ONLINE APPENDIX
Patents and Cumulative Innovation: Causal Evidence from the Courts

Alberto Galasso
University of Toronto and NBER

Mark Schankerman
London School of Economics and CEPR

1 Microfounding the JIP Measure

We develop a simple model of strategic voting, closely following Feddersen and Pesendorfer (1996). There are three judges $i \in \{1, 2, 3\}$ who must decide whether a patent is valid (V) or not invalid (N). Judges are uncertain about the validity of the patent and each judge gets a signal v or n that is correlated with the true state. Specifically we assume that

$$\Pr(v|V) = \Pr(n|N) = p_i.$$

The parameter $p_i \in [\underline{p}, \overline{p}]$ with $.5 < \underline{p} < \overline{p} < 1$ is the probability that a judge receives the correct signal. The parameter p_i can be interpreted as the 'complexity' of the case for judge i. The assumption that the signals are private information is standard in the literature on voting. Feddersen and Pesendorfer (1996) provide a number of reasons why the complete disclosure of private information may not occur. For example, some judges may have technical knowledge that is relevant for the case but difficult to communicate. Moreover, differences in preferences for patent validity may reduce the incentives to reveal private information in deliberations.

The judges vote simultaneously either to validate or invalidate and the decision is taken by majority voting. There are two outcomes: either the patent is invalidated (1) or not (0). We assume that each judge maximizes her expected utility and that preferences are given by $u(1, N) = u(0, V) = 0$ and $u(1, V) = -q_i$ and $u(0, N) = -(1 - q_i)$. The parameter q_i characterizes the judge's threshold of reasonable doubt. Let $\beta_i(n)$ denote the posterior probability for judge i that the patent is invalid, conditional on obtaining an invalidity signal and being pivotal, i.e that the other two judges, x and z, receive different signals from each other. Let $\beta_i(v)$ denote the posterior probability for judge i that the patent is invalid, conditional on obtaining a validity signal and being pivotal:

$$\begin{aligned} \beta_i(n) &= \frac{p_i(1-p_x)p_z}{p_i(1-p_x)p_z + (1-p_i)(1-p_x)p_z} = p_i \\ \beta_i(v) &= \frac{(1-p_i)(1-p_x)p_z}{p_i(1-p)p_z + (1-p_i)(1-p_x)p_z} = 1 - p_i. \end{aligned}$$

We assume that $\beta_i(v) < q_i < \beta_i(n)$ for each i. Feddersen and Pesendorfer (1996) show that if this assumption is satisfied each judge in equilibrium will vote according to his signal (i.e., what they call 'informative' voting). More specifically, a pivotal judge receiving an invalidity signal will vote for invalidity as long as her expected utility is higher from doing so:

$$\beta_i(n)0 - (1 - \beta_i(n))q_i \geq (1 - \beta_i(n))0 - \beta_i(n)(1 - q_i)$$

which is satisfied because we assumed $q_i < \beta_i(n)$. She will also vote for validity if she receives a validity signal because $\beta_i(v) < q_i$. Moreover, note that $\beta_i(v) = 1 - p_i$ and $\beta_i(n) = p_i$, so the condition for an informative equilibrium is always satisfied as long as $1 - \underline{p} < q_i < \underline{p}$.

We assume that the complexity of a case, p_i, is an *i.i.d.* draw from a distribution $F(p)$ with support $[\underline{p}, \overline{p}]$ and that $1 - \underline{p} < q_i < \underline{p}$. The ex ante probability that judge i will vote for invalidity will be $1 - F(q_i) \equiv f^i$ and the expected number of invalidity votes in the three judge panel will be equal to

$$JIP = f^1 f^2 f^3 + f^1 f^2 (1 - f^3) + f^1 (1 - f^2) f^3 + (1 - f^1) f^2 f^3.$$

Given the random allocation of judges to cases, the sample average of a judge's validity votes will be an unbiased estimator of her probability of voting for validity. Moreover, *JIP* is a consistent estimator of the number of validity votes in the three judge panel (it is not unbiased, as it is a nonlinear transformation of the f^i's).

2 Randomization and Heterogeneity of Judges: Robustness

As additional evidence on the randomization of judges to cases, Table A2 reports the correlations between *JIP* and various subsets of patent characteristics. In each regression we include a control for the year of the Federal Circuit decision. The table confirms that there is no significant correlation between *JIP* and the patent variables used in our analysis. Moreover, the point estimates of the coefficients are close to zero, indicating that the magnitude of the correlation is very small. Columns (1) to (4) present separate regressions for claims, citations, self-citations and age. In column (5) we present the correlation between the technology dummies and *JIP*. We cannot reject the hypothesis that the coefficients on the technology field dummies are jointly equal to zero (*p-value*=0.41). In column (6) we regress JIP on all these patent characteristics. We cannot reject the hypothesis that the coefficients on all these variables (except year dummies) are jointly equal to zero (*p-value* =of 0.73). Column (7) shows that we obtain similar results if we replace the linear control for age with a full set of age fixed effects. In this regression we cannot reject that the coefficients on the age dummies are jointly equal to zero (*p-value*=0.79) and the test that the coefficients on all patent characteristics (except year dummies) are zero (*p-value*=0.82).

