BOSTON UNIVERSITY SCHOOL OF LAW

WORKING PAPER SERIES, LAW AND ECONOMICS
WORKING PAPER NO. 07-08

THE PRIVATE COSTS OF PATENT LITIGATION

JAMES E. BESSEN AND MICHAEL J. MEURER

The Private Costs of Patent Litigation

Working Paper
Original version: March 2006
Current version: February 2008

By James Bessen and Michael J. Meurer*

Abstract: This paper estimates the total cost of patent litigation. We use a large sample of stock market event studies around the date of lawsuit filings for US public firms from 1984-99. Even though most lawsuits settle, we find that the total costs of lawsuits are large compared to estimated legal fees, estimates of patent value, and R&D spending. By the late 1990s, alleged infringers bore expected costs of over $16 billion per year. These estimates support the view that infringement risk should be a major concern of policy.

Keywords: patent, litigation, litigation cost, property rights

JEL Classifications: O31, O34, K41

*Research on Innovation and Boston University School of Law, and Boston University School of Law, respectively. Thanks for comments to Megan MacGarvie, Jesse Giummo, John Turner and participants at the IIOC and CELS conferences, and seminars at Harvard, Stanford, the NBER, the NBER Summer Institute. Thanks also to research assistance from Debbie Koker and Dan Wolf.

Contact: jbessen@bu.edu

1 Introduction

Like any regulatory mechanism, the patent system has benefits and costs, both private and social. Yet little empirical evidence exists about the magnitude of some of these costs, leaving policy analysts to sometimes rely on guesswork. For example, recent policy analysis of a patent opposition proceeding in the US (Levin and Levin 2002, Hall et al. 2004) has been based on rough estimates of the costs of patent litigation and the social costs of inappropriately-granted patents.

In contrast, there is a significant literature estimating benefits of the patent system, especially private benefits in the form of estimates of patent value[1] or of the patent premium (Arora et al. 2005). However, without comparable estimates of private and social costs, it is difficult to conduct either analysis of specific policy changes or a normative analysis of the patent system in comparison to other means of encouraging innovation. For example, Schankerman (1998) suggests that the ratio of aggregate patent value to R&D constitutes an upper bound measure of the subsidy that patents provide to R&D. He asserts that this ratio can be used to compare patents to other forms of appropriating returns on invention. But surely this is only an estimate of a *gross* subsidy against which private costs of patents need to be netted out.

This paper takes a step toward quantifying costs by estimating the private costs of patent litigation. Using event study methodology to analyze patent lawsuit *filing* we find the expected joint loss to the litigating parties is large, and probably much larger than the expected attorneys' fees. This result is a bit of a surprise because most patent lawsuits settle short of trial, and thus it might seem that average patent litigation costs would not be large.

But attorneys' fees can be high even when a patent lawsuit settles before trial, and the indirect costs of litigation can also be high before trial. Indirect business costs of patent litigation take many forms. Business can be disrupted as managers and researchers spend their time producing documents, testifying in depositions, strategizing with lawyers, and appearing in court. Litigation strains the relationship between the two parties and may jeopardize cooperative development of the patented technology or cooperation on some other front. Firms in a weak financial position might see their credit costs soar because of possible bankruptcy risk created by patent litigation.

Alleged infringers face additional costs. Preliminary injunctions can shut down production and sales while the litigation pends. But even without a preliminary injunction, customers may stop buying a product. Frequently, products require customers to make

[1] This literature began with Pakes and Schankerman (1984). See Bessen (2008) for a survey of this literature.

complementary investments; they may not be willing to make these investments if a lawsuit poses some risk that the product will be withdrawn from the market. Furthermore, patent owners can threaten customers and suppliers with patent lawsuits because patent infringement extends to every party who makes, uses, or sells a patented technology without permission, and sometimes to those who participate indirectly in the infringement. Some of these costs persist after settlement.

Even simple delay can impose large business costs. Consider, for example, litigation against Cyrix, a startup firm that introduced Intel-compatible microprocessors. Intel, the dominant microprocessor maker, sued Cyrix and the suit lasted a year and a half. During that time Cyrix had difficulty selling microprocessors to computer manufacturers, who were almost all also customers of Intel and who were reluctant to break ranks to go with a product that might be found to infringe. In the meantime, Intel responded by accelerating its development of chips that would compete against Cyrix's offerings. In the end, Cyrix won the lawsuit, but lost the war, having lost much of its competitive advantage. In effect, Cyrix lost the window of opportunity to establish itself in the marketplace. Litigation exacted a heavy toll indeed.

Although we explore the costs of litigation to both patent owners and alleged infringers in some detail, our chief interest is with the cost to alleged infringers. We choose this focus because innovators experience the patent system both as patent owners and as alleged infringers. Empirical methods that measure patent value by studying patent renewal or stock market valuation of patent portfolios account for the expected cost of enforcing patents through litigation.[2] Unfortunately, there are no studies that quantify the negative impact of patent litigation cost on alleged infringers.

To the extent that costly patent litigation is mainly the result of inadvertent infringement —and we argue elsewhere that it is (Bessen and Meurer 2005, 2006, 2008)—then the costs of defending against inadvertent infringement represent a disincentive to investing in innovation.[3] The risk of unavoidable infringement acts like a "tax" on innovation. We fear this tax has grown in recent years because we found that during the 1990s there was a dramatic increase in the hazard of patent litigation for publicly traded firms (Bessen and Meurer 2005).

The event study methodology has been used before to study litigation, beginning with Cutler and Summers (1988) in the context of litigation over a merger. Several papers have

[2] Nevertheless, it is useful to know how much patent value is eaten away by patent litigation, and what sort of reforms might reduce patent enforcement costs. Answers to those questions will have to wait for future research.

[3] These costs include the deadweight losses described above and also the settlement transfer from an innocent innovator/infringer to the patent owner.

performed event studies of patent litigation, both the event of the initial filing and the terminating event (settlement, judgment or verdict), including Bhagat, et al. (1994), Lerner (1995), Bhagat et al. (1998), Lunney (2004) and Haslem (2005).

These studies of initial filings, however, do not provide the best estimates from which to calculate the aggregate risk of infringement to the firms that perform R&D. They use small, selective samples and their estimates of wealth loss are not especially precise. Our contribution is to work with a much larger set of disputes: our sample covers most patent lawsuits filed against US public firms from 1984 through 1999, a sample responsible for the lion's share of R&D spending. This gives our results greater precision and also makes them highly representative of R&D-performing firms, permitting us to calculate a variety of cost and risk measures to inform policy. We find, in fact, that the estimates of wealth loss reported in some earlier studies appear to be overstated.

A key assumption of this literature is that the change in firm value that occurs around a lawsuit filing reflects investors' estimates of the direct and indirect effects of the lawsuit on the profits of the firm on average and do not systematically reflect any unrelated information. We show evidence below that the revelation of unrelated information does not overstate our estimates for defendants in infringement suits and that, therefore, we may associate the loss in wealth with the effective total cost of litigation for defendants.

We find that alleged infringers lose about half a percentage point of their stock market value upon being sued for patent infringement. This corresponds to a mean cost of $28.7 million in 1992 dollars (median of $2.9 million), much larger than mean legal fees of about half a million. In aggregate, infringement risk rose sharply during the late 1990s, exceeding $16 billion for US public firms. This amounts to 19% of these firms' R&D spending, a ratio that exceeds some estimates of the value of patents granted relative to R&D.

The next section describes the data and methods used for estimating cumulative abnormal returns. Section 3 reports average returns and some analysis of factors that affect returns. Section 4 calculates litigation cost, Section 5 calculates some broader measures of infringement risk and Section 6 concludes.

2 Data and Methods

2.1 Data Sources

Our research matches records from three data sources: lawsuit filings from Derwent's Litalert database, firm financial data from Compustat, and CRSP data on securities prices. In addition, we searched the electronic archives of the Wall Street Journal to locate any articles announcing lawsuit filings and also any announcements of other events that might confound our analysis.

Using these sources, we constructed two main samples. The first, small sample included just those lawsuits where we could identify one or more parties on both sides of the dispute as public firms. The second, large sample included all cases where the alleged infringer (defendant in an infringement suit or plaintiff in a declaratory action) was a publicly traded firm, but the patentee litigant need not be public.

Our primary source of information on lawsuit filings is Derwent's Litalert database, a database that has been used by several previous researchers (Lanjouw and Schankerman, 2004, Ziedonis, 2004, Bessen and Meurer 2005). Federal courts are required to report all lawsuits filed that involve patents to the U.S. Patent and Trademark Office (USPTO) and Derwent's data is based on these filings. Beginning with the Derwent data from 1984 through 2000, we removed duplicate records involving the same lawsuit as identified by Derwent's cross-reference fields. We also removed lawsuits filed on the same day, with the same docket number and involving the same primary patent. Sometimes firms respond to lawsuits by filing counter-suits of their own, perhaps involving other patents. Since our main focus is on initial disputes rather than on lawsuit filings *per se,* we also removed filings made within 90 days of a given suit that involved the same parties.

The Derwent data does not distinguish whether the suit filed is an infringement suit or a declaratory judgment suit. A firm threatened with an infringement suit can file a declaratory action which aims for a judgment that the patent is uninfringed or invalid. To classify each suit, we first identified whether the patent assignee of the main patent at issue matched one of the parties to the suit. If the assignee matched a plaintiff, the suit was classified as an infringement suit; if the assignee matched a defendant, the suit was classified as a declaratory action. We were able to match the assignee for 83% of the suits, and of these, only 17% were declaratory actions.[4]

[4] These numbers are quite similar to findings by Moore (2000) and Lanjouw and Schankerman (2004).

If the assignee did not match a party to the suit, then it was classified as an infringement suit because there are relatively few declaratory actions. This classification then allowed us to identify whether the subject firm was a "patentee litigant" (that is, plaintiff in an infringement suit or defendant in a declaratory action) or an "alleged infringer" (the reverse).

