

The Importance of Being Steve: An Econometric Analysis of the Contribution of Steve Jobs's Patents to Apple's Market Valuation

Daniel K.N. Johnson¹ & Sylvie Scowcroft¹

¹ Economics & Business Department, Colorado College, Colorado Springs, Colorado Springs, USA

Correspondence: Daniel K.N. Johnson, Schlessman Professor of Economics & Business, Colorado College, Colorado Springs, 14 E Cache La Poudre Street, Colorado Springs, CO 80903, USA.

Received: January 26, 2016

Accepted: February 11, 2016

Online Published: March 8, 2016

doi:10.5430/ijfr.v7n2p30

URL: <http://dx.doi.org/10.5430/ijfr.v7n2p30>

Abstract

This paper evaluates the contribution of patent-related events to changes in stock prices, proposing that economics has traditionally failed to find much effect for two reasons which we identify and correct. First, patents vary widely in quality so we use quantile analysis and alternative measures of patent quality to identify effects. Second, we permit the possibility that information leaks out into investor sentiment during the long and uncertain time until patent grant, so evaluate the stock price effect at four different dates in the life of each patent. As a case study to test this approach, track all patents over a 27-year period for Apple Inc., permitting design patents to have different effects than traditional utility-model patents, and isolate the effect that Steve Jobs's name on a patent has at each stage of a patent's life.

Keywords: patent, Apple, valuation, design patent, event analysis

1. Introduction

Firms know that patents represent valuable assets, as evidenced by the expenditures of resources to develop the innovation to a patentable stage, to apply for and protect them, and to value their intellectual property portfolios at times of merger or acquisition. However, frequently those valuations are performed on an *ad hoc* basis, not only because of the complex interconnectedness of technological relationships between firms, but because economic analysis has failed to provide a method that *a priori* can quantify the value of a patent or IP portfolio. We propose that this is due in part to two fundamental reasons: a) patents vary widely in value, and b) have a long timeline to make it to market (including uncertainty about their ultimate grant status, multiple dates at which information might leak out to the market before a final patent is granted, and complicated final documents and may not be immediately absorbed into market reactions until technological and legal experts interpret their significance).

This paper aims to improve our modeling of the stock market valuation of patents, using a nuanced version of event analysis. We model the timing of four distinct stages in the disclosure of information about the potential value of a patent, two of them private and two of them public, two of them preliminary and two of them legally binding. We test this insight using patents granted to Apple, Inc. over a 27-year period for several provocative reasons: to quantify the importance of patent-related events to a clearly technology-based firm's stock value, to compare the relative importance of utility-model (or standard) patents and design (or ornamental) patents, and to investigate the role that founder and CEO Steve Jobs had on the firm. This last motivation ties into a larger literature on executive compensation (e.g. Adams et al. (2005)), but we aim here to quantify only Jobs's activities as a co-inventor on over 300 patents granted over the sample time period, as separate from his roles as visionary, entrepreneur, manager or evangelist.

2. Literature

This research depends on two strands of literature: the methodology known as event analysis and the empirical evidence on patent valuation, both of which will be briefly reviewed here. There is a long literature on event analysis in the finance discipline, and we follow that established methodology. Fama and French (1996) outline the original four-factor model which sets the expected stock value in the absence of market-relevant news, and abnormal returns are calculated as deviations from this expected value. Other studies (e.g., those reviewed in Hall, 2000; Hall et al.,

2007) have found that patents do not explain as much market value as anticipated, faring worse than research expenditures in that regard but adding explanatory power beyond those expenditures alone.

The literature on patent valuation spans industrial and commercial analysis (much of which is done case-by-case for litigation or ownership transfer purposes) but also has a vibrant home in economics. Johnson and Popp (2003) document the empirical correlation evidence of many factors with patent value, building on the previous work of Pakes and Simpson (1989) and Trajtenberg (1990), Lanjouw et al. (1996), and Lanjouw and Schankerman (1999). All of these are predicated on characteristics of the patent document itself that convey or imply value, such as subsequent citations (a measure of scholarly or scientific value), patent families (a measure of how broadly international applicants protect their rights, a measure of potential market value), or renewal rates (a measure of how often the applicant pays to retain patent rights, a measure of potential duration of market value). Direct measures of market value are rare, usually relying on auction prices (e.g., Sneed and Johnson, 2009) but point to other measures of patent value such as originality and generality which we will adopt here.

