Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Racing to file: Patent priority rules, technological competition and firm patenting after the America Invents Act

Kirill Proskurin

koproskurin@edu.hse.ru

Southwestern University of Finance and Economics (SWUFE)

Research Article

Keywords: America Invents Act, First-inventor-to-file, Patent races, Technological competition, Difference-in-differences, Patent policy

Posted Date: November 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11293155/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Racing to file: Patent priority rules, technological competition and firm patenting after the America Invents Act

Kirill Proskurin[1,2]

[1] Southwestern University of Finance and Economics (SWUFE), Chengdu, Sichuan, China

[2] International College of Economics and Finance, HSE University, 11 Pokrovsky Boulevard, Moscow 109028, Russia

Corresponding author: Kirill Proskurin, e-mail: koproskurin@edu.hse.ru

ORCID: 0009-0005-5690-4185

Abstract In this paper I study how patent priority rules affect the way firms file patents. I use the Leahy–Smith America Invents Act (AIA), which in March 2013 replaced the US first-to-invent rule with a first-inventor-to-file rule. After the reform, the filing date decides who gets the patent. So the reform should matter most in technology areas where several firms work on similar inventions at the same time. I use a continuous difference-in-differences design and a firm–technology class–year panel of US publicly traded firms for 2007–2016, and I compare technology classes with different levels of patenting competition before the reform. Patent applications grew more after the reform in more competitive classes. Among active filers, moving from the 25th to the 75th percentile of pre-reform competition is associated with about 22% more applications. The response starts before the rule took effect, holds within the same firm and year, and also appears in granted patents, especially in physics and chemistry, and in electricity for applications. Average patent quality does not clearly improve, although the number of highly cited patents rises. Firms that are more exposed to competitive classes also enter fewer new technology classes and spend less on R&D. These results suggest that the reform made firms race harder to file within the technology areas where they were already active, rather than broadly increasing invention. For competition policy, they show that a procedural rule in patent law can change how firms compete in technology markets.

Keywords America Invents Act · First-inventor-to-file · Patent races · Technological competition · Difference-in-differences · Patent policy

JEL Classification O34 · O31 · O38 · L13

1 Introduction

Patents are one of the main tools of innovation policy. New knowledge is easy to copy, so inventors receive only part of the social benefit of their research, and markets tend to invest too little in it (Arrow 1962). Patents fix this problem by giving inventors a temporary right to exclude competitors. However, the rules that decide who receives this right also affect how firms compete. When several firms work on similar inventions, the priority rule decides who wins the race, and therefore how much effort firms put into being first.

In most countries, priority goes to the first applicant who files. The United States was an exception for a long time. Until 2013 it used a first-to-invent rule: a firm that could prove an earlier invention date could keep priority even if a rival filed first (Lemley and Chien 2003; Lo and Sutthiphisal 2009). The Leahy–Smith America Invents Act (AIA), signed in September 2011, replaced this rule with a first-inventor-to-file (FITF) system for applications

filed from 16 March 2013. It was the biggest change to US patent law since 1952. After the reform, apart from two narrow exceptions, the filing date alone decides who gets the patent.

In this paper I ask whether this change in the priority rule changed how firms patent, and where. The idea is simple. The FITF rule made it more costly to file later than a rival. This cost should matter most in technology areas where many firms work on similar inventions, because there the risk that a rival files first is highest. Simultaneous invention is common (Lemley 2012), and patent races are concentrated in some technology classes (Thompson and Kuhn 2020). So the reform should have a stronger effect in technology classes where patenting competition was already high before the reform.

I test this prediction with a continuous difference-in-differences design. The treatment intensity is the level of patenting competition in each Cooperative Patent Classification (CPC) subclass before the reform. I measure it as one minus the Herfindahl–Hirschman index (HHI) of firms' application shares over 2007–2010. The data combine firm-level patent records from DISCERN 2.0 (Arora et al. 2024), patent classification, claims and citation data from PatentsView, and financial data from Compustat. The main sample is a firm–CPC subclass–year panel of US publicly traded firms that do R&D. It compares 2007–2010 with 2013–2016 and contains 56,267 observations for 1,610 firms and 319 subclasses. All regressions include firm, class and year fixed effects, and standard errors are clustered by firm and by class.

The results support the prediction. After the reform, patent applications grew more in more competitive classes. Moving from the 25th to the 75th percentile of pre-reform competition is associated with about 22% more applications among active filers. The estimate hardly changes when I add firm × year fixed effects, which absorb all shocks to a given firm in a given year. So firms shifted their filing towards their more competitive classes. The response starts before the rule took legal effect: the first signs appear in 2010, when the direction of the reform became clear, and the response becomes stronger in 2011–2012. This is what we would expect from forward-looking firms. Granted patents show the same pattern, so the extra filings were not just applications that later failed examination. The effect is strongest in physics, chemistry and electricity, and small and statistically insignificant in more traditional fields such as mechanical engineering and textiles.

However, the extra patenting does not look like a broad increase in innovation. Average claims and citations per patent do not clearly increase, although the number of highly cited patents does. More than half of the grant response disappears when I exclude continuation, divisional and continuation-in-part filings. Firms that are more exposed to competitive classes become less likely to enter new technology classes and cut their R&D spending. Taken

together, these results suggest that the reform made firms race harder to file in the technology areas where they were already active.

The paper contributes to three groups of studies. First, it adds to the evaluation of the AIA. Most studies of the reform compare small and large applicants or firms with different exposure to litigation (Lerner et al. 2015; Chakravarty 2023; Driver 2025), or look at aggregate patenting (Dohse et al. 2023). The closest papers are Huang et al. (2021), who show that innovation-intensive firms disclosed less about their R&D after the AIA, and Valentine et al. (2025), who show that firms lagging behind rivals in a CPC subclass published more scientific articles after the AIA to create prior art against competitors, especially in more competitive classes. These studies show that competitive structure shaped how firms responded to the AIA, but they do not study patent filing itself. I show that the filing response also depended on technological competition, and I follow its consequences for grants, patent quality, technological scope and R&D investment.

Second, the paper is related to studies that use class-level differences in competition or patent rights to study innovation, such as Moser and Voena (2012), Poege (2022), Watzinger and Schnitzer (2022) and Hegde et al. (2023). In those studies the shock changes market structure, access to patented knowledge or the timing of disclosure. Here the main change is to the rule that decides who wins when rivals reach the same invention, which allows a closer focus on the priority rule. Third, the paper adds to the industrial organisation literature on patent races (Loury 1979; Lee and Wilde 1980; Aghion et al. 2005). Theory has long warned that races can make firms spend resources on being first instead of on new knowledge. My results are in line with this concern: the reform increased filing where races were most likely, but R&D and average patent quality did not rise with it. This matters for competition policy, because the design of patent procedures can push firms to compete by filing rather than by inventing.

The paper is organised as follows. Section 2 describes the reform and the related literature. Section 3 describes the data and the competition measure, and Section 4 the empirical strategy. Section 5 presents the results, Section 6 discusses mechanisms and policy implications, and Section 7 concludes.

2 Institutional background and related literature

2.1 The America Invents Act and the priority rule

The AIA was signed into law on 16 September 2011, and its provisions started at different dates (Table 1). Its main change was to replace the first-to-invent (FTI) priority rule with the first-inventor-to-file (FITF) rule. The FITF rule applies to applications with an effective filing date on or after 16 March 2013; earlier applications stayed under the old rules. The move to FITF made US rules similar to those of other major patent systems and was expected to make priority more predictable (Lemley and Chien 2003; Abrams and Wagner 2013).