To explore characteristics of firms involved in these lawsuits, we matched the listed plaintiffs and defendants to the Compustat database of U.S. firms from 1984-99 that report financials (excluding American Depository Receipts of foreign firms traded on US exchanges). These data were based on merged historical data tapes from Compustat and involved an extensive process of tracking firms through various types of re-organization and eliminating duplicate records for firms (e.g., consolidated subsidiaries listed separately from their parent companies).[5]

The lawsuit data were matched to the Compustat data by comparing the litigant names with all domestic firm names in Compustat and also a list of subsidiary names used in Bessen and Hunt (2007).[6] To check the validity and coverage of this match, we randomly selected a number of parties to suits and then checked them manually using various databases including PACER, LexisNexis, the Directory of Corporate Affiliations and the LexisNexis M&A databases. Although we were not able to definitively identify all parties, the rate of false positives was not more than 3% (no more than 5 of 165 parties were found to have been falsely matched) and the rate of false negatives was no more than 7% (no more than 34 of 502 public companies were not matched). The Compustat firms were then also matched to the CRSP file of daily security prices.

We identified 2,648 suits with sufficient data on alleged infringers, some of these having multiple alleged infringers, for a total of 2,887 events in our large sample. We also selected all those lawsuits where we could identify at least one party on each side as a publicly listed firm. This left us with asample of 750 plaintiffs and 747 defendants in lawsuits where public firms were parties on both sides.

Summary statistics of our samples are shown in Table 1 and further details from a closely related sample are reported in Bessen and Meurer (2005). Parties to patent lawsuits tend to be larger than average firms with large R&D budgets. Moreover, our large sample captures the

[5] This work was conducted by Bob Hunt and Annette Fratantaro at the Federal Reserve Bank of Philadelphia for an earlier project and we thank them for graciously sharing it with us.

[6] A software program identified and scored likely name matches, taking into account spelling errors, abbreviations, and common alternatives for legal forms of organization. These were then manually reviewed and accepted or rejected. Note that this match is based on the actual parties to litigation, not the original assignee of the patent at issue.

bulk of patent litigation against R&D performers. In 1999, US public firms in Compustat spent $150 billion on R&D, while total industrial R&D spending reported by the National Science Foundation was $160 billion.[7] Aside from issues of under-reporting, our large sample constitutes a comprehensive sample with which we can obtain a lower bound measure of the aggregate risk of infringement to R&D performers.

Finally, each lawsuit in the small sample was checked against the Wall Street Journal archive to identify those suits that were announced in the Journal within one month of the filing data and to identify possible confounding news about either party to the suit within one week of the filing date.

In Section 4 we discuss a supplemental dataset of lawsuits that reports legal fees.

2.2 Estimating Cumulative Abnormal Returns

To estimate the impact of a lawsuit filing on the value of a firm, we use event study methodology (see Mackinlay 1997 for a review). In particular, we use the dummy variable method described by Michael Salinger (1992).[8] This assumes that stock returns follow a market model,

(1) $$r_t = \alpha + \beta r_t^m + \epsilon_t$$

where r_t is the return on a particular stock at time t, r_t^m is the compounded return on a market portfolio, and ϵ_t is a stochastic error. If an event, such as a lawsuit filing, occurs on day T, then there may be an "abnormal return" to the particular stock on that day. This can be captured using a dummy variable,

(2) $$r_t = \alpha + \beta r_t^m + \delta I_t + \epsilon_t$$

where I_t equals 1 if $t=T$ and 0 otherwise. Equation (2) can be estimated using OLS for a single event. In practice, this equation is estimated over the event period and also over a sufficiently long pre-event window. In this paper we use a 200 trading-day pre-event window.[9] The

7 There are important differences in the scope of what was included in these two measures, nevertheless, they suggest that public firms account for the lion's share of R&D spending.

8 Salinger shows that this model is mathematically equivalent to the OLS market model described in Brown and Warner (1985) and widely used.

9 We also ran regressions with a 180 day pre-event window that ended 30 days before the lawsuit filing. Cumulative abnormal returns were very close to those with a 200 day window that last up to the day before the event window.

coefficient estimate of δ obtained by this procedure is then an estimate of the abnormal return on this particular stock. For different stocks, the precision of the estimates of δ will vary depending on how well equation (2) fits the data. The estimated coefficient variance from the regression provides a measure of the precision of the estimate of the abnormal return.

We want to obtain a representative estimate of the abnormal returns from lawsuit filings for multiple stocks, under the assumption that these represent independent events and that they share the same underlying "true" mean. Previous papers estimating abnormal returns from patent lawsuits have simply reported unweighted means for the group of firms. Although the unweighted mean is an unbiased estimator, it is not efficient. Since we are concerned with obtaining the best estimate to use in policy calculations (and not just testing the sign of the mean), we use a weighted mean to estimate the "average abnormal return," where the weight for each observation is proportional to the inverse of the variance of the estimate of δ for that firm.[10]

When we test our means against the null hypothesis that the true mean is zero, we report both the significance of *t*-tests using the weighted mean and also the significance of the *Z* statistic (see Dodd and Warner 1983), a widely used parametric test of significance that incorporates the variation in precision across events.[11] In any case, the significance test results are closely similar as are those of some non-parametric tests.

As Salinger (1992) notes, this procedure assumes that the returns for each event are independent of each other. However, when there are multiple defendants in a suit, returns may be systematically related. For example, one defendant may be a supplier to another or two defendants may be unequal rivals. Since the large sample has 188 lawsuits with multiple defendants, in these cases we estimate the returns for the defendants to each suit jointly, estimating common abnormal returns for this group of defendants.

Finally, (2) describes the abnormal return for a single day. It is straightforward to design dummy variables to estimate a "cumulative abnormal return" (CAR) over an event window consisting of multiple consecutive days. In the following, for instance, if the suit is filed on date $t=T$, then we may use a window from day T-1 to T+24.

---

[10] In any case, we find that for our entire sample, the weighted mean is quite close to the unweighted mean and also to the median. However, there are significant differences in the averages for sub-samples.

[11] The *Z* statistic is a joint test of the individual firm *t*-tests. We use a robust version described in Kramer (2001).

2.3 The Event

This paper also differs from previous research in the nature of the events we study. Previous studies have used the *announcement* of the lawsuit in a newspaper or wire service as the event. We use the *filing* of the lawsuit, instead. This may seem to be a minor difference, but it is significant for two reasons.

First, at the time of our sample, most patent lawsuits were not announced in newspapers or wire service reports at all. Various factors may influence whether a lawsuit is announced or not. Firms may choose to issue a press release or not. The SEC requires reporting of major lawsuits in quarterly and annual filings, but lawsuits will be reported separately only if they materially affect the profits of the firm. And news sources may not report all lawsuits even if the firms issue press releases.

We took a random sample of patent lawsuits against US public firms and searched LexisNexis for news stories that mention the lawsuits within one month of the filing date, before and after. We found that only 19% of the lawsuits were mentioned in the Dow Jones Newswire, one of the most comprehensive reporting services; only 7% were mentioned in the Wall Street Journal, which was used by several of the previous studies. Since one of our objectives is to tally the combined risk of lawsuits for public firms, clearly we cannot obtain comprehensive estimates relying only on announced lawsuits.

Moreover, announced lawsuits are a select group that may be qualitatively different from other lawsuits. That is, samples of announced lawsuits may suffer from sample selection bias. To test this, we performed a series of Probit regressions in our small sample on whether a lawsuit was reported in the Wall Street Journal (see Appendix). Among other things, we find that the probability of a Wall Street Journal announcement is strongly correlated with the defendant firm's stock market beta. This might reflect the editorial judgment of the Wall Street Journal that certain lawsuits are more newsworthy and more likely to affect a defendant's stock price or, perhaps, word of the lawsuit is already affecting the defendant's stock price. This, in turn, suggests that estimates made on a sample of announced lawsuits may have abnormal returns with a larger absolute magnitude than those from a more representative sample.

Below we compare estimates of abnormal returns from samples of lawsuits announced in the Wall Street Journal with estimates from our comprehensive sample. We find that our estimates from the announced sample are quite similar to those reported in the previous literature. However, these estimates are substantially larger in absolute magnitude than those for our comprehensive sample, suggesting considerable sample selection bias.

On the other hand, our estimates may be understated for another reason: investors may not receive news of the lawsuit within an event window around a filing date. With an announcement in a newspaper or major newswire, we can be reasonably sure that investors hear the news of the lawsuit within a day or two of the announcement. But we cannot be sure that investors hear the news about a legal filing in a district courthouse. Indeed, depending on how long it takes to serve papers, the defendant may not be aware of the lawsuit for a day or so after the filing date. In other words, news of an unannounced patent lawsuit filing may leak out more slowly and investors may not learn of a lawsuit within a specified event window.

We see evidence of this slower diffusion of information in the lawsuits that were announced in the Wall Street Journal. Figure 1 displays the frequency of these news stories relative to the actual court filing date. Event studies based on public announcements typically use an event window of two or three days (often one day before the announcement). Although many lawsuits are announced within two days of filing, such a small event window around a *filing date* would clearly miss a very large share of lawsuit announcements. Moreover, it seems likely (given the role of stock market beta in the likelihood of a Wall Street Journal article) that the lawsuits that are announced within a few days of the filing may be qualitatively different from those for which the news leaks out more slowly and are either announced later or not announced at all. Indeed, we find evidence within our data that stocks with beta above 1 react to the filing faster than lower beta stocks.[12] In order to have representative and comprehensive estimates, we use a longer event window than would be appropriate in an announcement event study. Specifically, we use a 25 day window (from T-1 to T+24), which, based on the data in Figure 1, should capture 96% of the announced events and, we hope, a large share of the unannounced filings. We show some CARs from shorter windows in the Appendix.