We also provide additional evidence on the heterogeneity of the propensity of judges

to invalidate. Specifically, we show that the estimated variation is inconsistent with a setup in which judge decisions reflect identical voting propensities plus random error. To do this, we construct a counterfactual where judges vote according to the same random process. We generate a simulated judge vote that takes into account the effect of observable patent characteristics on the probability of invalidation. To construct the simulated votes, we use the following procedure. First, we regress the votes of each judge on observable characteristics of the cases, without including judge fixed effects, and then construct the predicted probability of an invalidity vote for each judge j for patent p, based on these characteristics, ϕ_{jp}, and the regression residuals, e_{jp}. Second, we add to the probability ϕ_{jp} a random draw ω_{jp} from a normal distribution with mean and standard deviation equal to the mean and standard deviation of the distribution of the regression residuals. Finally, the simulated invalidity vote for judge j for patent p is set equal to one if the sum of the predicted invalidity and the random draw $(\phi_{jp} + \omega_{jp})$ is above one. We obtain very similar results using different thresholds. We use the simulated vote to estimate judge fixed effects and find that they are not statistically significant (p-$value$=0.66). In Figure A2 we compare the distribution of these fixed effects from simulated votes with the (statistically significant) fixed effects estimated using actual voting behavior. The difference between the two distributions is striking: the variance of the Federal Circuit fixed effects is much larger than the one we would observe if judges were voting on the basis of identical propensities plus random error.

3 Effect of Patent Invalidation: Robustness

In this section we describe a series of robustness checks on our main finding of positive effect of invalidation on follow-on innovation. Estimates for some of these regressions are reported in Table A3.

First, in the text we treated an invalidation judgement as the final verdict. However, parties to the dispute have the right to appeal the decision of the Federal Circuit to the Supreme Court (which retains discretion over whether to hear the case). This means that invalidation of a patent by the Federal Circuit retains some uncertainty, so that downstream innovators whom the patent blocked might not respond until this uncertainty is removed. In our context, this is equivalent to saying that our key variable, *Invalidation*, contains some measurement error. In theory, any such error should be taken care of by our instrumental variable estimation.

Nonetheless, as a further check we identified that the patent invalidity cases appealed to the Supreme Court in our data set. Only 23 Federal Circuit decisions were reviewed by the Supreme Court in the period 1982-2008 (Golden, 2009). Only 12 of these cases are in our dataset (the others involve issues other than patent validity).We drop these cases and re-estimate the model (by IV). In column 1 of Table A3 we show that the estimated effect is very close to the baseline coefficient of 0.410.

Second, the baseline model incorporates fixed effects for six broad (one-digit) technology fields. We also estimate a specification which uses a more refined technology classification – 32 two-digit subcategories from the NBER. Column 2 of Table A3 shows that the point estimate of the coefficient on *Invalidation* is nearly double the baseline estimate but less precise, 0.915 (standard error = 0.422), and we cannot reject the null hypothesis that the two estimated coefficients are the same (*p-value*=0.11). In the paper we retain the one-digit technology field dummies. We do this because the empirical analysis often involves using smaller subsamples split along various dimensions. As a robustness check, we re-estimate all of those regressions using the more detailed, two-digit technology field dummies and obtain qualitatively (and in most cases, quantitatively) similar results, but the estimates are less precise.

Third, the citations information obtained from the USPTO ends in 2010, so the latest years in the sample are subject to truncation. Moreover, citations are obtained from a dataset of patent grants (but note that they are linked to litigated patents using the application date) - this is always the case using the USPTO - and this may potentially amplify the truncation concerns. We run a series of robustness checks to assess whether truncation is an issue in our study. First, while our baseline specification controls for year effects and these mitigate the truncation problem, there may be concern that the truncation problems differ by technology field since citation patterns vary. To address this, in Table A3 we introduce interactions between the year and technology category dummy variables to allow truncation effects to be different across technology classes. Column 3 shows that the coefficient on invalidity is very similar to the one estimated in the baseline. Column 4 restricts the sample to patent decisions that take place before 2003, for which we have a complete 5-year time window of citations and thus expect truncation to be less severe. Results are robust and the coefficient on invalidity is slightly larger than, but not statistically different from, the one in our baseline regression. We also estimate the invalidity effect focusing on patents litigated before 2000. This is equivalent

to dropping about half of the sample. Also in this case the result is robust and the coefficient is equal to 0.361 (std. error = 0.158), which is very similar to our baseline estimate.

As a complementary approach to correct more directly for truncation, we adjusted citation totals for Federal Circuit decisions in which we observe only a portion of the five year post-decision interval. To perform this correction, we exploit the "quasi-structural" citation lag distribution estimated in Hall, Jaffe and Trajtenberg (2001). They provide an estimate of the distribution of citations received over the life of patents across different technology classes.[1] Their estimates allow us to inflate the citations received by patents for which we observe only a fraction of the five year window, taking into account the propensity of patents in that technology field to be cited at the specific age in which data are missing.