Under FTI, competing claims to the same invention were settled in interference proceedings. In these proceedings the patent office had to reconstruct when each inventor first had the full idea, when the invention was shown to work, and whether the inventor kept working on it in between. These proceedings were complex and costly, but they protected inventors who filed second. Lemley and Chien (2003) find that in more than 40% of interference cases the first inventor was the last to file. The AIA replaced interference proceedings with narrower derivation proceedings. In a derivation proceeding, a later filer can challenge an earlier application only by showing that the invention was taken from them. The AIA also kept a one-year grace period, during which an inventor's own public disclosure does not count as prior art against them. So the reform did not create a pure first-to-file system, but outside these two cases the filing date alone decides priority.

Table 1 Main implementation dates of the America Invents Act

Date	Provision
16 September 2011	Enactment; inventor oath requirement (§4); higher threshold for inter partes reexamination (§6); fee-setting authority and micro-entity fees (§10); failure to disclose best mode no longer invalidates a patent (§15)
16 September 2012	New post-grant trial proceedings at the patent office (§6); third-party submission of prior art (§8); transitional review programme for business method patents (§18)
16 March 2013	First-inventor-to-file rule takes effect (§3); derivation proceedings replace interference proceedings (§3)

Section numbers refer to the Leahy–Smith America Invents Act (Pub. L. 112–29)

The timing is important for the empirical design. The FITF rule started eighteen months after the AIA was signed, so firms could expect it and change their filing before March 2013. Bills with a filing-based priority rule had been discussed in Congress since 2009, and by March 2010 the Senate Judiciary Committee had agreed on a draft that was close to the final

law. For this reason I treat 2011–2012 as a transition period, exclude it from the baseline comparison, and study it separately.

2.2 Related literature

There is little evidence on changes in patent priority rules, because few countries have moved from first-to-invent to first-to-file. The closest case is Canada, which made this change in 1989. Lo and Sutthiphisal (2009) find that the Canadian reform did not increase R&D or overall patenting, but shifted patenting towards large corporations. Abrams and Wagner (2013) find that it reduced the share of patents granted to individual inventors by about 14%, without improving patent quality. These studies suggest that first-to-file rules change who gets patent rights more than how much is invented.

For the AIA itself, the average effect is hard to identify, because the reform applied to all applicants at the same time. A simple comparison of the periods before and after the reform mixes the reform with changes in patent examination and in economic conditions. Lerner et al. (2015) find no clear short-run effect on small patent-intensive firms. Dohse et al. (2023) find more applications and fewer grants after the AIA, but this comparison is descriptive. Later studies therefore use differences in exposure to the reform. Chakravarty (2023) finds that small entities filed fewer patents than large entities after the AIA, especially when they were more exposed to litigation. Driver (2025) finds that the reform reduced innovation investment among smaller firms that value patents and increased it among larger ones. Helmers and Love (2024) study the new post-grant review rules and find that exposure to Patent Trial and Appeal Board proceedings is associated with more R&D and more patent filings. Huang et al. (2021) and Valentine et al. (2025), discussed above, show that firms changed what they disclose in response to the reform, and more so in competitive technology areas.

The mechanism I study comes from the theory of competition and innovation. In the escape-competition logic of Aghion et al. (2005), firms that are close to their rivals have strong incentives to innovate to get ahead of them. Bloom et al. (2013) show that closeness in technology space is a separate type of rivalry, different from competition in product markets. The FITF rule works directly on this type of rivalry: it makes it more valuable to file before technologically close rivals. Under FTI, a firm that filed later could still win by proving an earlier invention date. This option was most valuable where simultaneous invention was most likely, and FITF removed it. At the same time, the patent-race literature shows that races can make firms duplicate effort and spend resources only on being first (Loury 1979; Lee and

Wilde 1980). So whether a stronger race to file leads to more innovation, or mainly moves effort towards filing, is an empirical question.

2.3 Predictions

This discussion gives three predictions. First, after the FITF rule, patent filing should increase more in technology classes with higher competition before the reform. Second, because the rule was known before it took effect, part of the response may appear in the 2011–2012 transition period. Third, if the response comes from a race to file and not from more invention, it should be stronger for patent counts than for patent quality, research investment or technological breadth.

3 Data and measurement

3.1 Data sources and sample

I use three data sources. Firm-level patent records come from DISCERN 2.0 (Arora et al. 2024). This database links USPTO patents to their ultimate owners among US publicly traded firms that do R&D, over 1980–2021, and it takes subsidiaries and ownership changes into account. Patent classification, claims and citation data come from PatentsView. Financial data come from Compustat (industrial, consolidated, annual) and are merged to DISCERN firms through a dynamic CRSP–Compustat identifier link.

I build two outcomes. Patent applications are pre-grant publications in DISCERN matched to the PatentsView CPC classification of pre-grant publications. Granted patents are DISCERN grant records matched to the PatentsView CPC classification of granted patents. Granted patents are dated by the year in which the application was filed, not by the grant year, so both outcomes follow the same filing timeline. The unit of technology space is the CPC subclass (the four-character CPC code). One patent can belong to several subclasses, so I aggregate the patent–class links to firm–subclass–year counts. Claims and forward citations are available for granted patents and come from the PatentsView grant files.

The baseline comparison uses 2007–2010 as the pre-reform period and 2013–2016 as the post-reform period. I exclude the transition years 2011–2012 from the baseline regressions and study them separately. The application sample contains 56,267 firm–subclass–year observations for 1,610 firms and 319 CPC subclasses. The granted-patent sample comes from a separate PatentsView file and is larger: 90,072 observations, 1,737 firms and 529 subclasses. All samples keep only observations with non-missing outcomes, treatment and

controls, and singletons are dropped. I estimate linear regressions with high-dimensional fixed effects.

3.2 Measuring technological competition

I measure technological competition in CPC subclass c with the HHI of firms' application shares, following Poege (2022) and Watzinger and Schnitzer (2022):

$$HHI_{ct} = \sum_i \left(\frac{Apps_{ict}}{\sum_j A\,pps_{jct}} \right)^2 \tag{1}$$

where $Apps_{ict}$ is the number of applications filed by firm i in class c in year t. The index equals one when one firm files all applications, and $1/N$ when N firms have equal shares. The competition index is one minus the HHI, averaged over the pre-reform years so that it cannot react to the reform:

$$Competition_c^{pre} = \frac{1}{|T^{pre}|} \sum_{t \in T^{pre}} (1 - HHI_{ct}), \qquad T^{pre} = \{2007, \ldots, 2010\} \tag{2}$$

A value of zero means that only one firm filed in the class, and values close to one mean that many firms had similar shares. The index measures active patenting competition. Firms that did not file in a class are not included, so potential entrants are not captured.

I compute the index from the publicly traded firms in DISCERN, because the paper studies how these firms respond, and their closest rivals are firms of the same type. To check that this choice does not change the ranking of classes, I compute the same index from all USPTO assignees in PatentsView over 2007–2010, including small firms, individual inventors and foreign applicants. The Spearman rank correlation between the two measures across subclasses is 0.885 (Fig. A1 in the Appendix). Fig. A1 also shows one difference. Subclasses in which only one DISCERN firm filed have a DISCERN-based value of zero, but often a much higher value when all assignees are counted. For these classes the DISCERN-based measure overstates concentration.

Fig. 1 shows the distribution of $Competition_c^{pre}$ across observations in the application sample. Most observations lie between 0.4 and 0.9, and there is a separate group at zero, which corresponds to the single-filer subclasses described above. The mean is 0.558 and the standard deviation is 0.279 (Table 2). The 25th and 75th percentiles are 0.464 and 0.753, so the interquartile range is 0.289.