There are two possible concerns with such a longer window. First, the longer window introduces more "noise" into the estimation reducing precision and possibly attenuating the estimates. Because we have a much larger sample size than earlier studies, this is not such a significant concern and our estimates are reasonably precise, although they may be slightly attenuated. Second, research on long horizon event studies—that is, studies with *multi-year* event windows—find certain biases that arise for a variety of reasons (Barber and Lyon 1997, Kothari

[12] At day 2, the higher beta stocks for defendant firms have a CAR that is significantly lower than the CAR for lower beta stocks (at the 5% level) and the lower beta CAR is not significantly different from zero. At day 24, the CARs for these two groups are not significantly different and both are significantly different from zero. One explanation for the faster speed of diffusion for high beta stocks is that the opportunities for investors to make returns from the information about the lawsuit filing are relatively greater for these stocks.

and Warner 1997).[13] However, it seems highly unlikely that these concerns can exert a substantially greater influence in a 25 day window than they exert in a three day window.

In summary, restricting the study to events announced in news services likely introduces substantial sample selection bias. Our estimates, based on a larger window following the filing of the lawsuit, are smaller, although they might be biased toward zero.

3 Empirical Findings

3.1 Estimates of cumulative abnormal returns

Since previous studies have used samples where parties on both sides of a lawsuit were public firms and the suits were reported to newspapers or wire services, we begin by exploring a sub-sample. Table 2 shows estimates of cumulative abnormal returns for just those suits from our small, matched sample that were reported in the Wall Street Journal. In this table, we exclude suits that had a potentially confounding news story in the Wall Street Journal within a month of the suit filing date. Two previous studies have reported on event study estimates on announcements of patent lawsuits filings. Bhagat et al. (1998) examine lawsuits filed between 1981 and 1983 (51 plaintiffs and 33 defendants) and Lerner (1995) obtains estimates for 26 biotech lawsuits from 1980 to 1992.[14] To maintain consistency with the previous literature, in this Table (but not in the next) we report simple unweighted means of cumulative abnormal returns.[15] The mean and median values are reported for two different event windows, one around the Wall Street Journal publication date, the other a longer window around the actual suit filing date reported in court records (these dates occasionally differed significantly).

Consistent with most of the previous literature on litigation, we find that patentee litigants do not show a positive response to a lawsuit filing. Bhagat et al. (1998) report a CAR of -0.31%, and we find a similar value. For defendants (alleged infringers), we find a substantial loss in market value of around 2%. Bhagat et al. report a loss of 1.50%. For the combined loss of wealth, we find a mean of 2.5-2.6%, although smaller median values. Bhagat et al. (1994) report a mean loss of 3.13% and Lerner (1995) reports a mean loss of 2.0%. All three results are

[13] These reasons include: 1.) with a long window, the composition of the market index may change with the addition of new entrants or from rebalancing, 2.) compounding of returns leads to a highly skewed distribution, 3.) not all firms survive to the end of a long event window, and, 4.) the market model or its variance may change or may be sensitive to specification errors over long windows. We find that our measured returns are not highly skewed and there are few cases of firms failing to survive the event window.

[14] Bhagat et al. (1998) includes the data from the Bhagat et al. 1994 paper, so we do not list that separately. Lerner searched the Wall Street Journal as well as news wire services for announcements. The other studies just used articles in the Wall Street Journal.

[15] For this reason, this table does not report standard errors or significance tests.

broadly similar and quite substantial. Lerner reports a mean absolute loss of shareholder wealth of $67.9 million, a median loss of $20.0 million. In general, there does not appear to be a major difference between the results reported in the event window around the Wall Street Journal publication date and the longer window around the filing date.

As noted above, estimates for this sub-sample may be unrepresentative of most patent litigation, however, because most lawsuits are not reported in the Wall Street Journal. Table 3 reports cumulative abnormal returns for all lawsuits in the matched sample (top) as well as those for the large sample (bottom). The base result for the matched sample uses a 25 day event window (T-1 to T+24) and excludes possibly confounding events. The table also reports CARs for suits that were positively identified as infringement suits (that is, the plaintiff was the patent assignee), and for a sample that included lawsuits with possibly confounding news events. The reported means and standard errors use weights based on the variance of the dummy variable coefficient in the event regression. Several results stand out.

First, the estimated percentage losses for alleged infringers are substantially less than those for lawsuits reported in the Wall Street Journal in Table 2. We cannot tell, however, whether the percentage loss estimates in the Journal are larger because of a selection effect or because of the greater information conveyed by publication in the Journal. Even though some learning takes place, we suspect that in most lawsuits, investors remain relatively uninformed compared to those cases where an announcement is published in the Wall Street Journal. The SEC requires reporting of major lawsuits in quarterly and annual filings, but lawsuits will be reported separately only if they materially affect the profits of the firm. For a handful of suits, we checked published sources and typically found no mention of the suit. For this reason, estimates for the non-Journal sample should be interpreted as lower bound estimates of defendant firms' loss of wealth—significant numbers of investors likely became informed about the suit either after our event window or, if there were pre-filing interactions, before.

Second, patentee litigants/plaintiffs appear to suffer some losses as well. These losses are smaller than those for alleged infringers/defendants, but they are statistically significant.[16] This is

[16] It might seem puzzling that the average market response when a patent holder files a lawsuit is negative. Individual rationality implies the patent holder only files lawsuits that have positive expected value. If this is the only relevant information, then plaintiff CARs should be positive. As we explain in more detail in section 4.2, the event of filing may reveal information to investors about more than just the lawsuit. Filing might reveal private information that the patent holder's patent is stronger than investors believed, or that the patent holder has better technology or better entry prospects than investors believed. These possibilities provide additional reasons that the patent holder's share value should rise with the filing of a lawsuit. In contrast, filing might reveal private information of patent weakness, or that a tacit industry agreement not to file patent lawsuits has broken down. These possibilities suggest share value should fall upon lawsuit filing. Thus, a negative CAR might be explained as follows: When a pharmaceutical firm files a patent suit investors perceive the suit has positive expected value, but they also perceive that a key patent was not as strong as they thought and did not deter entry by a potential competitor. Alternatively, when a semiconductor firm files a patent

consistent with previous research and it indicates that lawsuits do not represent simple transfers of wealth on average. Instead, there is dissipation of wealth to consumers, to rivals or to deadweight loss.

Finally, the magnitudes of returns for definite infringement suits are generally larger than for those of all suits and they show a higher level of statistical significance. This may be because among those cases where we could not match the patent to one of the parties, some plaintiffs are mistakenly classified as defendants and vice versa. Or it could be because declaratory actions may be more likely when the stakes at issue are smaller.

The bottom of Table 3 reports results for our large sample. The CARs for alleged infringers are similar to those obtained from the smaller sample—a loss of 0.5% to 0.6%—but here they have statistical significance at the 1% level, except for those lawsuits involving multiple defendants.

When multiple defendants are involved the returns are negligible, suggesting that something is fundamentally different about these estimates. There are several possible explanations for this. It may be that suits naming multiple defendants are more frivolous, so that investors do not expect serious losses. Alternatively, some defendants may have been contractually indemnified, diluting the estimates. A higher percentage of defendants in lawsuits with multiple defendants come from retail and wholesale industries, suggesting that these suits more frequently involve downstream resellers who have less at stake. Costs may be shared among multiple defendants, reducing the individual firm costs.

The estimates in the lower portion of the table do not control for possibly confounding events. However, we find that excluding observations with possibly confounding events does not seem to substantially alter the mean estimated CARs in the top portion of the table (the matched parties sample). To check this further, we repeated the estimates for the large sample of all alleged infringers, but we terminated the pre-event window 30 days prior to the filing of the lawsuit. This made little difference in our estimates, suggesting that confounding events may add noise, but do not bias our estimates.[17]

Figure 2 shows histograms for the cumulative abnormal returns for all lawsuits from the matched sample. The curve for alleged infringers/defendants clearly falls to the left of the curve for patentee litigants/plaintiffs, but both curves are quite diffuse. The distributions are

suit investors perceive the suit has positive expected value, but they also perceive the patent holder plans to exit the industry or has become less forward-looking for some reason, and therefore, the firm is willing to deviate from a no-lawsuit equilibrium. Further research is required to resolve this puzzle.

[17] For example, the estimate for single defendants was 0.608% (0.176%) for the full 200 day pre-event window and 0.609% (0.178%) for the truncated window.

significantly leptokurtic (kurtosis of 7.2 and 9.7 for plaintiffs and defendants, respectively), meaning that they have long tails. This suggests that outliers may be influential. To make sure that our results are not driven by outliers, we also conducted non-parametric tests (the binomial probability test and the Wilcoxon signed rank test) on the large sample and several sub-samples. All of these tests rejected the null hypothesis of a CAR of zero at either the 5% or the 1% level of statistical significance. In addition, the close correspondence between the means and the medians suggests that our mean estimates for alleged infringers are representative.

3.2 Factors affecting Abnormal Returns

Tables 4 and 5 explore factors that might influence the magnitude of investors' reactions to lawsuit filings by comparing means of different sub-groups. We test differences in the means of different sub-groups using one-tailed *t*-tests, allowing unequal variances between the sub-groups and calculating the degrees of freedom using Satterthwaite's approximation (1946). We conduct these comparisons both for the subject firm's characteristics as well as characteristics of its opposing party in the lawsuit. We also ran regressions with various combinations of the variables in Table 4 (or continuous equivalents) on the right hand side. However, given the noisiness of our data, little conclusive could be drawn from these regressions and where significant results were found, they matched the results found with simple *t*-test comparisons of means.

For patentee litigants, we find that firms with high liabilities relative to assets (and to a lesser extent, firms with high current liabilities to current assets) have much more negative returns from initiating lawsuits. One explanation is provided by Haslem (2005), who observes that lawsuit settlements, including patent settlements, are associated with a decline in firm value, on average. Following Jensen and Meckling (1976), he argues that poorly governed firms will tend to settle lawsuits too soon (from the perspective of shareholders) because that allows managers to expend less effort. Firms with low debt have more leeway for managerial discretion. He finds that these firms experience greater declines in value on settlement. By similar logic, firms with low debt may have more discretion about which lawsuits to file. Therefore, they may choose to file just the most profitable lawsuits while managers in more debt-laden companies may be driven to file more marginal lawsuits, leading to relatively lower CARs.