This correction generates a negligible increase in the sample average of the citations received in the five years after invalidation (from 8.70 to 8.77). Nonetheless, for patents with truncated post-decision windows, the average external cites received increase by 40 percent (from 1.08 to 1.52), which highlights the impact of truncation. Nonetheless, in column 5, we show that replacing our main citation measure with the truncation-adjusted measure does not affect our baseline results. The estimated coefficient is 0.437, which is only slightly larger than the one estimated in our baseline regression. As an additional test, we exploit the estimates by Hall, Jaffe and Trajtenberg (2001) to inflate the citations received in the five year window after invalidation in order to obtain a predicted 10-year post-decision citations total for each patent. Not surprisingly, with this correction, the sample average of post-decision citations increases substantially, from 8.70 to 14.76. Nonetheless, the estimated impact of invalidation is robust – using this measure, the estimated coefficient is 0.460 (standard error =0.229). The estimated effect based on the post-decision 10-year window is very similar to the one estimated for the five year window when we adjust for truncation. This is consistent with our finding in the paper on the timing of the effect – where we estimated that most of the effect takes place in the third to sixth years following invalidation.

Fourth, one may be concerned that citations have a geographic component and the effect of invalidation could be different for foreign and domestic patentees (remember that all our

[1] The estimates are made under the assumptions of (i) proportionality (the shape of the distribution is independent of the total number of citations received), (ii) stationarity (the distribution does not depend on the cohort of the patent) and (iii) patents do not received citations after 35 years of age.

sample patents are in the U.S.). There are at least two reasons why the effect of invalidation of a patent in the U.S. might affect foreign follow-on innovators less strongly. First, if licensing barriers represent only a small fraction of the overall costs associated with entry into the U.S. market, removal of patent protection may not have a major impact on the decisions of later foreign innovators. Second, the patent in the U.S. confers protection only in the U.S., while the decision to innovate depends on profits that can earned in the global market. Thus the impact of invalidation of the U.S. patent will depend on the relative importance of the U.S. market to the follow-on innovators and this may differ for domestic and foreign innovators.

To address this concern empirically, we exploit the USPTO assignee type codes information and constructed new measures that distinguish between citations originating from domestic (U.S.) innovators citations and foreign innovators and then examine whether the invalidation effect is different for the two groups. On average, 27 percent of the citations received by the patents in our sample before the Federal Circuit decision belong to foreign innovators. We re-estimate our baseline regression using each of these two dependent variables and found that the invalidation effect is exclusively driven by an increase in citations by *domestic* follow-on innovators. Column 6 in Table A3 shows the estimated effect for subsequent citations by domestic patentees. For citations by foreign owners, the estimated coefficient is an order of magnitude smaller and not statistically significant. This result is consistent with the idea that it is costly for foreign innovators to patent in U.S. and that licensing frictions represent only a fraction of the total cost sustained (or possibly that the U.S. market is relatively less important for foreign innovators).

Additional Unreported Tests

Finally, we conduct a set of additional robustness tests that are not reported in Table A3. First, the baseline specification incorporates a full set of patent age fixed effects. However, the age distribution of citations may vary across technology fields (for evidence, Jaffe and Trajtenberg, 2002). To allow for this, we extend the specification by including a full set of interactions between the technology field and age dummies. The estimated coefficient on *Invalidation* is 0.401 (standard error = 0.192), which is nearly identical to the baseline estimate.

The second robustness check involves how to treat patents that receive no citations before the Federal Court decision (4 percent of the sample) and those that receive no cites in the five year widow after the decision (23 percent of the sample). In our baseline specification we 'fix'

this problem by using $log(PostCites + 1)$, which is common practice but may introduce bias. We re-estimate the baseline model adding dummy variables for patents that received no cites before the Federal Circuit decision and for patents that receive no cites after the decision. The results are robust – the point estimate on $Invalidation$ is 0.449 (standard error = 0.167). We get similar results if we drop these patents from the sample entirely. We get similar results if we use the number of citations without logarithmic transformation as the dependent variable (once suitably transformed for comparability). Finally, we also estimated a Poisson count model by instrumental variables (using the predicted probability of invalidation $\widehat{P}$ as the instrument). The point estimate is 0.638 (standard error = 0.321) which is larger than, but not statistically different from, the baseline coefficient. In the analysis in the paper, we do not use the Poisson model because the econometric techniques that we use to estimate the heterogenous marginal treatment effect of patent invalidation have only been developed for linear models.