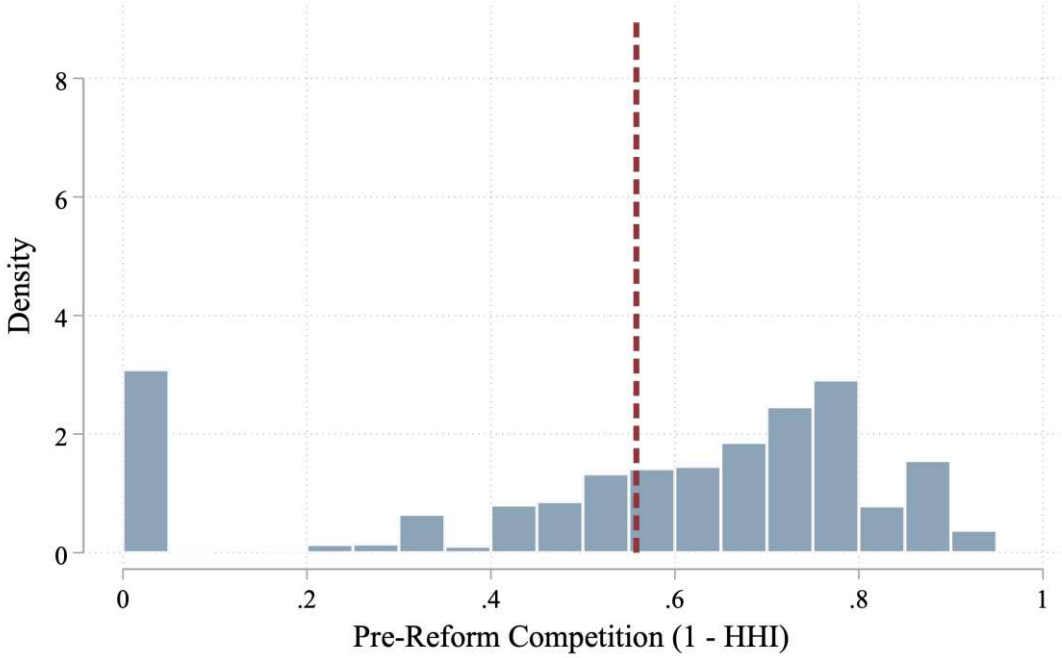


Fig. 1 Distribution of pre-reform technological competition, $Competition_c^{pre}$, across firm–CPC subclass–year observations in the application sample. The dashed line marks the sample mean (0.558)

3.3 Descriptive statistics

Table 2 summarises the application sample. The outcome, log patent applications, has a mean of 0.928 and a median of 0.693 (two applications). It is right-skewed: most firm–class–year cells contain few applications, and a few contain very many. The firms are large and research-intensive. In the average observation, R&D equals 8.0% of assets, capital expenditure 3.5%, debt 25.6% and cash 21.9%. Financial controls are winsorised at the 1st and 99th percentiles.

Table 2 Summary statistics, application sample

Variable	Mean	SD	P25	Median	P75	Min	Max
Outcome							
log(Patent applications)	0.928	1.178	0.000	0.693	1.609	0.000	8.645
Treatment							
$Competition^{pre}$	0.558	0.279	0.464	0.648	0.753	0.000	0.902
Firm controls							
log(Assets)	8.717	2.419	7.202	8.907	10.521	2.078	13.108
R&D intensity	0.080	0.137	0.018	0.038	0.086	0.001	4.323
Capex intensity	0.035	0.031	0.017	0.028	0.043	0.000	0.577
Leverage	0.256	0.231	0.130	0.237	0.348	0.000	8.461
Cash ratio	0.219	0.210	0.075	0.150	0.279	0.001	0.998
Return on assets	−0.002	0.263	0.015	0.054	0.091	−6.643	1.174
log(Market capitalisation)	9.028	2.290	7.532	9.222	10.864	2.837	13.177

The sample contains 56,267 firm–CPC subclass–year observations for 2007–2010 and 2013–2016. $Competition^{pre}$ is the mean of (1 − HHI) in the CPC subclass over 2007–2010. Variable definitions are given in Table A1

4 Empirical strategy

4.1 Specification

The main specification is a difference-in-differences regression with a continuous treatment:

$$\log(Apps_{ict}) = \alpha_i + \alpha_c + \alpha_t + \beta\left(Post_t \times Competition_c^{pre}\right) + \gamma' X_{it} + \varepsilon_{ict} \quad (3)$$

where $Apps_{ict}$ is the number of applications filed by firm i in CPC subclass c in year t, $Post_t$ equals one in 2013–2016 and zero in 2007–2010, α_i, α_c and α_t are firm, class and year fixed effects, and X_{it} contains the firm controls from Table 2. Standard errors are clustered two-way, by firm and by CPC subclass (Cameron et al. 2011). The coefficient β shows how much more patenting increased after the reform in classes with higher competition before the reform. I estimate the same specification with log granted patents as the outcome.

The panel is built from observed filings, so it contains only firm–class–year cells with at least one application. The estimates therefore describe the intensive margin: how much active filers changed their filing in classes where they were already present. Entry of new firms into a class is not part of the identifying variation. I come back to this point in Section 5.6.

I also report two extensions of Eq. (3). The first one replaces the separate firm and year effects with firm × year fixed effects. In this case, identification comes only from differences across classes within the same firm and year. In other words, it asks whether a firm increased its filing more in its competitive classes than in its other classes. Shocks that affect the whole

firm, such as a change in its patent budget, are absorbed. The second extension adds the transition years with a separate interaction $Transition_t \times Competition_c^{pre}$, where $Transition_t$ equals one in 2011–2012. To follow the dynamics, I also estimate an event-study version with year-specific interactions:

$$\log(Apps_{ict}) = \alpha_i + \alpha_c + \alpha_t + \sum_{k \neq 2009} \beta_k \left(\mathbb{1}[t = k] \times Competition_c^{pre}\right) + \gamma' X_{it} + \varepsilon_{ict} \quad (4)$$

on the full 2007–2016 sample, with 2009 as the omitted year.

4.2 Identification

The identifying assumption is that, without the reform, patenting in more and less competitive classes would have followed parallel paths. With a continuous treatment, β compares the response to the reform across classes with different levels of competition. Callaway et al. (2024) show that a causal reading of such comparisons needs a stronger version of parallel trends. So β is best read as the average difference in the post-reform change per unit of pre-reform competition. Several features of the design support the assumption.

First, $Competition_c^{pre}$ is measured before the reform and does not change over time, so post-reform filing cannot affect it. Second, the transition years are excluded from the baseline comparison, which reduces the risk that anticipation affects the pre-period. Third, the event study shows no differential trends in the earliest pre-reform years. The interaction coefficients for 2007 and 2008, relative to 2009, are jointly insignificant ($F(2,318) = 0.30$, $p = 0.740$). The formal test uses only 2007 and 2008, because the coefficient starts to rise in 2010, when the direction of the reform became clear. The pre-period is short, so this test has limited power. For this reason I also check that the result holds when the competition measure is built only from 2007–2009 and the years 2010–2012 are excluded (Section 5.6).

Another concern is that other provisions of the AIA, especially the new post-grant review proceedings, may affect the estimates. These provisions change the value of a patent after it is granted. There is no clear reason why they should cause a larger filing response in classes with higher pre-reform competition, which is the variation I use. The FITF rule, in contrast, acts directly at the filing stage and most strongly in such classes. I therefore read β as the differential response of more competitive classes to the AIA, with the race to file as the most likely channel, and not as an isolated effect of the FITF provision alone.

5 Results

5.1 Patent applications

Table 3 shows the estimates of Eq. (3) for patent applications. Column (1) is a simple interaction without controls or fixed effects, and the next columns add firm controls and fixed effects. The coefficient is positive and significant at the 1% level in every specification. It grows from 0.361 in column (1) to 0.679 in my preferred specification in column (5), which includes firm, class and year fixed effects and firm controls. This suggests that the result is not driven by fixed differences between firms or by how firms are spread across classes.