Another explanation might arise if some industries have a "mutual forbearance" repeated game type equilibrium—firms mutually avoid suing each other because they recognize that if

they initiate a suit, they may be punished in the future with retaliatory suits. However, a failing firm may have limited future prospects, hence little to fear from future retaliation. So failing firms, which have high liabilities, may be more likely to initiate suits, including less profitable suits.

For alleged infringers, we find five statistically significant differences. First, if the parties to the lawsuit are in different industries, then the alleged infringer suffers a substantially larger loss, which is statistically significant at the 1% level. Suits from outside the industry may be more of a surprise to investors and may be more indicative of inadvertent infringement. Alternatively, when disputes occur within a narrow industry, the parties may have greater latitude to craft a settlement that benefits both jointly, including, perhaps, collusive settlements.

Second, if the patentee litigant is a newly public firm, the alleged infringer makes out better. This might be because newly public firms are less able to pursue sustained litigation, posing less of a threat to the alleged infringer. Or, perhaps, a suit by an entrant firm provides a signal that the technology may be more profitable than investors previously realized (see the discussion of signaling below).

The remaining three differences from the large sample, shown in Table 5, are statistically significant at the 5% level. First, small firms seem to have substantially more negative returns. This result appears robust to alternative cutoff points below 500 employees, but we found no significant variation in returns among firms larger than that. One explanation for this is that legal costs are relatively higher for small firms, creating a “floor” on the costs of litigation. Second, we find limited evidence that R&D intense firms suffer more negative returns, however, this result seems sensitive to the specific cutoff used. Finally, we also find some evidence that returns were worse during the 1990s compared to the 1980s. Note that the lower returns for alleged infringers do not appear to be matched by greater returns to patentee litigants (top of the table). In other words, this evidence of greater losses does not suggest a greater transfer of wealth to patent holders.

4 The Costs of Patent Litigation

4.1 Legal Costs

We first look at attorneys' fees in patent litigation using supplemental data we collected from legal records. We then estimate the total costs of litigation to alleged infringers based on our event study estimates.

Public documents in certain U.S. patent lawsuits record attorneys' fees because American patent law gives judges the discretion to shift fees in exceptional cases. Patentees usually get fee awards based on a finding of willful infringement, and alleged infringers usually get fee awards based on a finding the patent suit was frivolous or vexatious. We searched Westlaw for all patent cases from 1985-2004 that discussed fee-shifting. We found 352 cases in which one of the parties requested fees (about 100 patent cases go to trial per year). The request was granted in 137 (or 38.9%) of these cases. From this set of 137 cases we were able to determine the magnitude of the fees in 87 cases (63.5% of awards) from judicial opinions or from documents filed by the parties available through the PACER system.

Table 6 shows the median and mean amounts of the fee awards in millions of year 1992 dollars. Mean fees for cases that went through trial ranged were $1.04 million for patentee litigants and $2.46 million for alleged infringers. For cases that were decided prior to trial, the mean fees were $0.95 million for patentee litigants and $0.57 million for alleged infringers.[18] Median values tend to be smaller because the distribution is skewed. In the most extreme case, a $26 million fee was awarded to Bristol-Myers Squibb in conjunction with a successful defense against a pharmaceutical patent suit brought by Rhone-Poulenc. The next largest award was about $7 million.

Our fee-shifting data is in line with survey information collected by the American Intellectual Property Law Association (AIPLA). AIPLA asked patent litigators to estimate the fees associated with patent lawsuits under six different scenarios. Specifically, the survey question divided cases into three different intervals based on stakes, and asked for estimates for cases that concluded at the end of discovery, and cases that reached trial. Their 2001 report indicates the estimated cost through trial was $499,000 when the stakes are less than $1 million, $1.499 million when the stakes are between $1 million and $25 million, and $2.992 million when

[18] We included cases that ended in summary judgments, one case that settled, one case that was a default judgment, and one case that ended in a motion to dismiss.

the stakes are over $25 million.[19] The estimated cost through discovery was $250,000 when the stakes are less than $1 million, $797,000 when the stakes are between $1 million and $25 million, and $1.508 million when the stakes are over $25 million.[20]

The expected legal cost associated with the filing of a patent lawsuit depends on the frequency of the different ways a lawsuit may be terminated. Kesan and Ball (2005) analyze patent lawsuit termination data available from the Administrative Office of the federal judiciary. Examining 5,207 lawsuits that were filed in 1995, 1997, and 2000, they find that most cases terminate short of trial, summary judgment, or other substantive court rulings. In particular, 4.6% of lawsuits reached trial, 8.5% of lawsuits terminated with a summary judgment, dismissal with prejudice, or confirmation of an arbitration decision, and the remaining 86.9% of cases terminated earlier in the process.

Kesan and Ball construct two proxies for legal fees in patent lawsuits: number of days until the suit terminates, and number of documents filed. Their data show that suits that go to trial last about 1.5 times as many days as suits that end with a summary judgment, and suits that end with a summary judgment last about 1.5 times as many days as all other suits. Further, their data shows that suits that go to trial generate about 2.5 times as many documents as suits that end with a summary judgment, and suits that end with a summary judgment generate about 2.5 times as many documents as all other suits.[21] If we assume that the expected legal cost in a suit that ends before summary judgment is one-half of the cost of suit that reaches summary judgment, then using our data in Table 6 we have estimates of $410,000 for the alleged infringer, and $624,000 for the patentee. A similar calculation using AIPLA data for stakes between $1 million and $25 million yields an estimate of $483,000.

4.2 Firm value and patent lawsuits

Using our CAR estimates, we can calculate the loss of wealth that occurs upon a lawsuit filing. From this, we can then infer a cost to alleged infringers. Multiplying the estimated CAR for each firm by the value of its outstanding shares of common stock immediately prior to the lawsuit filing, we obtain a mean loss of wealth in 1992 dollars of $83.7 million. This is an

[19] These amounts increased substantially in the 2003 and 2005 AIPLA reports.

[20] The AIPLA estimate of costs through discovery should be larger than the fees shifted at the summary judgment stage to the extent that discovery continues after summary judgment.

[21] We derive these ratios from their Tables 10-12.

unbiased estimate of the mean loss of wealth, however, it is not the most efficient estimate. We can do better by multiplying the *mean* CAR by each firm's capitalization.[22]

Using means for three categories (suits with multiple defendants, those with single defendants with more than 500 employees and those with single defendants with 500 or fewer employees), we obtain a mean estimated loss of $52.4 million in 1992 dollars and a median loss of $4.5 million.[23] These estimates are somewhat smaller than Lerner's estimate for biotech companies of a mean loss of $67.9 million and a median loss of $20.0 million.

This loss of wealth corresponds to the associated drop in investors' expected profits. But does this loss of wealth correspond to the cost of litigation? There are two reasons why it might not. First, the filing of a lawsuit might reveal information that causes investors to revalue the firm for reasons other than the direct and indirect costs of litigation. We explore these possibilities in this section. In the next section, we consider how much investment the firm must undertake in order to restore its investors' wealth—this might not equal the loss of wealth itself.

News of a lawsuit causes investors to re-evaluate their expectations of the discounted profit flow expected from the defendant firm for several different reasons. We assume that the Efficient Market Hypothesis holds, implying that investors incorporate all publicly available information into their valuation of the firm. Consider defendant firm i at time $t = 0$, before the lawsuit filing, and at $t = 1$, immediately after the news of the filing has been made public. At $t = 0$, investors' expected value of the firm based on publicly available information, V, is

$$V_i(0) = \pi_i(0) - p_i(0)\, C \tag{3}$$

where π represents the discounted expected profits of the firm (excluding litigation), p is the expected number of times the firm will be sued for patent infringement and C is the total expected cost to the firm of a patent lawsuit. This expected cost of litigation includes:

- Legal costs.

---

22 The first estimator is $\frac{1}{N}\sum_{i=1}^{N}(r+e_i)\,x_i$ where N is the number of firms, r is the true CAR, e is the error in measuring the ith firm's CAR, and x is the ith firm's market capitalization. The second estimator is $\frac{1}{N}\left(r+\frac{\sum_{i=1}^{N} e_i}{N}\right)\sum_{i=1}^{N} x_i$. It is straightforward to show that both are unbiased but that the latter has smaller variance assuming that e and x are uncorrelated.

23 Specifically, we multiply the common stock capitalization by .00012 for firms in cases with multiple defendants, by .00564 for single defendants with more than 500 employees, and by .0208 for small single defendants.

- Indirect costs, such as management distraction, loss of market share during the lawsuit, and loss of lead-time advantage.
- Financial costs arising from greater risk, including risk of bankruptcy. These include possibly higher costs of funds and also the loss of wealth associated with a higher risk-adjusted discount rate applied to the stream of future expected profits.[24]
- Costs of expected outcomes including those associated with a settlement agreement and trial outcome. Investors take expectations over all possible outcomes and also over the length of time and cost incurred before outcomes are reached.

Then at time $t = 1$,

$$V_i(1) = \pi_i(1) - p_i(1)\,C - C \tag{4}$$

Comparing (2) and (1) and taking expectations over all lawsuits, the mean CAR should equal

$$E[\Delta \mathrm{V}] = E[\Delta\pi] - E[\Delta \mathrm{p}]\,C - C \tag{5}$$

The first term represents the change in investors' expectations about the future profit stream based on new information made public by the lawsuit filing. The second term in (5) represents investors' re-assessment of the risk of future litigation. This occurs if the lawsuit provides information that the firm is somehow more prone to litigation than originally expected. Clearly, if the sum of the first two terms is non-zero, then the change in firm value provides a biased estimate of the cost of litigation.

There are two sources of information from the filing that might affect these two terms:

1. Information revealed by the filing documents themselves (and any associated press releases, etc.), and,
2. any information revealed by the event as a signal of the patentee's beliefs. For example, because litigation is costly, the lawsuit may signal that the patent holder believes that the opportunity at stake is particularly valuable—otherwise the suit might not be worth the cost. Note that the documents may reinforce this signal, e.g., the claim for damages may also be large, but with a signal the claim may become credible.