Third, there is a concern that some Federal Circuit decisions may involve rulings that limit the scope of patentable subject matter rather than simply assessing the validity of the focal patent. Such decisions could reduce subsequent citations for the *entire* technology field, leading us to underestimate the true blocking effect of patent rights. To address this, we identified the most important Federal Circuit decisions that relate to patentable subject matter during our sample period (the main sources are Dolmeage, 2006 and Kappos et. al., 2008). We obtained a list of 14 Federal Circuit decisions that are concentrated in the areas of software, business methods and biotechnology, of which only three are in our sample. There are very few cases in our sample because most of the key Federal Circuit decisions on patentable subject matter do not involve granted patents but only patent applications. Moreover, because of their importance, some of these cases are decided 'en banc' by the entire court and not by a panel of three judges. We excluded such special cases from our sample. Dropping the three decisions that were in our sample and re-estimating the model, we obtain coefficients that are nearly identical to the baseline estimates.

Finally, we examine robustness of our results when controlling for competition in two ways. First, the size of the patentee is likely to shape the levels of product market and technology competition with other firms. We control for the size of the patentee, measured as the size of its patent portfolio (number of patents granted in the five years preceding the Federal Circuit decision) and obtain an estimated invalidation effect that is very similar to the baseline.

Second, competitive pressure is also likely to depend on the concentration of firms operating in a technology area. Controls for concentration with the share of the four largest patentees (or the Herfindahl index of patent ownership) in the two-digit technology sub-category of the litigated patent yields very similar results.

We also look at the difference in invalidation effect across different samples comprising patents with different level of pre-invalidation citations. To conduct these exercises, we create a new variable equal to the (logarithm) of pre-invalidation citations filtered by age effects, year effects, technology effects, number of claims and pre-invalidation self-citations. We find no evidence of a larger (or smaller) effect of invalidation for patents that receive more pre-invalidation (normalized) citations.

4 Intensive and Extensive Margins: Robustness

We conduct extensive robustness checks on the regressions in Table 8. We report the estimates for some of these regressions in Table A4. First, we vary the thresholds for defining 'small' firms ($\leq$ 1, 10 and 15 patents), and for defining 'large' firms ($\geq$ 75, 110 and 150 patents). Second, we re-estimate the effects of patent invalidation by splitting the samples between large and non-large patentees. We also break down the category of non-large patentees into two groups, small and medium sized firms. In all of these experiments, the pattern that emerges in Table 8 is extremely robust. In every case the effect of invalidation is concentrated on the subsequent citations by small innovators to focal patents held by large firms, and it is predominantly an extensive margin effect.

Because of sample size, we cannot allow the effect of invalidation to vary with technology field in these regressions (we do allow for additive field effects, however). If citations from small citers to large patentees are overrepresented in fragmented and complex technology fields, where we found blockage was more likely, our finding that blocking effect of invalidation is limited to the large patentee-small citing firm category could be simply a technology field composition effect. To check this concern, we examined the percent of citations in each technology field accounted for by citations by small to large patentees. The technology fields where invalidation has a statistically significant blocking effect (medical instruments, electronics and computers) are not those with the largest fraction of citations from small to large patentees – the mean fraction of sample citations from small to large patentees is 7.4 percent in these fields, as

compared to 9.9 percent in the other fields. We conclude that our empirical finding is not due to a technology field composition effect.

5 Using Non-Patent Data to Measure Follow-on Innovation: Data Construction and Robustness

5.1 Medical Instruments

Data Construction

Following the Medical Device Amendments Act passed by Congress in 1976, the U.S. Food and Drug Administration (FDA) has primary authority to regulate medical devices sold in U.S. These products are subject to a regulatory process that requires detailed product information and evidence of safety from clinical trials. The FDA releases data on approval requested for medical instruments. We construct a new measure of cumulative innovation by linking FDA approval requests to the medical instrument patents in our sample.

Specifically, there are two types of medical instrument approval processes that can take place at the FDA. The first and more common approval is the 510k premarketing submission made to FDA to demonstrate that the device to be marketed is safe and effective. This procedure is for 'low risk' and 'moderate risk' devices. The second, less common and more stringent approval is the Premarket Approval (PMA) application for 'high risk' devices. The FDA data indicate that, for the period 1981-2013, there were about 125,000 requests for 510k approvals and 26,000 requests for PMA approval.

In our sample of patents litigated at the Federal Circuit, there are 121 cases involving medical devices (NBER sub-categories 32 and 39). To link patents with FDA approval requests, for each patent we identify a set of keywords related to the technology and search for FDA approval requests that contain such keywords. We carefully read the abstract and title of each litigated patent and collected one 'primary' keyword and up to two 'secondary' keywords for each of the patents in our sample. Exploiting these keywords, we construct our first two measures of cumulative innovation. The first one, *Approvals_1*, is equal to the number of FDA approval requests (PMA and 510k) for which the product name contains the primary keyword. The second one, *Approvals_2*, is equal to the number of FDA approval requests (PMA and 510k) for which the product name contains at least one of the keywords (primary or secondary) of the patent. As in our baseline analysis, the dependent variables are constructed focusing on

a five year window following the Federal Circuit decision. The regressions control for the total number of applications in the period between the grant of the patent and the Federal Circuit decision, as well as for the other controls employed in our baseline regression.