The effect is economically important. Moving from the 25th to the 75th percentile of $Competition^{pre}$ (from 0.464 to 0.753) is associated with an additional increase in applications of $\exp(0.679 \times 0.289) - 1 \approx 22\%$ after the reform, among firm–class pairs with active filing.

Column (6) replaces the separate firm and year effects with firm × year fixed effects. The coefficient is 0.670, almost the same as in the baseline. So, holding fixed everything that happens to a firm in a given year, the firm still increased its filing more in its more competitive classes. Firm-level shocks, such as changes in profits, financing or the total patent budget, cannot easily explain this result.

Column (7) adds the transition years 2011–2012 with a separate interaction. The transition coefficient is 0.631 and the post-reform coefficient is 0.676, and both are significant at the 1% level. So firms increased their filing in competitive classes already before the FITF rule took effect, which is what forward-looking firms would do.

Table 3 Post-AIA change in patent applications by pre-reform competition

Dependent variable: log(Patent applications)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post × $Competition^{pre}$	0.361***	0.443***	0.458***	0.638***	0.679***	0.670***	0.676***
	(0.092)	(0.103)	(0.096)	(0.147)	(0.136)	(0.103)	(0.131)
Transition × $Competition^{pre}$							0.631***
							(0.121)
Firm controls	No	Yes	Yes	Yes	Yes	No	Yes
Firm FE	No	No	No	Yes	Yes	No	Yes
Class FE	No	No	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	No	Yes	No	Yes
Firm × year FE	No	No	No	No	No	Yes	No
Observations	56,267	56,267	56,267	56,267	56,267	55,510	67,119
Within R^2	0.018	0.116	0.110	0.022	0.001	0.001	0.001

The unit of observation is a firm–CPC subclass–year. Columns (1)–(6) use 2007–2010 and 2013–2016; column (7) adds 2011–2012 and estimates a separate transition interaction, with 2007–2010 as the base period. The mean of the dependent variable is 0.928. Firm controls are log assets, R&D intensity, capex intensity, leverage, cash ratio, return on assets and log market capitalisation; their coefficients are omitted for brevity. Standard errors, clustered by firm and by CPC subclass, are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Fig. 2 shows the raw data behind these estimates. It plots mean log applications for the top and bottom quintiles of $Competition^{pre}$. In 2007–2010 both groups grow, and the gap between them widens only slowly, from about 0.16 to 0.29. Both groups rise in 2011, but the high-competition group rises further in 2012 and then stays at about 1.1, while the low-competition group falls back. After 2012 the gap is about 0.42–0.52. It is largest in 2013 and becomes slightly smaller by 2016.

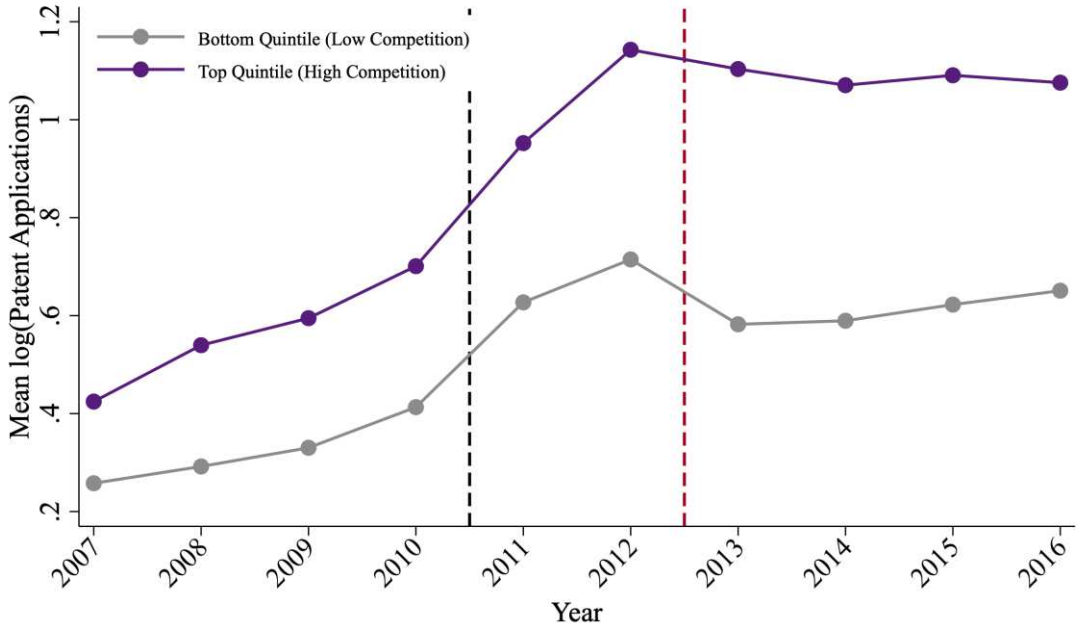
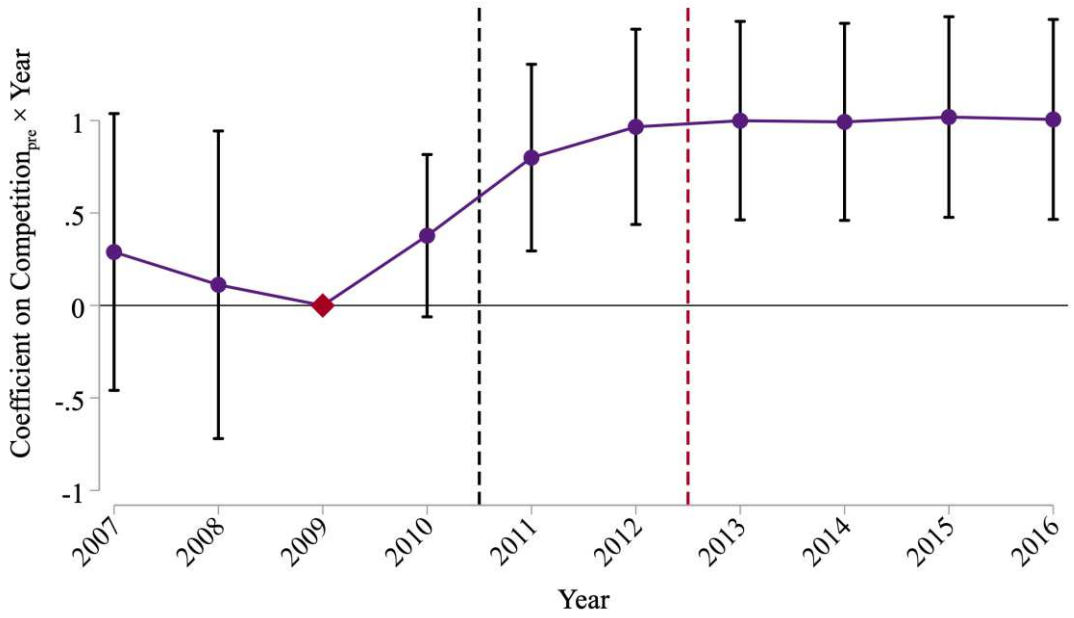


Fig. 2 Mean log patent applications for firm–CPC subclass pairs in the top and bottom quintiles of pre-reform competition, 2007–2016. The dashed lines separate the pre-reform period (2007–2010), the transition period after the AIA was signed (2011–2012) and the period under the first-inventor-to-file rule (2013–2016)

Fig. 3 shows the event study from Eq. (4). The coefficients for 2007 and 2008 are close to zero and not significant. The coefficient starts to rise in 2010, grows further in 2011 and 2012, and stays at about 1.0 in 2013–2016. The early rise fits the legislative timeline: by 2010 the direction of the reform was clear, and firms that expected a filing-based rule had reasons to file earlier. At the same time, the rise in 2010 means that the pre-reform period is not completely clean. This is why I also run a robustness check that drops 2010 (Section 5.6).