[24] Implying that π includes the discounted profit stream evaluated at the original discount rate. This interpretation is consistent with our definition of the cost of litigation being the level of investment necessary to restore the wealth of the firms' investors to the level just prior to the lawsuit.

In order for either source to cause investors to revalue the firm, the lawsuit filing must somehow reveal information that was not previously public knowledge—under the Efficient Market Hypothesis we assume that investors correctly incorporate all public knowledge. In other words, the patent holder or the defendant firm must have some *private* knowledge that is revealed in the filing documents or by the signal generated by the filing.

Therefore, if the first two terms in (5) are to affect the mean CAR substantially, there must be a *systematic* reason for the patent holder or the alleged infringer to have private information that is revealed by the lawsuit filing. The documents in the lawsuit filing typically reveal relatively little hard information other than the fact of the filing, often exaggerated claims of damages, and possible allegations of bad behavior by the defendant (something we discuss below). The patents themselves, of course, are necessarily public information before the suit is filed. But we can identify three reasons why the parties might have private information that is revealed by the filing:

1. Private information about the quality of the technology. For well-known reasons, managers have private information about the quality of their technology. A lawsuit may signal that the patent holder knows that the defendant's technology is of better quality than investors previously realized, hence the market potential is greater, hence a lawsuit may be more profitable. Note that in this case, $E[\Delta\pi] > 0$.
2. Private information about entry plans. If a patent holder plans on entering the defendant firm's market, then the lawsuit might reveal this knowledge, causing investors to revalue the defendant firm downwards because they expect greater competition for the firm. Note that in order for this factor to substantially affect our average CARs, such prospective entrants must initiate a substantial number of patent lawsuits. Also, the prospective entrants cannot have revealed any information about their entry plans prior to filing the lawsuit. This strikes us as a rather odd business strategy—one would think a superior strategy would be to enter the market *before* filing a lawsuit so as to capture market share from those customers who want to avoid the defendant firm. Nevertheless, we will look at empirical evidence regarding this story below. In this case, $E[\Delta\pi] < 0$.
3. Private information about managerial quality or level of effort. For well-known reasons, managers keep private information about their abilities and about the level of effort that they exert. Lawsuits might tend to indicate that managers at the defendant firm were not sufficiently diligent in clearing patent rights or, worse, that they copied

technology rather than developing their own. If this tends to be true and if managers tend not to correct their behavior following a lawsuit, then investors might revalue future profits downwards. This occurs both because investors might expect more patent litigation in the future (the second term of (5)) and because poor managerial quality might also reduce profits generally (the first term in (5)).

Several empirical observations lead us to discount the second and third explanations, however. If lawsuit filings revealed news about previously unknown entrants, we might expect this to be particularly true for plaintiffs that had recently gone public. These plaintiffs might not be widely known and therefore, on average, defendant firms might lose greater value when sued by newly public firms. However, we find that defendants' CARs are significantly more positive when the plaintiff is a newly public firm (see Table 4).[25]

In addition, if news about entry is a significant factor affecting average CARs, then we would expect to find that a significant portion of plaintiffs were: a.) not known as market rivals to the defendant firm prior to the lawsuit, but, b.) became market rivals subsequently. Using Compustat's market segment data, we found that this fact pattern is actually rather uncommon. Compustat reports SIC codes for each firm's major market segments. Of the plaintiffs who had no market segments in common with defendants prior to the lawsuit, we found that only 5% entered a market segment in common with the defendant during the three years following the lawsuit filing.[26] Thus it seems unlikely that a substantial part of defendants' CARs can be explained by revelation of previously unknown entrants.

Other evidence leads us to discount the significance of any news about managerial quality or effort revealed by the lawsuit. Managerial quality is less likely to be of significance in lawsuits that are filed the same year that the patent is granted—often these patents contain claims that were not previously publicly known, so there is less that managers could have done to avoid infringement and managerial quality is less of an issue. For this reason, lawsuits on these patents cannot reveal as much about managerial quality. If revelations about managerial quality explain a large portion of the defendants' CARs, we would expect the CARs to be more positive for

[25] The increase could be because startup firms are less able to pursue sustained litigation and therefore a lawsuit from a startup poses less of a threat. Alternatively, a lawsuit by an entrant may indicate that the technological opportunity is greater than investors previously realized.

[26] This figure compares SIC market segments at the 4 digit level. A comparable calculation using 3-digit industry classifications finds a 6% entry rate. This comparison only concerns major market segments, so some entry is unrecorded in minor segments, however, rivalry in minor market segments is only likely to have a minor effect on firm value.

patents issued the same year as the lawsuit. In fact, we find that the CARs are more negative for these patents, although the difference is not statistically significant.

Also, we would expect that the managerial quality explanation is much more significant the first time a firm is sued. That is, if a lawsuit reveals significant information about managerial quality, we would expect the second lawsuit to reveal less, and the fifth or tenth lawsuit to reveal even less. We would expect investors to learn and, for this reason, we would expect that, on average, CARs would reflect less revelation of information about managerial quality for, say, the fourth through tenth lawsuit than for the first three.[27] We compared defendant CARs depending on the number of lawsuits the firm had in our sample or on the sequence of the lawsuit. We found no significant differences between CARs for a wide range of different comparisons. E.g., firms with only one lawsuit in our sample had CARs that were on average only .0008 (standard deviation of .0047) less than the CARs for firms sued multiple times. Thus revelations about managerial quality do not seem to explain much of the average loss in firm value from the filing of a lawsuit.

We have little empirical evidence bearing on the role of revelations about technological quality other than anecdote.[28] In Table 4, we saw that defendants do better when the lawsuit is filed by a newly public firm. One explanation is that suits by newly public firms reveal information about technological quality, but this is not the only possible explanation. However, as we noted above, for revelation about technological quality, $E[\Delta\pi] > 0$. Given this, we conclude that $E[\Delta\pi] \geq 0$ and $E[\Delta p] \approx 0$, so that $C \geq -E[\Delta V]$. That is, the cost of litigation is likely at least as large as the loss in firm market value.

4.3 Investment level costs

If we want to know how much litigation "taxes" investment in innovation, then we need to calculate something other than the loss of wealth. That is, we define the "cost of litigation" as the amount that the firm has to invest in order to increase its value to the level it had just prior to the lawsuit, all else equal. This does not necessarily equal the amount of wealth the firm loses because firms are not necessarily operating at the long-run steady state. They may, instead, be undergoing dynamic adjustment. Then changes in investment will be larger or smaller than the associated changes in firm value. In particular, assuming constant returns to scale, an additional investment of one dollar should increase firm value by an amount equal to Tobin's Q.

27 This assumes, of course, that management is not entirely replaced between lawsuits.

28 A tech industry joke on hearing that someone has been sued is, "Congratulations, you must be doing something right!"

Following this logic, to calculate the cost of litigation, we divide the estimated loss of wealth by Tobin's *Q*.[29] This gives us a mean cost of litigation to alleged infringers of $28.7 million and a median cost of $2.9 million in 1992 dollars.

These estimates are clearly much larger than the estimates of direct legal costs. Most of the cost of litigation to firms appears to arise from expected settlement payments and business costs such as loss of market share, management distraction, and increased financial costs from greater risk. These costs are incurred even if the suit does not proceed to trial, as happens most often.

It is interesting to compare our estimate to data from cases that proceed to trial. For the small number of reported cases that go to trial, are won by the patentee, and which award damages to the patentee, we can compare the magnitude of these damages. Mean reported lawsuit damages from 1991-2005 are $10.7 million in 1992 dollars.[30] This number does not include the business cost of the injunction to the infringer, which is often much larger than the damages. For example, the court found damages of $53.7 million in NTP v. RIM, but because of the injunction, NTP eventually settled for $612 million. This mean also does not include the costs of pursuing the litigation, both direct payment of legal costs and indirect business costs. Nevertheless, it is reassuring that this figure is of the same order of magnitude as our mean estimate.

5 The Risk of Infringement for Public Firms

These cost estimates can be summed over all the observed lawsuit filings to obtain measures of firm risk. Table 7 shows three related measures.

The first column lists the annual cost of litigation obtained by summing the cost over all the events in our large sample in each year of the sample. During 1996-99, this averaged $14.9 billion in 1992 dollars. This number is large compared to estimates of patent value. Bessen (2008), using renewal data to estimate patent value, reports the aggregate value of patents issued to *all* US patentees (not just public firms) in 1991was about $4.4 billion.

Moreover, this figure has varied considerably over time, increasing dramatically from $2.0 billion in 1984 to $16.1 billion in 1999. Figure 3 shows the annual time series. The rise began in the early 1990s and closely follows the increasing frequency of litigation (Bessen and

29 We calculate Tobin's *Q* as the aggregate value of firms divided by the inflation-adjusted value of the aggregate sum of accounting assets and R&D. For details on the computation of these quantities, see Bessen (2006b).

30 This figure is the mean of deflated annual means reported in Pricewaterhouse Coopers (2006).

Meurer 2005). But other factors contributed as well, including the increase in R&D spending and firm capitalization. Below we look at infringement risk normalized by R&D. The absolute cost of litigation was borne almost entirely by large firms and nearly half by firms in the computer, electronics and software industries.

Note that this series may be substantially understated because, as is well-known, the Derwent Litalert data under-report lawsuits (Lanjouw and Schankerman 2004, Bessen and Meurer 2005). In our 2005 working paper using this sample, we find that only about 64% of lawsuits are reported in Derwent. We have left the first column uncorrected, since it reports a simple sum for our sample. However, the second and third columns compare litigation cost to numbers of firms and to R&D spending, respectively, so to make the appropriate comparisons, we correct these for under-reporting by dividing by 0.64.[31]

On the other hand, this series may slightly overstate the aggregate cost of patent litigation *per se* because some of the suits listed involved more than just charges of patent infringement and validity. For example, sometimes patent owners will combine allegations of patent infringement with allegations that other rights (including other intellectual property rights) have been violated. Some of the suits of this sort might occur even if there were no patent infringement at issue, so it might not be appropriate to include all of the costs associated with these suits in an aggregate estimate of patent litigation costs. However, for two reasons we do not think this is a serious problem. First, searching published court decisions between 1991 and 1999, only 11% of patent infringement and validity suits also involved claims involving trade secrets, trademarks, copyright, false advertising, unfair competition or noncompete clauses.[32] Second, in Table 4 we observed that the alleged infringer's losses are much greater for inter-industry suits than for intra-industry suits. Since most of the cases involving these additional legal issues occur between rivals in the same industries, these suits do not contribute much to aggregate litigation costs. So it seems unlikely that our aggregate cost estimates overstate the costs of patent litigation by more than a few percent.