The correlation between citations received in the five years after the decision and Approvals_1 is 0.3; the correlation with Approvals_2 is 0.1. If we exclude potentially unmatched patents (patents receiving citations before invalidation but that could not be matched to any FDA approval requests), the correlation between cites and Approvals_1 increases to 0.4 and the one between cites and Approvals_2 increases to 0.3.

As a second approach to exploit FDA approval data to measure cumulative innovation, we assign each litigated patent to a set of product codes among the roughly 6,000 product codes in which the FDA classifies medical devices. Also in this case we distinguish between a 'primary' product code and up to two 'secondary' codes. We construct two additional measures of cumulative innovation counting the number of FDA approval requests (PMA and 510k) in the same product code of the litigated patent. The first measure only counts applications in the primary product code, whereas the second one includes primary and secondary product codes. The correlation with the citations received before invalidation is roughly equal to 0.1 for both of these measures of cumulative innovation.

Both measures based on "primary" and "extended" sets of keywords and product classes may be associated with two distinct types of measurement errors. Focusing on primary keywords and product classes may lead to under-counting follow-on innovation, if applications related to the patented technologies do not include the keyword or are in a different product classes. However, enlarging the number of keywords or patents may lead to over-counting follow-on innovation if we include products that are not related to the patented technology.

Results

Column 2 in Panel A of Table 9 in the paper shows the effect of invalidation using Approvals_2 as dependent variable, the measure constructed exploiting an extended set of keywords. The estimated effect using Approvals_1 as dependent variable, the measure based only on the primary keyword is very similar – the point estimate is 1.116 (standard error =0.617). Column 3 in the same panel uses the measure based on the applications in all the product classes (primary and secondary) assigned to the patent. The estimated coefficient using as the dependent variable the number of applications in the primary product class assigned

to the patent is 0.378 (std. err. = 0.635). The sign of the coefficient suggests a positive effect of invalidation on follow on innovation, but the estimate is not statistically significant. As we discussed above, measures based on primary and extended sets are likely to be subject to countervailing measurement errors. The differences in magnitude, and statistical insignificance using the measure based on products in the primary class, suggest that this measure may substantially under-count the extent of cumulative innovation (and thus be subject to attenuation bias).

5.2 Drugs

Data Construction

In our sample 167 patent cases involve drug patents (patents belonging to the NBER category 3 but not in the sub-categories of Medical Instruments and Biotechnology). As an alternative to citations, in this technology field we measure cumulative innovation by identifying the subsequent clinical trials that are related to each Federal Circuit patent. To this end, we match each patent to the trade name of a drug, recover information on the active ingredients of the drug, and collect data on clinical trials that refer to the active ingredients.

Our data source for clinical trials is the website ClinicalTrials.gov, which is a registry database of publicly and privately supported clinical studies of human participants. ClinicalTrials.gov was created as a result of the Food and Drug Administration Modernization Act of 1997 that required the U.S. Department of Health and Human Services to establish a registry of clinical trials information. The site was made available to the public in 2000 and only reports clinical trials from 2000 onwards. Therefore, in order to have at least two years of post-litigation clinical trials data for each patent, most of our analysis will focus on patents litigated after 1997 (140 patent cases).

For each of the Federal Circuit drug patent in our sample, we identified the trade name of the drug protected by the patent. Such information was obtained from a multiplicity of sources: the text of the court decision, the FDA Orange Book (Approved Drug Products with Therapeutic Equivalence Evaluations) and specialized web-sites (in particular, DrugPatentWatch.com). We were able to match a patent with a drug for 94 patent cases, about 68 percent of our post-1997 sample. Cases in which the drug was not identified mostly involved patents related to DNA recombination, gene expressions and methods for conducting clinical tests. For

the cases in which we were not able to match a patent with a specific drug, we collected up to three keywords describing the technology, after careful reading of the patent title and abstract.

For all these matched drugs, we identify the active ingredients. Also in this case the information was obtained from a multiplicity of sources: the text of the court decision, the Orange Book and specialized web-sites (especially from DrugPatentWatch.com or the Merck Index Online). Only 10 of the 94 patents have multiple active ingredients. Finally, for each of the matched patents, we obtained information on the clinical trials related to the active ingredient conducted in the five years following the Federal Circuit decision. The information was obtained from a text search in ClinicalTrials.gov with the active ingredient as keyword. In the case of multiple active ingredients, we search for each ingredient separately and measure clinical trials as the sum of trials involving any of the ingredients.

The average number of clinical trials in the five years that follow invalidation is 134.41 (standard deviation = 125.76). The correlation between citations received in the five years after the Federal Circuit decision and clinical trials is 0.15, and 0.18 in the subsample of drugs with only one active ingredient.

Results

Column 2 in Panel A of Table 9 in the paper shows the invalidation effect using the number of trials as the dependent variable. This column focuses only of the sample of matched drugs, those for which we identified a trade name associated with the patent. Results are consistent with those obtained using the patent citation measure. The coefficient is positive but statistically insignificant and its magnitude is very similar.