328

329 **Fig. 3** Event study: year-specific coefficients on $\mathbb{1}[t = k] \times Competition^{pre}$ for log patent
330 applications, Eq. (4), with 2009 as the base year. Bars show 95% confidence intervals based
331 on standard errors clustered by firm and CPC subclass. Dashed lines as in Fig. 2

332 *5.2 Granted patents*

333 More applications do not automatically mean more patents. If firms filed weaker
334 applications just to get early priority dates, many of them could fail examination. So Table 4
335 repeats the analysis with log granted patents, dated by filing year. The baseline coefficient is
336 0.610. It is significant at the 1% level and close to the application estimate of 0.679. With firm
337 × year fixed effects it is 0.710. The transition coefficient is 0.402, smaller than for
338 applications, which fits the lag between filing and grant. Because the magnitudes are similar,
339 the extra filing in competitive classes does not seem to consist mainly of applications that
340 failed examination.

Table 4 Post-AIA change in granted patents by pre-reform competition

Dependent variable: log(Granted patents)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post × $Competition^{pre}$	0.269***	0.334***	0.312***	0.635***	0.610***	0.710***	0.593***
	(0.085)	(0.087)	(0.087)	(0.100)	(0.099)	(0.105)	(0.100)
Transition × $Competition^{pre}$							0.402***
							(0.088)
Firm controls	No	Yes	Yes	Yes	Yes	No	Yes
Firm FE	No	No	No	Yes	Yes	No	Yes
Class FE	No	No	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	No	Yes	No	Yes
Firm × year FE	No	No	No	No	No	Yes	No
Observations	90,072	90,072	90,072	90,072	90,072	89,290	115,833
Within R^2	0.035	0.119	0.104	0.049	0.002	0.002	0.002

The unit of observation is a firm–CPC subclass–year; granted patents are dated by application filing year. The mean of the dependent variable is 0.817. Sample periods, controls and clustering as in Table 3. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

5.3 Heterogeneity by technology field

If the response comes from a race to file, it should be strongest in fields where simultaneous invention is common. Table 5 estimates Eq. (3) separately for each CPC section. For granted patents, the coefficient is largest in Physics (section G, 0.948) and Chemistry and Metallurgy (C, 0.838). It is also significant at the 1% level in Performing Operations and Transporting (B, 0.573), and positive but only marginally significant in Electricity (H, 0.753). It is small and not significant in Human Necessities (A), Textiles and Paper (D), Fixed Constructions (E) and Mechanical Engineering (F). The application estimates show the same pattern with one difference. In Electricity, the application coefficient (2.272) is much larger than the grant coefficient. So in electricity, which includes telecommunications and electronics, grants did not fully keep up with the filing response. Possible reasons are examination lags, higher rejection rates or more continuation filings. Section Y groups unrelated cross-field technologies, gives an uninformative estimate, and I do not interpret it.

Table 5 Post-AIA effect by CPC section

CPC section	Applications	Obs.	Granted patents	Obs.
A Human necessities	0.396 (0.259)	6,162	0.218 (0.310)	8,711
B Performing operations; transporting	0.810*** (0.227)	9,096	0.573*** (0.137)	16,031
C Chemistry; metallurgy	0.854*** (0.195)	8,184	0.838*** (0.181)	12,178
D Textiles; paper	0.319 (0.356)	522	0.388 (0.380)	1,204
E Fixed constructions	0.289 (0.517)	811	0.267 (0.267)	1,797
F Mechanical engineering; lighting; heating; weapons	0.019 (0.221)	5,258	0.289 (0.185)	9,182
G Physics	0.879** (0.388)	13,672	0.948*** (0.208)	19,384
H Electricity	2.272*** (0.460)	11,251	0.753* (0.386)	15,820
Y Cross-field technologies	–	–	0.168 (1.349)	5,765

Each cell reports the coefficient on Post × $Competition^{pre}$ from Eq. (3), estimated on the observations in one CPC section, with standard errors clustered by firm and CPC subclass in parentheses. All regressions include firm, CPC subclass and year fixed effects and firm controls. Section Y has too few observations in the application sample. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Fig. 4 compares the section estimates for granted patents with the average value of $Competition^{pre}$ across each section's observations in the estimation sample. Sections with more competition tend to have larger responses, and the technology-intensive sections C, G and H are in the upper right part of the figure. The figure uses only eight sections and is descriptive, but it agrees with the race-to-file mechanism. Table A2 in the Appendix gives the unweighted distribution of pre-reform concentration across subclasses in each section. It ranks the sections differently, because it gives every subclass the same weight.

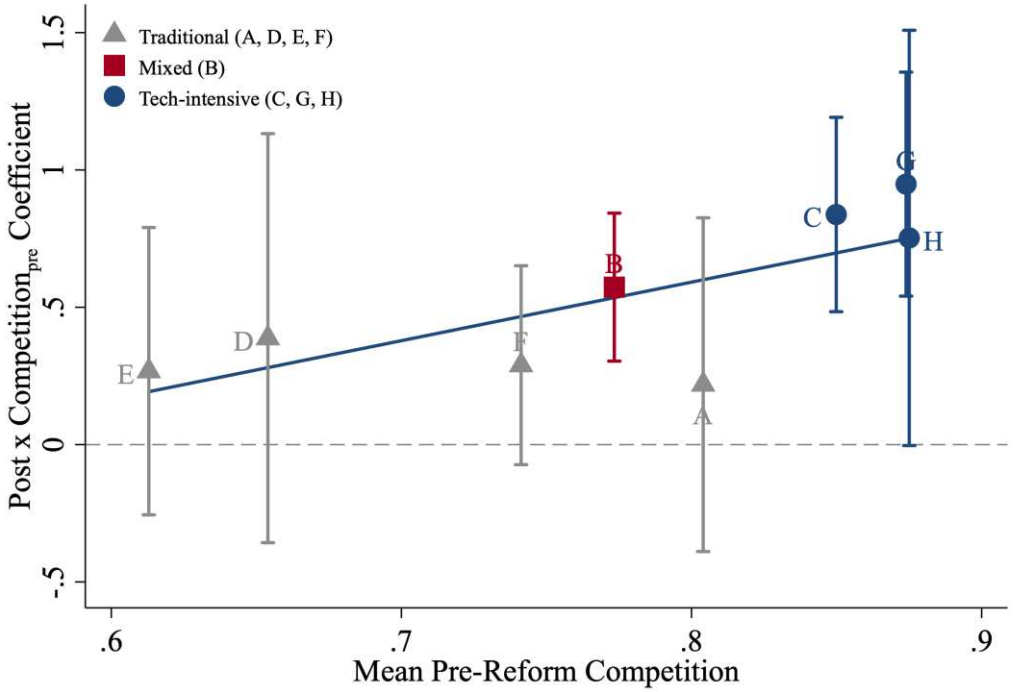


Fig. 4 Section-level estimates of the post-AIA effect on granted patents (Table 5) plotted against the average $Competition^{pre}$ of each section's observations. Bars show 95% confidence intervals; the line is a linear fit. Section Y is excluded

5.4 Has innovation increased?

More patents can mean more invention, but they can also mean more strategic filing of the same inventions. Table 6 therefore looks at outcomes that are closer to the value of patents: the number of claims, which measures patent scope, and forward citations, which measure technological impact. For each outcome, it reports the preferred specification and the firm × year specification.

