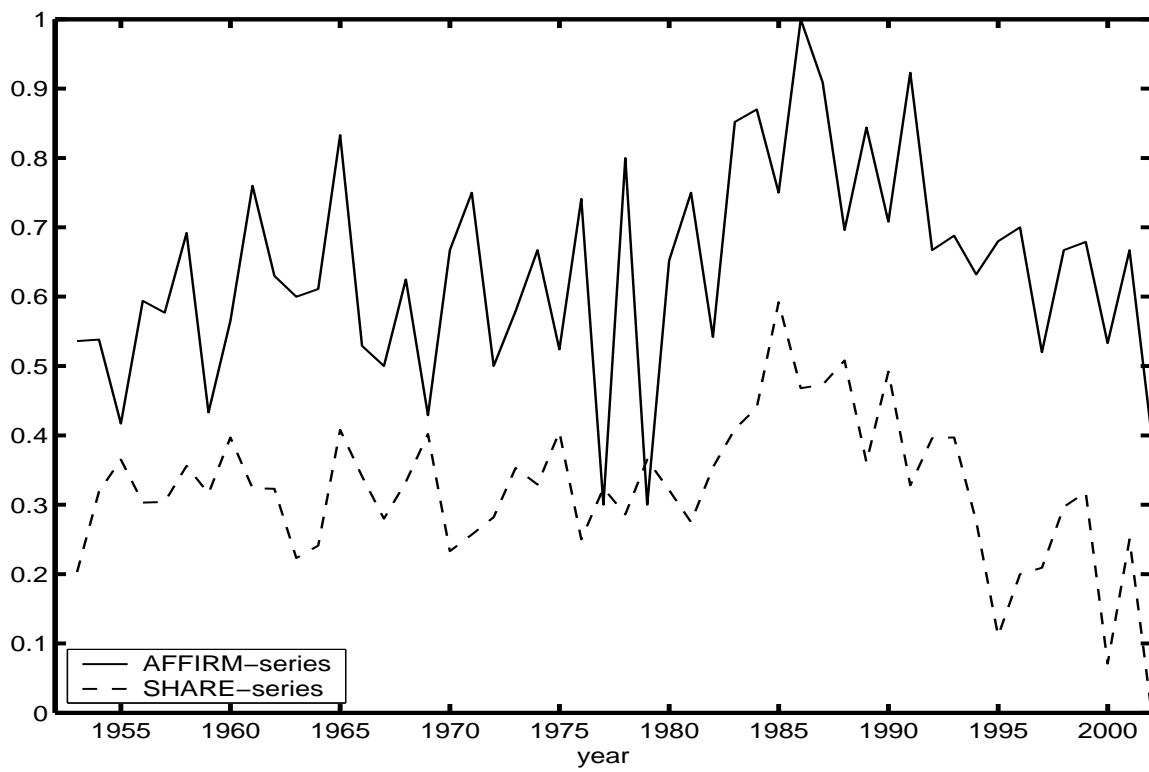


Figure 3: “Valid and Infringed” Series