The second column displays the annual firm infringement risk. This is the mean expected cost of litigation for a firm from patent infringement lawsuits (or related declaratory actions). It averaged $4.5 million during 1996-99 and it shows a similar pattern of distribution.

31 Lanjouw and Schankerman (2004) found no significant differences between the characteristics of the reported and unreported lawsuits.

32 Based on a search of case synopses in the Westlaw FIP-CS database.

The third column shows the ratio of annual litigation cost to annual aggregate R&D. This averaged 14.0% during 1996-99. This relative rate also increased from 1984 to 1999, more than tripling to 19.3% (roughly in line with the growth of the litigation hazard), but this increase was not as rapid as for the quantity in column 1. Note that relative to R&D, litigation risk is low for small firms and for firms outside of the chemical, pharmaceutical and tech industries.

It is tempting to compare this ratio with the "equivalent subsidy rate" for patents, that is, the aggregate value of patents divided by the value of the corresponding R&D. Schankerman (1998) suggests that this ratio represents an upper bound on the subsidy that patents provide to invest in innovation. But, as we argued above, this is clearly a gross subsidy that can be offset by litigation risk if innovators risk inadvertent infringement and by other costs. Several papers calculate this ratio by comparing the value of a nation's patents (estimated using patent renewal data) to R&D (calculated by allocating national R&D spending to the patents obtained in the subject country). Lanjouw et al. (1998) review this literature and report that most subsidy rates are on the order of 10-15%. Arora et al. (2005), use survey data to obtain a comparable estimate of 17%.

However, these numbers are not directly comparable to our estimates of relative litigation risk for at least three reasons. First, because of the way these studies allocate global R&D, they effectively report the subsidy provided by worldwide patents, not patents in a single country.[33] However, the litigation cost is only for US litigation and does not include the costs of litigation in other countries nor the costs of other dispute resolution such as opposition proceedings. An "apples-to-apples" comparison would include these costs as well.

Second, the subsidy rate calculations based on patent value use the value of all of the nation's patents, including patents from individual inventors and small firms. The litigation risk estimates are only for public firms; these are the firms that conduct the lion's share of R&D. A more appropriate comparison would calculate subsidy rates using patents values only for public firms (see Bessen 2008 for comparable figures). In any case, public firms may experience both different subsidy rates and different litigation costs than other firms.

Finally, the litigation costs are estimated for the current year, but the value of patents granted reflects a stream of profits in *future* years. Ideally, we would want to compare litigation

33 That is, using trade data, they allocate a share of the R&D performed in every OECD country to, say, French patents when they calculate the subsidy rate using the value of French patents. The apparent assumption behind this allocation is that subsidy rates are the same across nations and that the share of trade is proportional to each nation's share of worldwide patent value. Then the calculated subsidy rate will represent the return from worldwide patents. Similarly, Arora et al. use US patents as a right hand variable, but this proxies for each firm's worldwide patents.

costs to the profits from patents on the same cohort of technologies that were litigated. Some of these profits are realized prior to the time of litigation. Since both litigation costs and patent values are trending up, this use of current patent values understates the significance of litigation costs.

All three of these considerations suggest that a direct comparison of reported subsidy rates to US litigation risk overstates the relative positive value of patents. At the very least, these estimates suggest that litigation risk is quite large compared to the private benefits of patents, especially in recent years.

6 Conclusion

Using a large set of event studies, we estimate the total cost that patent litigation imposes on firms and we estimate the risk of infringement litigation. We find that, contrary to what is sometimes assumed, the business costs of litigation far exceed the direct legal costs. And we find that by the late 1990s, patent litigation risk was of the same order as, if not larger than, estimates of the private benefits firms receive from patents. Moreover, consistent with the previous literature, the losses to alleged infringers do not correspond to a transfer of wealth to patent holders; instead there is a substantial joint loss of wealth. Our estimates concern private costs rather than the social costs of litigation, nevertheless these estimates tell us something about the effectiveness of patents as a policy tool to encourage investment in innovation.

In the best case, this suggests that the patent system is at present an inefficient form of subsidy or regulation. Thomas Hopkins estimates the total 1992 cost of general regulatory compliance is $389,911 per firm (in 1995 dollars).[34] But the costs of complying with the patent system—annual infringement risk of $4.5 million—are much larger.

In the worst case, the net effect of patents today may be to reduce the profits of public firms and to possibly impose disincentives on innovative activity as well. Exploration of the possible causes and the significance of this for policy and for normative analysis are beyond the scope of this paper, however. Nevertheless, our analysis indicates that infringement risk should be an important consideration in the formulation of patent policy.

34 Thomas D. Hopkins The Changing Burden of Regulation, Paperwork, and Tax Compliance on Small Business, http://www.sba.gov/gopher/Legislation-And-Regulations/Burden/burd7.txt

7 References

American Intellectual Property Law Association. 2001. "Report of the Economic Survey."

Arora, Ashish, Marco Ceccagnoli and Wesley Cohen. 2005. "R&D and the Patent Premium," working paper.

Barber, Brad M. and John D. Lyon. 1997. "Detecting Long-run Abnormal Stock Returns: The empirical power and specification of test statistics," *Journal of Financial Economics,* v. 43, pp. 341-72.

Bessen, James (2006b), "Estimates of Firms' Patent Rents from Firm Market Value," Boston Univ. School of Law Working Paper No. 06-14.

Bessen, James (2008 forthcoming), "The Value of U.S. Patents by Owner and Patent Characteristics," *Research Policy*.

Bessen, James and Robert M. Hunt. 2007. "An Empirical Look at Software Patents," *Journal of Economics and Management Strategy* 16, no. 1, pp. 157-89.

Bessen, James and Michael J. Meurer. 2005. "The Patent Litigation Explosion," B.U.S. L. Law and Economics Working Paper Series, No. 05-18.

Bessen, James and Michael J. Meurer (2006), "Patent Litigation with Endogenous Disputes," *American Economic Review*, 96, no. 2, pp. 77-81.

Bessen, James and Michael J. Meurer (2008 forthcoming), *Do Patents Work?* manuscript.

Bhagat, Sanjai, John M. Bizjak, and Jeffrey L. Coles, 1998. "The shareholder wealth implications of corporate lawsuits," *Financial Management* 27, pp. 5–27.

Bhagat, Sanjai, James A. Brickley, & Jeffrey L. Coles. 1994. "The Costs of Inefficient Bargaining and Financial Distress: Evidence from Corporate Lawsuits," *Journal of Financial Economics*, v. 35, p. 221.

Brown, Stephen J. and Jerold B. Warner. 1985. "Using Daily Stock Returns: The case of event studies," *Journal of Financial Economics,* v. 14, pp. 3-31.

Cutler, David and Lawrence Summers. 1988. "The costs of conflict resolution and financial distress: Evidence from the Texaco-Pennzoil Litigation," RAND Journal of Economics, 19, pp. 157-72.

Dodd, Peter and Jerold B. Warner. 1983. "On Corporate Governance: A study of proxy contests," *Journal of Financial Economics*, v. 11, pp. 401-38.

Hall, Bronwyn H., Stuart Graham, Dietmar Harhoff, David C. Mowery, 2004. "Prospects for Improving U.S. Patent Quality via Postgrant Opposition," Innovation policy and the economy. Volume 4, Cambridge and London: MIT Press for the National Bureau of Economic Research, pp. 115-43.

Haslem, Bruce. 2005. "Managerial Opportunism during Corporate Litigation," Journal of Finance, 60, no. 4, pp. 2013-41.

Jensen, Michael C., and William H. Meckling, 1976, "Theory of the firm: Managerial behavior, agency costs and ownership structure," Journal of Financial Economics 3, 305–360.

Kesan, Jay P. and Gwendolyn G. Ball. 2005. "How Are Patent Cases Resolved? An Empirical Examination of the Adjudication and Settlement of Patent Disputes," U. Illinois Law & Economics Research Paper No. LE05-027.

Kothari, S. P. and Jerold B. Warner. 1997. "Measuring Long-horizon security price performance," *Journal of Financial Economics,* v. 43, pp. 301-39.

Kramer, Lisa A. 2001. "Alternative Methods for Robust Analysis in Event Study Applications," in Lee, C. F., ed., *Advances in Investment Analysis and Portfolio Management,* v. 8, pp. 109-32.

Lanjouw, Jean O., Ariel Pakes and Jonathan Putnam. 1998. "How to Count Patents and Value Intellectual Property: The uses of patent renewal and application data," *Journal of Industrial Economics* 46 405-32.

Lanjouw, Jean O. and Mark Schankerman. 2004. "Protecting Intellectual Property Rights: Are Small Firms Handicapped?," *Journal of Law and Economics*, 47, pp. 45-74.

Lerner, Josh. 1995. "Patenting in the Shadow of Competitors," *Journal of Law and Economics*, 38, no. 2, pp. 463-95.

Levin, R. C. and J. Levin. 2002. "Patent Oppositions." working paper.

Lunney, Glynn S. Jr. 2004. "Patent Law, the Federal Circuit, and the Supreme Court: A Quiet Revolution," Supreme Court Economic Review, 11, pg. 1.

Mackinlay, A. Craig. 1997. "Event Studies in Economics and Finance," *Journal of Economic Literature*, 35(1), pp. 13-39.

Meurer, Michael J. 1989. "The Settlement of Patent Litigation," *RAND Journal of Economics*, Spring 1989, v. 20, iss. 1, pp. 77-91.