In column 3 we enlarge the sample including unmatched patents. The number of clinical trials for these patents is equal to the number of trials referring both to the primary and secondary keywords of the patent. The coefficient remains insignificant and its magnitude is slightly higher than the one in the previous columns. If we measure clinical trials for unmatched drugs by counting the clinical trials, referring only to the primary keyword, the estimated coefficient is 0.364 (standard error = 1.128) and also in this case the estimate is statistically insignificant.

6 Substitute Patents: Data Construction

We define a U.S. patent as related to the litigated patent if it has been granted before the decision date and it appears in the top ten related patent documents listed by Google. We collect a maximum of five related patents for each patent litigated at the Federal Circuit. The Google algorithm identified at least one related patent for 699 Federal Circuit decisions (about 52 percent of our sample). About 27 percent of these observations have only one matched related patent, 20 percent have two related patents and 53 percent have three or more related patents. For the Federal Circuit patents for which we were able to find at least one match, we first confirm that the estimated causal effect of invalidation on citations is similar to the one obtained in the full sample. Specifically, the IV coefficient is 0.541 (std. error = 0.257), implying an increase in citations of approximately 70 percent.

We then re-estimated the baseline model but using as the dependent variable the citations to the related patents. In doing this, we experiment with a variety of alternative samples (e.g., balanced samples with top four or five related patents, as well as unbalanced samples where we keep all of the related patents identified by Google). All the regressions show a negative relationship between the post-decision citations to related patents and the invalidation of the focal patent, providing some support for the substitution hypothesis. Nonetheless, the estimates tend to be small and often statistically insignificant, confirming that the substitution hypothesis cannot explain much of the increase in citation caused by Federal Circuit invalidation.

References

[1] Dolmeage, Brianna (2006), "The evolution of patentable subject matter in the United States," *Whittier Law Review,* 27: 1023-1045

[2] Feddersen, Timothy and Wolfgang Pesendorfer (1996) "The Swing Voter Curse," *American Economic Review,* 86: 408-424

[3] Golden, John (2009), "The Supreme Court as "Prime Percolator": A Prescription for Appellate Review of Questions in Patent Law," *UCLA Law Review,* 56: 657-724

[4] Hall, Bronwyn, Adam Jaffe and Manuel Tratjenberg (2001) "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," NBER Working Paper 8498

[5] Kappos David, John Thomas and Randall Bluestone (2008), "A Technological contribution requirement for patentable subject matter: Supreme Court Precedent and Policy," *Northwestern Journal of Technology and Intellectual Property,* 6: 152-170

Table A1. Federal Circuit Judges

Judge	*Active Service*	*Validity Decisions 1982-2008*	*Percentage of Decisions in which the Judge voted for Invalidation*
Randall Ray Rader	1990-	242	39.6
Daniel Mortimer Friedman	1982–1989	112	21.2
Pauline Newman	1984-	309	26.9
Glenn Leroy Archer, Jr.	1985–1997	170	34.7
Haldane Robert Mayer	1987–2010	269	42.4
S. Jay Plager	1989–2000	153	35.3
Alan David Lourie	1990-	293	46.8
Raymond Charles Clevenger III	1990–2006	232	37.9
Alvin Anthony Schall	1992–2009	248	37.5
William Curtis Bryson	1994-	238	44.1
Arthur J. Gajarsa	1997–2011	164	41.5
Richard Linn	1999–	111	43.2
Timothy B. Dyk	2000-	131	37.4
Sharon Prost	2001-	106	40.6
Kimberly Ann Moore	2006-	21	76.2
Giles Sutherland Rich	1982–1999	152	40.8
Arnold Wilson Cowen	1982-2007	59	33.9
Oscar Hirsh Davis	1982–1988	70	50.1
Philip Nichols, Jr.	1982-1990	38	26.3
Byron George Skelton	1982–2004	56	33.9
Phillip Benjamin Baldwin	1982-1991	54	25.9
Howard Thomas Markey	1982–1991	138	49.3
Marion Tinsley Bennett	1982–2000	57	57.9
Shiro Kashiwa	1982-1986	34	38.2
Jack Richard Miller	1982-1994	35	42.9
Edward Samuel Smith	1982-2001	91	36.3
Paul Redmond Michel	1988–2010	245	41.6
Helen Wilson Nies	1982–1996	89	38.2
Jean Galloway Bissell	1984–1990	41	24.4

TABLE A2. Exogeneity of Judge Panels (OLS Regressions)