Table 6 Patent scope and citation outcomes

Outcome	Granted: baseline	Granted: firm × year FE	Applications: baseline	Applications: firm × year FE
log(Count)	0.610*** (0.099)	0.710*** (0.105)	0.679*** (0.136)	0.670*** (0.103)
log(Average claims + 1)	−0.026 (0.036)	−0.028 (0.036)	–	–
log(Total claims + 1)	0.581*** (0.101)	0.678 (–)	−0.092 (0.199)	−0.022 (0.181)
log(Average forward citations + 1)	0.135* (0.070)	0.148** (0.067)	−0.007 (0.103)	0.130 (0.104)
log(Total forward citations + 1)	0.680*** (0.125)	0.762*** (0.133)	0.020 (0.189)	0.159 (0.193)
log(Top-10% cited patents + 1)	0.532*** (0.058)	0.553*** (0.052)	0.075 (0.056)	0.099 (–)

Each cell reports the coefficient on Post × $Competition^{pre}$. Baseline columns include firm, CPC subclass and year fixed effects and firm controls; firm × year columns include firm × year and CPC subclass fixed effects. Top-10% cited patents are those at or above the 90th percentile of forward citations within the same CPC subclass and filing year. Claims are available only for granted patents; application-linked claims and citations are defined for applications that were later granted. A dash in parentheses means that the two-way clustered standard error could not be computed because the variance matrix is singular. Standard errors clustered by firm and CPC subclass. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

These results give only limited support for a broad improvement in quality. Average claims per granted patent do not change ($\hat{\beta} = -0.026$), so individual patents did not become broader. Average citations per granted patent increase only a little ($\hat{\beta} = 0.135$, $p < 0.10$; 0.148, $p < 0.05$, with firm × year fixed effects). The number of highly cited patents rises strongly ($\hat{\beta} = 0.532$), which does not fit a story in which firms only filed low-value patents. However, total claims, total citations and top-cited counts all rise automatically when more patents are granted, so they mix quantity and quality. Citation measures for later cohorts are also truncated, because these patents have had less time to be cited. For application-linked outcomes, none of the claim or citation measures changes significantly. Overall, the reform increased the number of patents, including some important ones, but did not clearly improve average patent quality.

Table 7 looks at firms' technological breadth and research investment at the firm–year level. Here the treatment is a firm's exposure to competition: the average of $Competition^{pre}$ over the classes in which the firm filed before the reform, weighted by its pre-reform filing shares:

$$Exposure_i = \sum_c \frac{Apps_{ic}^{pre}}{Apps_i^{pre}} Competition_c^{pre} \quad (5)$$

Firms with higher exposure became less likely to enter new CPC subclasses after the reform ($\hat{\beta} = -0.394$) and reduced their R&D spending ($\hat{\beta} = -0.609$). Both effects are significant at the 1% level. The share of new classes is zero by construction in the pre-reform period, so column (1) compares post-reform entry across firms with different exposure and should be read as descriptive. Together, the two results do not support the idea that the reform increased research effort or pushed firms into new technology areas. Instead, more exposed firms focused on the classes where they were already active.

Table 7 Technological scope and R&D investment (firm–year level)

	(1) Share of new CPC subclasses	(2) log(R&D expenditure)
Post × Exposure	−0.394***	−0.609***
	(0.083)	(0.195)
Firm FE	Yes	Yes
Year FE	Yes	Yes
Observations	6,572	7,132
Within R^2	0.021	0.343

The unit of observation is a firm–year. Exposure is the firm's pre-reform filing-share-weighted average of $Competition^{pre}$, Eq. (5). In column (1) the outcome is the share of the firm's active CPC subclasses in the post-reform period that were not active in the pre-reform period (zero in the pre-reform period). In column (2) the outcome is log R&D expenditure from Compustat. Standard errors clustered by firm. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

5.5 Continuation filings

Firms can secure an early priority date and later extend or change their protection with continuation, divisional and continuation-in-part filings. These strategies become more useful when filing speed matters. When I exclude such filings from the granted-patent sample, the coefficient falls from 0.610 to 0.242 but stays significant at the 5% level (Table 8, row 6). So continuation-type filings explain a large part of the grant response, but not all of it. This agrees with the strategic portfolio building described by Hall and Ziedonis (2001): firms use patents not only to protect separate inventions, but also to keep their options open and defend their technology space against close rivals.

5.6 Robustness

Table 8 reports robustness checks for both outcomes. Row (2) builds $Competition^{pre}$ only from 2007–2009 and drops 2010–2012 from the sample, so the early rise in 2010 cannot

affect either the treatment or the pre-period. The coefficients stay positive and significant, and the application coefficient becomes larger (1.113), which fits the idea that 2010 already contains some anticipation. Row (3) uses the standardised number of distinct filers in a class before the reform instead of the HHI-based measure; classes with more filers also show a stronger response. Row (4) keeps only firm–class pairs that are active in both periods, so the result cannot come from changes in which firm–class pairs are observed. Row (5) uses a leave-one-out version of the competition measure that excludes the firm's own filings. This removes the mechanical link between a firm's own activity and its measured competition. The coefficient stays positive and significant in every check.

Table 8 Robustness checks

Specification	Applications	Obs.	Granted patents	Obs.
(1) Baseline	0.679*** (0.136)	56,267	0.610*** (0.099)	90,072
(2) $Competition^{pre}$ from 2007–2009; 2010–2012 excluded	1.113*** (0.260)	33,412	0.684*** (0.126)	82,594
(3) Number of pre-reform filers (standardised)	0.129*** (0.041)	56,267	0.066*** (0.018)	90,072
(4) Firm–class pairs active in both periods	0.885*** (0.237)	4,903	0.579*** (0.154)	40,111
(5) Leave-one-out competition	0.787*** (0.189)	5,124	0.405*** (0.096)	44,078
(6) Excluding continuation, divisional and CIP filings	–	–	0.242** (0.096)	90,072

Each cell reports the coefficient on Post × treatment, where the treatment is $Competition^{pre}$ except in row (3). All regressions include firm, CPC subclass and year fixed effects and firm controls. Row (6) excludes continuation, divisional and continuation-in-part (CIP) filings identified from PatentsView related-document records. Standard errors clustered by firm and CPC subclass in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

These checks also show the main limitations of the design. The estimates describe the intensive margin, because the panel contains only cells with positive filing. If the reform made firms start filing in competitive classes where they were not active before, the baseline would understate the total response. The balanced-pair and leave-one-out samples are much smaller than the baseline sample, especially for applications, so these checks are less precise and use a narrower set of firm–class pairs. Finally, the outcome is the logarithm of positive counts. A count-data model estimated on a panel with zeros, such as Poisson pseudo-maximum likelihood (Santos Silva and Tenreyro 2006), would capture both margins and is a natural next step.

6 Discussion

6.1 Interpreting the response

Overall, the results fit a race-to-file interpretation. The response is strongest where the FITF rule should matter most: in classes with many active rivals and in technology-intensive fields such as physics and chemistry. This is in line with Thompson and Kuhn (2020), who find that patent races are concentrated in information-technology fields. The response starts before the rule took effect, with first signs in 2010 when the direction of the reform became clear, and it becomes stronger after the AIA was signed. This is what we would expect if firms adjusted their filing in advance. It also holds within the same firm and year, so it is not explained by firm-wide changes in patenting.

The other results suggest that the response mainly consists of faster and more intensive filing in areas where firms were already active. Part of the grant response comes from continuation-type filings, which let firms secure priority early and extend protection later. Average patent quality does not clearly improve. More exposed firms enter fewer new classes and spend less on R&D. Using the exploration–exploitation distinction of March (1991), the reform moved firms towards exploiting the technology positions they already had. For a single firm this response makes sense. After the reform, filing quickly in familiar classes gives faster and more certain returns than building a presence in new fields, where the firm has no prior art, no patent portfolio and no prosecution experience.