Moore, Kimberly A. 2000. "Judges, Juries and Patent Cases – An Empirical Peek Inside the Black Box, *Michigan Law Review* 99, pp. 365-406.

Pakes, Ariel and Mark Schankerman. 1984. "The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources," in Griliches, Zvi, ed., *R&D, Patents and Productivity*, chicago: University of Chicago Press for the NBER.

PricewaterhouseCoopers (2006), "2006 Patent and Trademark Damages Study."

Salinger, Michael. 1992. "Standard Errors in Event Studies," *The Journal of Financial and Quantitative Analysis* 27(1), pp. 39-53.

Satterthwaite, F. E. 1946. “An approximate distribution of estimates of variance components,” Biometrics Bulletin, 2. pp. 110-14.

Schankerman, Mark. 1998. "How Valuable is Patent Protection: Estimates by Technology Fields," *RAND Journal of Economics*, 29, no. 1, pp. 77-107.

Ziedonis, Rosemarie Ham. 2004. “Patent Litigation in the US Semiconductor Industry,” in Cohen, Wesley M. and Stephen Merrill, eds., *Patents in the Knowledge-Based Economy*, National Research Council.

Tables and Figures

Table 1. Summary Statistics

	Matched Sample				All Alleged Infringers	
	Patentee Litigants		Alleged Infringers			
	Mean	Median	Mean	Median	Mean	Median
Sales ($ million)	7,020.4	1,267.7	6,186.7	1,022.7	8,604.0	1,368.1
Employees (1000s)	40.2	9.2	36.1	6.7	46.3	9.3
R&D / Sales	9.4%	5.4%	18.9%	5.3%	13.9%	5.0%
No R&D reported	6.1%		9.0%		18.4%	
No. observations	771		720		2887	

Table 2. Cumulative Abnormal Returns from Suits Announced in Wall Street Journal, 1984-99

Event window	WSJ article T-2 to T+1	Suit filing T-1 to T+24	Bhagat et al. (1998)
Patentee Litigant (Plaintiff)			
mean	-0.3%	-0.1%	-0.31%
median	0.0%	0.9%	
no. of observations	86	86	
Alleged Infringer (Defendant)			
mean	-2.6%	-1.8%	-1.50%
median	-1.4%	-1.9%	
no. of observations	82	82	
Combined (matched parties)			
mean	-2.6%	-2.5%	
median	-1.8%	-0.5%	
no. of observations	80	80	
Addendum: mean combined abnormal returns			
Bhagat et al. (1994)	-3.13%		
Lerner (1995)	-2.0%		

Note: Events with possibly confounding news are excluded. Average cumulative abnormal returns are simple unweighted means.

Table 3. Cumulative Abnormal Returns

	Mean CAR		Median CAR	Robust *Z* Statistic	Observations
Sample: Matched Parties					
Patentee Litigants					
Base	-0.38%	(0.30%)	0.00%	-1.51	667
Definite infringement suits	-0.63%	(0.37%)*	-0.45%	-2.18*	412
Alleged Infringers					
Base	-0.62%	(0.33%)*	-0.97%	-1.55	661
Definite infringement suits	-0.77%	(0.42%)*	-0.83%	-1.70*	407
With possibly confounding events	-0.45%	(0.31%)	-0.57%	-1.32	743
Sample: All alleged infringers					
Base	-0.50%	(0.16%)**	-0.51%	-3.24**	2,887
Single defendants	-0.61%	(0.18%)**	-0.54%	-2.94**	2,460
Multiple defendants	-0.01%	(0.39%)	-0.39%	-1.38	427
Single defendants, definite infringement cases	-0.63%	(0.27%)**	-0.42%	-2.37**	1,108

Note: Standard errors in parentheses. Single asterisk indicates statistical significance at the 5% level; double asterisk indicates 1% significance. Average cumulative abnormal returns (CARs) are weighted means, with weights proportional to the inverse of the estimated variance of each return. In matched sample, events with possibly confounding news are excluded, except where noted. Event window is 25 days (T-1 to T+24). Cumulative abnormal returns are estimated using OLS except for cases with multiple defendants (in large sample), which are estimated jointly. The robust *Z* statistic is a joint test of the individual firm *t* statistics (Kramer 2001).

Table 4. Differences in Mean CARs by Characteristics

Sample: Matched Parties

	Alleged Infringer	Patentee Litigant
Firm characteristic		
Employees < 500	-3.20% (2.32%)	-3.18% (2.45%)
R&D / Sales > .15	0.22% (2.16%)	-0.53% (1.22%)
Total liabilities / Total Assets > .5	1.40% (0.87%)	-2.35% (0.75%)**
Capital / Employee > $100,000	-0.02% (0.93%)	-1.02% (0.74%)
Current Assets / current liabilities < 1.5	0.94% (1.00%)	-1.91% (0.87%)*
Newly public firm	-0.94% (1.78%)	-1.92% (2.56%)
Rival characteristic		
Employees < 500	1.06% (1.19%)	-1.37% (1.07%)
R&D / Sales > .15	0.23% (1.62%)	0.81% (0.97%)
Total liabilities / Total Assets > .5	-0.15% (0.86%)	-0.35% (0.80%)
Capital / Employee > $100,000	-0.99% (0.95%)	1.02% (0.74%)
Current Assets / current liabilities < 1.5	1.69% (1.11%)	1.19% (0.86%)
Newly public firm	3.77% (1.51%)**	0.32% (1.05%)
Other Characteristics		
Year > 1989	-0.15% (0.82%)	0.09% (0.77)%
Firms in same SIC4 primary industry	2.67% (1.16%)**	-0.11% (0.78%)

Note: Standard errors in parentheses. Single asterisk indicates difference is statistically significant at the 5% level; double asterisk indicates 1% significance (one-tailed test allowing unequal variances and using Satterthwaite's calculation for degrees of freedom). Average cumulative abnormal returns are weighted means, with weights proportional to the inverse of the estimated variance of each return. Comparisons are for cases where infringement is known and no possibly confounding events have been found.

Table 5. Differences in Mean CARs by Firm Characteristics
Sample: All alleged infringers

Employees < 500	-1.70% (0.92%)*
R&D / Sales > .15	-1.79% (0.80%)*
Total liabilities / Total Assets > .5	0.05% (0.33%)
Capital / Employee > $100,000	-0.26% (0.44%)
Current Assets / current liabilities < 1.5	0.11% (0.34%)
Year > 1989	-0.56% (0.32%)*
Patentee is public firm	-0.12% (0.35%)
Industry	
SIC = 28 (chemicals, inc. pharma)	-0.41% (0.41%)
SIC = 35,36,73 (electronics, computer,sw)	0.06% (0.38%)
Other manufacturing	0.16% (0.33%)

Note: Standard errors in parentheses. Single asterisk indicates difference is statistically significant at the 5% level; double asterisk indicates 1% significance (one-tailed test allowing unequal variances and using Satterthwaite's calculation for degrees of freedom). Average cumulative abnormal returns are weighted means, with weights proportional to the inverse of the estimated variance of each return.

Table 6. Attorneys' Fees Awarded in Patent Lawsuits (in millions of year 1992 dollars)

	Mean	Median	Observations
Patentee Litigant			
Summary Judgment	.95	.40	8
Verdict	1.04	.78	51
Alleged Infringer			
Summary Judgment	.57	.30	10
Verdict	2.46	.98	18

Table 7. Measures of Infringement Risk, Public Firms

	Aggregate Annual Cost of Litigation to Alleged Infringers (billion $92)	Annual Firm Infringement Risk (million $92)	Aggregate Risk / R&D
1984	2.0	1.3	4.9%
1999	16.1	7.0	19.3%
1996 - 99			
All firms	14.9	4.5	14.0%
Small firms (employees <500)	0.1	0.1	1.3%
Large firms (employees>=500)	14.8	9.8	14.9%
SIC = 28 (chemicals, inc. pharma)	3.4	9.7	14.1%
SIC = 35,36,73 (electronics, computer, software)	6.8	5.7	14.8%
Other manufacturing	1.7	2.3	5.3%

Note: Annual cost of litigation is the mean CAR times the market capitalization of each firm's common stock divided by a GDP deflator and by the aggregate Tobin's *Q* (market value divided by replacement value of capital including R&D). Firm infringement risk is the expected annual cost of litigation. Column 1 includes all events in the large sample (2,887) with separate means for small firms and lawsuits with multiple defendants. Columns 2 and 3 have been adjusted for under-reporting of lawsuits (see Lanjouw and Schankerman 2004 and Bessen and Meurer 2005).