	1	2	3	4	5	6	7
Dependend Variable	JIP	JIP	JIP	JIP	JIP	JIP	JIP
log(Claims)	-0.001					-0.001	-0.001
	(0.002)					(0.002)	(0.002)
log(PreCites)		-0.001				0.001	0.001
		(0.002)				(0.002)	(0.002)
log(PreSelfCites)			-0.001			-0.001	-0.001
			(0.002)			(0.002)	(0.002)
Age				-0.001		-0.001	
				(0.001)		(0.001)	
Chemicals					0.003	0.003	0.003
					(0.006)	(0.006)	(0.006)
Computers and Communication					-0.007	-0.007	-0.008
					(0.006)	(0.006)	(0.006)
Drugs and Medical					0.004	0.004	0.004
					(0.005)	(0.005)	(0.005)
Electrical and Electronics					0.005	0.005	0.006
					(0.006)	(0.006)	(0.006)
Mechanicals					0.007	0.007	0.007
					(0.007)	(0.007)	(0.007)
Year Effects	YES***	YES***	YES***	YES***	YES***	YES***	YES***
Age Effects	NO	NO	NO	NO	NO	NO	YES
Fed. Circuit Decisions	1357	1357	1357	1357	1357	1357	1357
F-stat for all patent characteristics (except year dummies) being zero					F=1.00 p=0.41	F=0.67 p=0.73	F=0.77 p=0.82

NOTES: * significant at 10 percent, ** significant at 5 percent and *** significant at 1 percent. Robust standard errors are reported in parentheses. Invalidated=1 if Federal Circuit invalidates at least one claim of focal patent. PreCites = cites from patents of other assignees received before Federal Circuit decision. PreSelfCites = cites received from patents owned by same patentee of focal patent before Federal Circuit decision. Claims = total number of claims listed in focal patent. Age = age in years from filing date of patent at Federal Circuit decision. Year= year of Federal Circuit Decision. JIP= propensity to vote for patent invalidity of judge panel constructed from invalidity votes of judges in other sample cases. We add one to all citation measures to include patents with zero cites.

Table A3. Impact of Invalidation on Citations - Robustness (IV Regressions)

	1	2	3	4	5	6
Dependent Variable	log(PostCites)	log(PostCites)	log(PostCites)	log(PostCites)	log(PostCites) Truncation-Adjusted	log(PostCites) Domestic Citations
Invalidated	0.394**	0.915**	0.392**	0.484**	0.473**	0.529***
	(0.197)	(0.422)	(0.180)	(0.214)	(0.201)	(0.206)
Refined (2-digit) Tech dummies		YES***				
Tech x Year Effects			YES***			
Sample	Drop Supreme Court Appeals	Full	Full	Decisions up to 2003	Full	Full
Fed. Circuit Decisions	1345	1357	1357	1001	1357	1357

NOTES: * significant at 10 percent, ** significant at 5 percent and *** significant at 1 percent. Robust standard errors are reported in parentheses. All regressions control for log(PreCites), log(PreSelfCites), log(Claims), age and year effects. PostCites = cites from patents of other assignees in 5 year window after Federal Circuit decision. Invalidated=1 if Federal Circuit invalidates at least one claim of focal patent. Technology field effects use 6 categories in columns 1 and 32 subcategories in column 2 (for details see Hall et al. 2001). We add one to all citation measures to include patents with zero cites.

Table A4. Intensive and Extensive Margins - Robustness (IV Estimates)

Dependent Variable	Total Effect log(PostCites)			Extensive Margin log(Number of Distinct Assignees)		
	1	**2**	**3**	**4**	**5**	**6**
	Citing Patents in Small Portfolios (< 5 patents)	Citing Patents in Small Portfolios (< 2 patents)	Citing Patents in Small Portfolios (<2 patents)	Citing Patents in Small Portfolios (< 5 patents)	Citing Patents in Small Portfolios (< 2 patents)	Citing Patents in Small Portfolios (<2 patents)
Invalidated	0.046	0.125	0.128	0.015	0.076	0.088
	(0.179)	(0.168)	(0.165)	(0.152)	(0.143)	(0.141)
Invalidated X Large Patentee (> 75 patents)	2.552**		2.248*	1.842**		1.390*
	(1.360)		(1.277)	(0.951)		(0.745)
Invalidated X Large Patentee (> 102 patents)		1.769**			1.216**	
		(0.752)			(0.550)	

NOTES: * significant at 10 percent, ** significant at 5 percent and *** significant at 1 percent. Robust standard errors are reported in parentheses. All regressions control for log(PreCites) in the size group, log(PreSelfCites), log(Claims), age and year effects. PostCites = cites from patents of other assignees in 5 year window after Federal Circuit decision. Invalidated=1 if Federal Circuit invalidates at least one claim of focal patent. Invalidated and its interactions are instrumented by the Probit estimates of the probability of invalidation and its interactions. We add one to all citations measures to include patents with zero cites.