The fall in R&D should be read with care. It is a firm-level result without controls for other changes that may have affected more exposed firms, and R&D spending is measured for the whole firm, not by technology class. Still, it is the opposite of what we would expect if the reform had increased the return to research in competitive fields. A more likely explanation is that firms moved resources from research towards securing and defending patent positions.

6.2 Implications for competition and patent policy

The welfare effect of the FITF rule is not clear, but the results show which margins matter. On the positive side, the rule probably simplified priority disputes, removed costly interference proceedings and may have encouraged earlier disclosure of inventions. The rise in highly cited patents is consistent with valuable inventions entering the patent system sooner. On the negative side, theory predicts that patent races can waste resources, because firms spend them on being first rather than on new knowledge (Loury 1979; Lee and Wilde 1980). The concentration of the response in contested classes, the role of continuation filings,

the lack of a clear improvement in average quality and the fall in R&D are all consistent with some waste of this kind.

For competition policy, the main lesson is that the procedural rules of the patent system affect how firms compete in technology markets. A filing-based priority rule changes the form of rivalry in contested fields: firms compete more on the speed and volume of filing and on building patent portfolios in the areas where they are already present. Earlier work on first-to-file reforms found that such rules shift patent rights towards large organisations with more legal resources (Lo and Sutthiphisal 2009; Abrams and Wagner 2013), and recent work finds that the AIA affected firms of different size differently (Chakravarty 2023; Driver 2025). My results add that the reform increased filing most where technological rivals were closest. So competition authorities that look at patent thickets, portfolio strategies or the market power of incumbents in technology markets should take into account how patent procedures shape these strategies. Patent offices, in turn, should look at how procedural reforms affect the intensity and form of technological competition, and not only the number of patents.

6.3 Limitations

This study has several limitations. First, the estimates describe the intensive margin of filing by publicly traded firms that were already active in a class. The main panel does not capture new firm–class pairs or the response of small firms and individual inventors. Second, the competition measure is based on publicly traded firms. It ranks classes in a similar way to a measure based on all applicants, but it overstates concentration in classes with a single DISCERN filer. Third, the event study shows anticipation from 2010, so the clean pre-period is short. Fourth, the AIA changed several rules at the same time. The variation across classes that I use is most closely linked to the FITF rule, but other provisions cannot be fully ruled out. Fifth, citation measures are affected by truncation, and the outcomes are logarithms of positive counts rather than count models estimated with zeros. Finally, the R&D and technological-scope results are estimated at the firm level with simple specifications and should be seen as suggestive.

7 Conclusion

In this paper I study how the move from first-to-invent to first-inventor-to-file under the America Invents Act changed the way firms patent. I use pre-reform patenting competition across CPC subclasses as a measure of exposure to the new incentive to file first. I find that patent applications and granted patents increased more after the reform in more competitive

technology classes. The response is economically important, starts before the rule took effect, holds within the same firm and year, and is concentrated in technology-intensive fields.

However, the extra patenting does not look like a broad increase in innovation. Average patent quality improves little, part of the response comes from continuation-type filings, and more exposed firms enter fewer new technology classes and cut R&D spending. So the reform changed the timing and organisation of patenting more clearly than the amount of invention. Filing-based priority rules can make priority simpler and disclosure faster, but they can also make close technological rivals race harder to file. Future research could study the extensive margin, including entry by new and small firms, and use count-data models and longer post-reform periods to measure the effect on follow-on innovation.

Statements and Declarations

Competing interests The author has no competing interests to declare that are relevant to the content of this article.

Funding The author did not receive support from any organization for the submitted work.

Data availability The patent data are from DISCERN 2.0 (Arora et al. 2024) and PatentsView (https://patentsview.org), which are publicly available. Financial data are from Compustat (S&P Global Market Intelligence), available by subscription and not shareable under its licence terms. The code used to construct the sample and estimate the results is available from the author on reasonable request.

Ethics approval Not applicable.

Author contributions A.B. is the sole author and was responsible for the study design, data collection and analysis, and writing of the manuscript.

Use of generative AI Generative AI tools (ChatGPT and Claude) were used to help format tables and equations, to proofread and edit the language, and to help identify coding errors. The author reviewed all output and is responsible for the content.

References

Abrams DS, Wagner RP (2013) Poisoning the next Apple? The America Invents Act and individual inventors. Stanford Law Rev 65(3):517–564

Aghion P, Bloom N, Blundell R, Griffith R, Howitt P (2005) Competition and innovation: an inverted-U relationship. Q J Econ 120(2):701–728. https://doi.org/10.1093/qje/120.2.701

548 Arora A, Belenzon S, Cioaca L, Sheer L, Shin HM, Shvadron D (2024) DISCERN 2: Duke
549 Innovation & SCientific Enterprises Research Network, version 2.0.1 [dataset]. Zenodo.
550 https://doi.org/10.5281/zenodo.13619821
551 Arrow KJ (1962) Economic welfare and the allocation of resources for invention. In:
552 Universities-National Bureau Committee for Economic Research (ed) The rate and
553 direction of inventive activity: economic and social factors. Princeton University Press,
554 Princeton, pp 609–626. https://doi.org/10.1515/9781400879762-024
555 Bloom N, Schankerman M, Van Reenen J (2013) Identifying technology spillovers and
556 product market rivalry. Econometrica 81(4):1347–1393.
557 https://doi.org/10.3982/ECTA9466
558 Callaway B, Goodman-Bacon A, Sant'Anna PHC (2024) Difference-in-differences with a
559 continuous treatment. NBER Working Paper 32117. https://doi.org/10.3386/w32117
560 Cameron AC, Gelbach JB, Miller DL (2011) Robust inference with multiway clustering. J
561 Bus Econ Stat 29(2):238–249. https://doi.org/10.1198/jbes.2010.07136
562 Chakravarty S (2023) The America Invents Act and innovation by small entities. Job market
563 paper, Department of Economics, University of North Carolina at Greensboro.
564 https://satyaki4.github.io/files/JMP_satyaki.pdf
565 Dohse D, Goel RK, Saunoris JW (2023) Patenting uncertainty and its impact on innovation:
566 evidence from the United States. J Technol Transf 48(5):1839–1859.
567 https://doi.org/10.1007/s10961-023-10014-0
568 Driver J (2025) Property rights, firm size and investments in innovation: evidence from the
569 America Invents Act. Center for Economic Studies Working Paper CES-WP-25-31, US
570 Census Bureau, Washington, DC
571 Hall BH, Ziedonis RH (2001) The patent paradox revisited: an empirical study of patenting in
572 the U.S. semiconductor industry, 1979–1995. RAND J Econ 32(1):101–128.
573 https://doi.org/10.2307/2696400
574 Hegde D, Herkenhoff K, Zhu C (2023) Patent publication and innovation. J Polit Econ
575 131(7):1845–1903. https://doi.org/10.1086/723636
576 Helmers C, Love BJ (2024) Patent law reform and innovation: an empirical assessment of the
577 last 20 years. Int Rev Law Econ 79:106210. https://doi.org/10.1016/j.irle.2024.106210
578 Huang R, Li L, Lu LY, Wu H (2021) The impact of the Leahy-Smith America Invents Act on
579 firms' R&D disclosure. Eur Account Rev 30(5):1067–1104.
580 https://doi.org/10.1080/09638180.2020.1806896