Figure 1. Frequency of Wall Street Journal Stories Relative to Court Filing Date

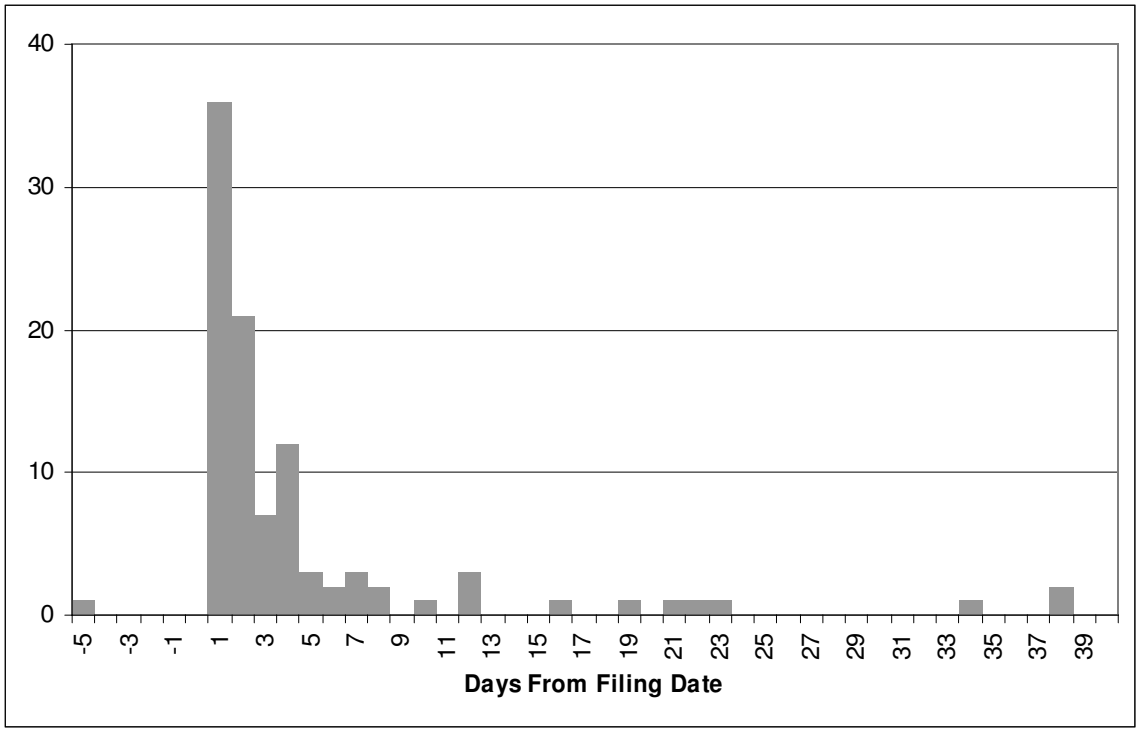

Figure 2. Histograms of Cumulative Abnormal Returns

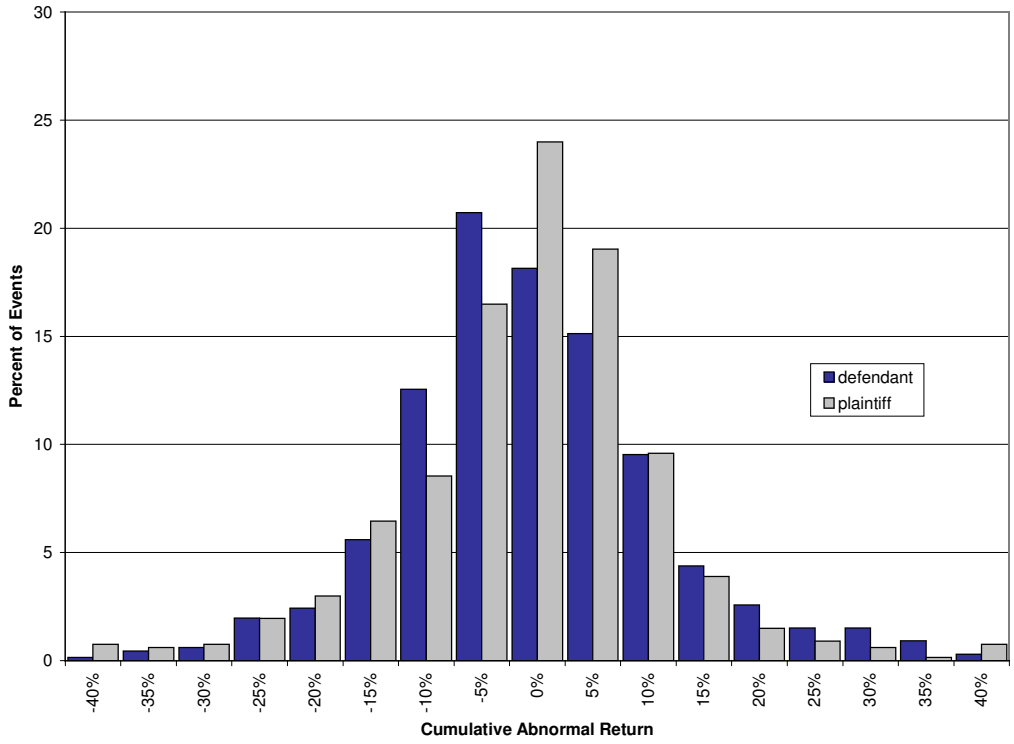

Figure 3. Aggregate annual cost of patent litigation to alleged infringers

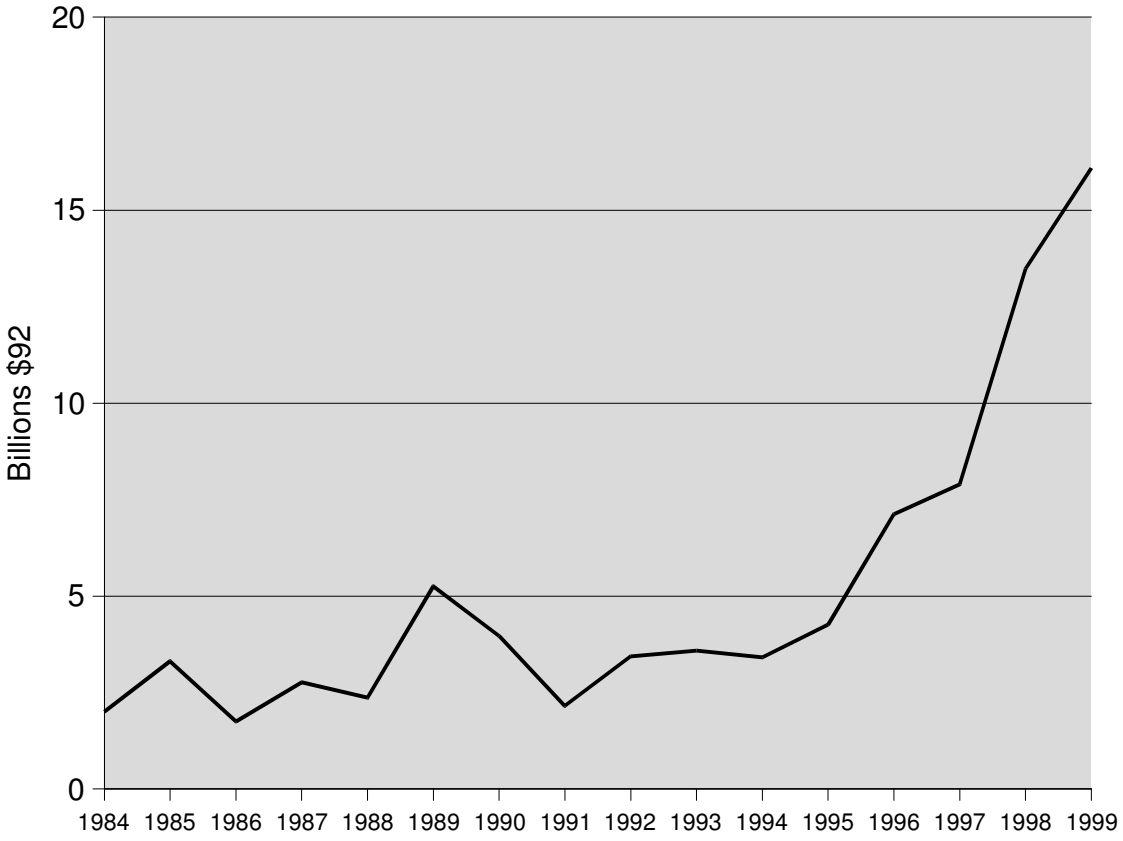

8 Appendix

This appendix explores further our choice of a window around the lawsuit filing date rather than an announcement in a newspaper or wire service. First, we explore whether a sample based on Wall Street Journal articles is likely to suffer sample selection bias. Table A1 shows Probit regressions on whether a lawsuit in our matched sample received mention in the Wall Street Journal. The patentee litigant's capital intensity and the alleged infringer's stock beta are both highly significant (at the 1% level) predictors of a Wall Street Journal article. Because high beta stocks are likely to have a larger reaction to news of a lawsuit, this suggests that samples based on Wall Street Journal articles may have significant bias. We find, in fact, that our estimates from our sub-sample of lawsuits announced in the Wall Street Journal do have much more negative CARs.

Because we do not have information on whether a suit is an infringement suit or a declaratory action in all cases, we likely mis-identify some plaintiffs and defendants, possibly diluting our estimates for alleged infringers. One way to correct this would be to limit our sample to cases of definite infringement, but this might also introduce a selection bias. The last two columns of Table A1 explore characteristics that may affect whether the suit is an infringement suit or a declaratory action. It appears that newly public patentees may be a bit more aggressive in filing suits, while larger alleged infringers may be more likely to end up in an infringement suit. Large firms may avoid filing declaratory actions, waiting for evidence that the patent owner has the resources to conduct a lawsuit. Because there may be a selection bias, we report CARs both for the entire sample and also for cases that we know are infringement suits.

Finally, as discussed in the text, because news of a lawsuit filing leaks out more slowly than a newspaper announcement, we use a 25 day event window. Figure A1 shows the mean CARs we would obtain using shorter event windows. Note that the unweighted mean and the median CARs both react more sharply in the days after the filing. This is because high beta stocks respond more quickly after the filing (they are the ones where investors may have the greater incentive to obtain such news). Because the CARs for low beta stocks are estimated more precisely and their response is slower, the weighted mean responds more slowly. However, all three averages are roughly equal by the end of our 25 day window.

Table A1. Suit Announcement and Type

	Wall Street Journal Article		Infringement Suit	
	1	2	3	4
Plaintiff/patentee litigant				
Ln employment		0.05 (.03)	.02 (.03)	.01 (.04)
New firm		-.25 (.23)	**.63** (.29)	.62 (.33)
Stock Beta	.13 (.12)	.15 (.11)		.20 (.13)
Capital / employee	**1.01** (.38)	**1.12** (.40)		-.64 (.49)
Defendant/alleged infringer				
Ln employment		-.01 (.03)	.06 (.03)	**.07** (.03)
New firm		.28 (.20)	-.01 (.20)	-.05 (.22)
Stock Beta	**.35** (.13)	**.35** (.13)		.05 (.14)
Capital / employee	.05 (.36)	.11 (.36)		-.95 (.51)
No. of observations	637	637	507	475
Pseudo-R-squared	.049	.062	.023	.057

Note: Probit regressions. Robust standard errors in parentheses. Bold estimates are significant a the 5% level or better. Regressions include industry dummies (not shown).

Figure A1. Average Abnormal Cumulative Returns Over Time