Figure A1. Age Distribution of Litigated Patents

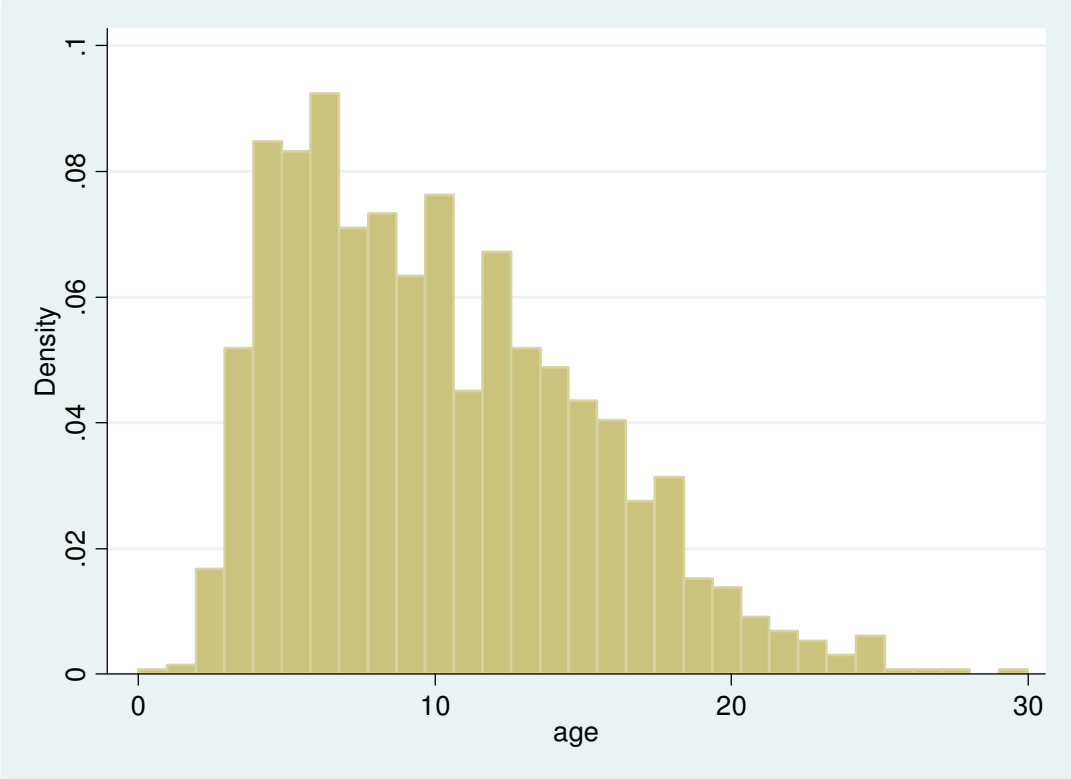

Figure A2. Simulated and Estimated Judge Fixed Effects

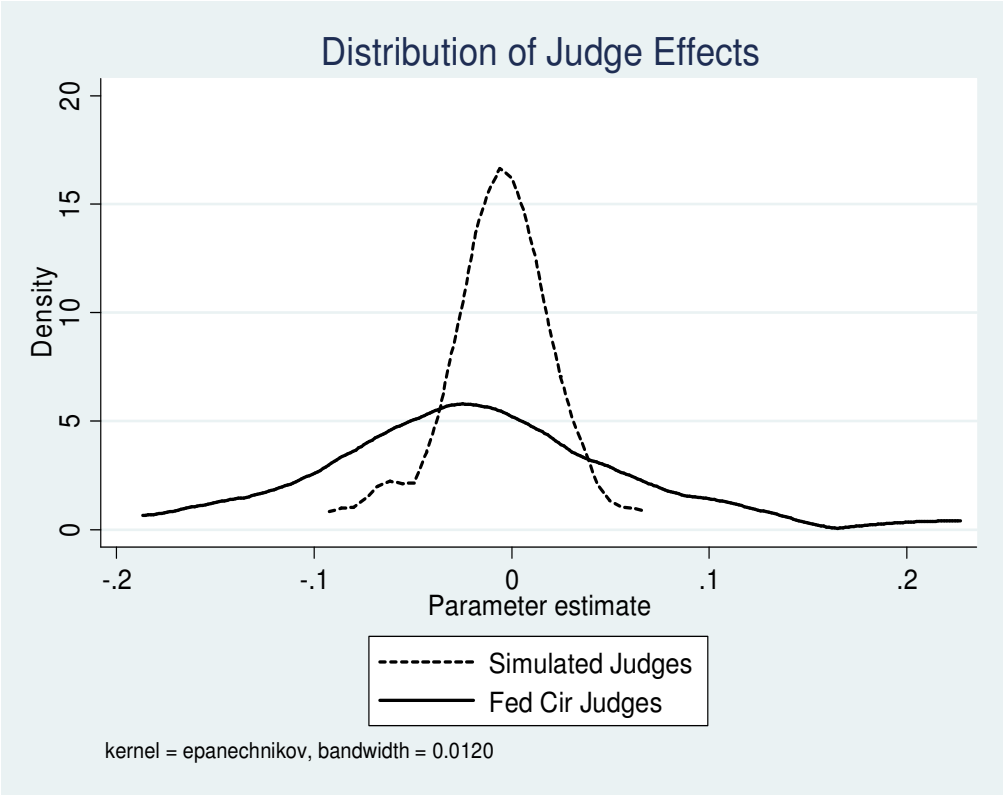

NOTES: The Figures compares the distribution of judges fixed effects in our data (solid line) with the distribution obtained from a simulation in which each judge votes with the same random process.