581 Lee T, Wilde LL (1980) Market structure and innovation: a reformulation. Q J Econ
582 94(2):429–436. https://doi.org/10.2307/1884551
583 Lemley MA (2012) The myth of the sole inventor. Mich Law Rev 110(5):709–760
584 Lemley MA, Chien CV (2003) Are the U.S. patent priority rules really necessary? Hastings
585 Law J 54(5):1299–1334
586 Lerner J, Speen A, Leamon A (2015) The Leahy-Smith America Invents Act: a preliminary
587 examination of its impact on small businesses. Report RS429, Office of Advocacy, US
588 Small Business Administration, Washington, DC
589 Lo S-t, Sutthiphisal D (2009) Does it matter who has the right to patent: first-to-invent or
590 first-to-file? Lessons from Canada. NBER Working Paper 14926.
591 https://doi.org/10.3386/w14926
592 Loury GC (1979) Market structure and innovation. Q J Econ 93(3):395–410.
593 https://doi.org/10.2307/1883165
594 March JG (1991) Exploration and exploitation in organizational learning. Organ Sci 2(1):71–
595 87. https://doi.org/10.1287/orsc.2.1.71
596 Moser P, Voena A (2012) Compulsory licensing: evidence from the Trading with the Enemy
597 Act. Am Econ Rev 102(1):396–427. https://doi.org/10.1257/aer.102.1.396
598 Poege F (2022) Competition and innovation: the breakup of IG Farben. IZA Discussion Paper
599 15517, Institute of Labor Economics (IZA), Bonn
600 Santos Silva JMC, Tenreyro S (2006) The log of gravity. Rev Econ Stat 88(4):641–658.
601 https://doi.org/10.1162/rest.88.4.641
602 Thompson NC, Kuhn JM (2020) Does winning a patent race lead to more follow-on
603 innovation? J Leg Anal 12:183–220. https://doi.org/10.1093/jla/laaa001
604 Valentine K, Zhang JL, Zheng Y (2025) Strategic scientific disclosure: evidence from the
605 Leahy–Smith America Invents Act. J Account Res 63(4):1723–1755.
606 https://doi.org/10.1111/1475-679X.12605
607 Watzinger M, Schnitzer M (2022) The breakup of the Bell System and its impact on US
608 innovation. CEPR Discussion Paper 17635, Centre for Economic Policy Research,
609 London

Appendix

Table A1 Variable definitions

Variable	Definition	Source
log(Patent applications)	Log number of pre-grant publications filed by firm *i* in CPC subclass *c* in year *t*; cells with zero filings are not in the panel	DISCERN 2.0; PatentsView
log(Granted patents)	Log number of granted patents of firm *i* in subclass *c*, dated by application filing year; cells with zero patents are not in the panel	DISCERN 2.0; PatentsView
log(Average claims + 1)	Log of mean claims per granted patent in the cell, plus one	PatentsView
log(Total claims + 1)	Log of total claims of granted patents in the cell, plus one	PatentsView
log(Average forward citations + 1)	Log of mean forward citations per granted patent in the cell, plus one; missing citations set to zero	PatentsView
log(Total forward citations + 1)	Log of total forward citations of granted patents in the cell, plus one	PatentsView
log(Top-10% cited + 1)	Log number of granted patents at or above the 90th percentile of forward citations in the same subclass and filing year, plus one	PatentsView
Share of new CPC subclasses	Share of the firm's post-reform active subclasses that were not active before the reform; zero in the pre-reform period	DISCERN 2.0; PatentsView
log(R&D expenditure)	Log R&D expenditure (xrd); non-positive or missing values excluded	Compustat
$Competition^{pre}$	Mean of (1 − HHI) of application shares in the subclass over 2007–2010, Eqs. (1)–(2)	DISCERN 2.0; PatentsView
Leave-one-out competition	$Competition^{pre}$ computed without firm *i*'s own applications	DISCERN 2.0; PatentsView
Number of pre-reform filers	Standardised mean number of distinct DISCERN firms filing in the subclass per year, 2007–2010	DISCERN 2.0
Exposure	Pre-reform filing-share-weighted average of $Competition^{pre}$ across the firm's subclasses, Eq. (5)	DISCERN 2.0; PatentsView
Post; Transition	Post = 1 in 2013–2016, 0 in 2007–2010; Transition = 1 in 2011–2012	–
log(Assets)	Log total assets (at)	Compustat
R&D intensity	R&D expenditure (xrd) / total assets (at)	Compustat
Capex intensity	Capital expenditure (capx) / total assets	Compustat
Leverage	(Short-term debt (dlc) + long-term debt (dltt)) / total assets	Compustat
Cash ratio	Cash and short-term investments (che) / total assets	Compustat
Return on assets	Net income (ni) / total assets	Compustat
log(Market capitalisation)	Log market value of equity (mkvalt)	Compustat

Financial variables are winsorised at the 1st and 99th percentiles

Table A2 Pre-reform concentration (HHI) of CPC subclasses by section, 2007–2010

Section	Mean	Median	SD	Min	Max	Subclasses
A Human necessities	0.507	0.417	0.314	0.027	1.000	77
B Performing operations; transporting	0.517	0.444	0.336	0.043	1.000	151
C Chemistry; metallurgy	0.411	0.256	0.334	0.047	1.000	79
D Textiles; paper	0.539	0.468	0.309	0.177	1.000	26
E Fixed constructions	0.541	0.426	0.295	0.177	1.000	30
F Mechanical engineering; lighting; heating; weapons	0.425	0.353	0.257	0.070	1.000	94
G Physics	0.400	0.250	0.350	0.031	1.000	77
H Electricity	0.301	0.168	0.296	0.047	1.000	54
Y Cross-field technologies	0.282	0.141	0.336	0.021	1.000	13

The unit of observation is a CPC subclass observed in 2007–2010. Higher HHI means higher concentration and lower competition

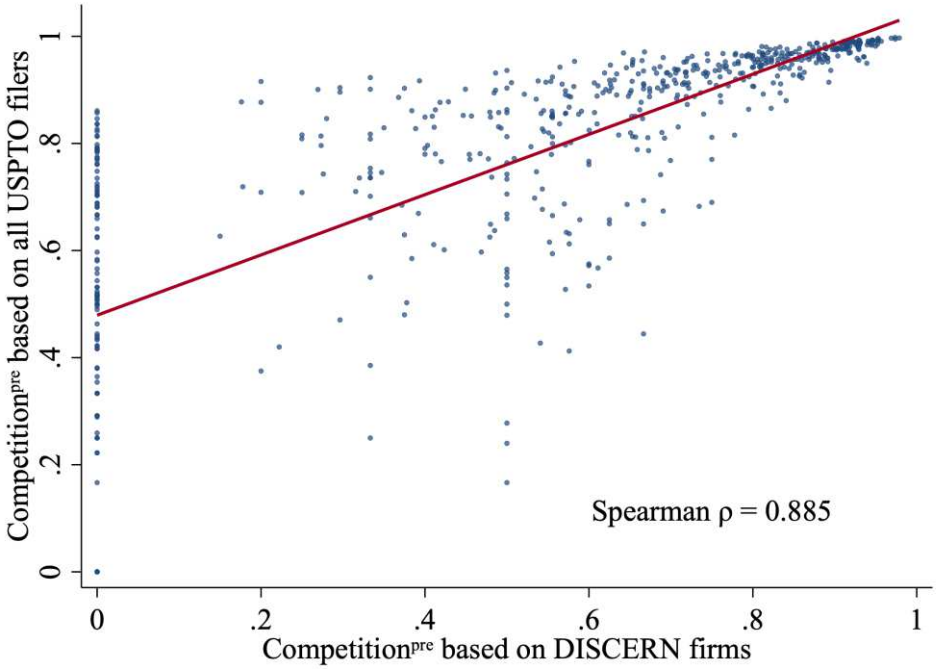


Fig. A1 Pre-reform competition by CPC subclass computed from DISCERN firms (horizontal axis) and from all USPTO assignees in PatentsView (vertical axis), 2007–2010. Each point is a subclass; the line is a linear fit. Spearman ρ = 0.885

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Fig1.tif
- Fig2.tif
- Fig3.tif
- Fig4.tif
- FigA1.tif