3. Data and Methodology

We consider all 9,861 market days between 1985 and 2011 as potential event dates, tracking 13,618 separate patent-related events for 3,481 distinct patent documents (2,903 utility-model or traditional patents and 578 design patents). Those patent events occurred on 2,457 distinct dates, with as many as 54 events on a single day. For the 265 events which occurred on days when stock markets were closed, or for which our defined event window would have occurred on closure days, we use the market closing price from the previous day (if the closure occurred before the patent event) or subsequent day (if the closure occurred on or after the patent event).

Using only our control group of observations, the 7,404 dates on which no patent-related event occurred, we regress each event window's observed rate of return on French's (2013) data for the four-factor model as proposed in Fama and French (1996):

$$R_t = \alpha_0 + \alpha_1 T_t + \alpha_2 M_t + \alpha_3 S_t + \alpha_4 B_t + e_t \quad (1)$$

where R_t is the rate of return on Apple stock;

T_t is the implicit rate of return on one-month Treasury bonds;

M_t is the rate of return on the market portfolio;

S_t is the difference between the rate of return on a portfolio of small-cap stocks and the rate of return on a portfolio of large-cap stocks;

B_t is the difference between the rate of return on a portfolio of high book-to-market value stocks and a portfolio of low book-to-market value stocks;

e_t is the residual.

The primary analysis of this control group uses a two-week rate of return, to compare with pre-event-to-post-event closing prices in our experimental group, but we also perform sensitivity analysis on week-long and two-day-long rates of return. The estimated coefficients of this model, White-corrected to address heteroskedasticity, are presented in Table 1 for the four time event windows we consider.

Notice that the simple four-factor model is highly satisfactory according to standard diagnostics (t-statistics and F-statistic), but explains only a small fraction of the variation in returns, a result which was largely anticipated as we model a single firm over time. Interestingly, the model performs best over the longest window (one week before to one week after), moderately well over the shortest window (one day before to one day after) and worst over intermediate windows.

For that reason, we will use the two-week, longest window as our primary result while testing the alternatives during our sensitivity tests.

Table 1. Estimated coefficients for four-factor model on Apple stock returns

Coefficient	Short and Quick (1 day before until 1 day after)	Short and Slow (1 day before until 7 days after)	Long and Quick (7 days before until 1 day after)	Long and Slow (7 days before until 7 days after)
α_0	0.314 (4.43) ^{***}	0.868 (5.71) ^{***}	1.087 (6.88) ^{***}	1.575 (9.54) ^{***}
α_1	0.705 (21.03) ^{***}	0.440 (12.50) ^{***}	0.227 (5.00) ^{***}	1.039 (35.83) ^{***}
α_2	-3.512 (2.30) ^{**}	-3.075 (2.61) ^{**}	-5.582 (4.50) ^{***}	-4.931 (6.02) ^{***}
α_3	-0.011 (0.16)	0.783 (11.14) ^{***}	-0.216 (2.80) ^{***}	0.220 (3.76) ^{***}
α_4	-0.483 (8.28) ^{***}	-0.184 (2.70) ^{***}	-0.230 (3.73) ^{***}	-0.646 (11.74) ^{***}
F-statistic	134.84 ^{***}	81.19 ^{***}	14.63 ^{***}	422.73 ^{***}
R-squared	0.139	0.064	0.014	0.197

Significance is indicated by ^{***} for 99%, ^{**} for 95%, ^{*} for 90%, and t-statistics are included in parentheses.

Based on these results, we calculate the expected rate of return \widehat{R}_t for every date in the sample, for each of the four event windows. Abnormal returns are therefore calculated as the standardized difference between actual and expected returns, $AR_t = [R_t - \widehat{R}_t]/s$ where s is the standard deviation of the expected rates of return in the control group. The Dickey-Fuller test rejects the null hypothesis of a unit root in this series well above the 99% level for each event window.

We consider four events in the life of a patent document--- application, publication, status acceptance and grant--- and only consider successful patents (i.e., those which we know to have been granted by the end of 2012). While this definition serves to eliminate some heterogeneity from our sample, by eliminating unsuccessful and abandoned applications from consideration, it opens our analysis to criticisms of survival bias. Further, investors do not know in advance which applications will be successful, so may react to applications not represented in our sample. However, details of unsuccessful patents were not public information for more than half of our sample period, as explained below, so we choose to omit them entirely from consideration.

When an inventor (or their employer) applies for a patent, their potential protection period begins at the date of application. However, the information recorded in the application is entirely private, as is the fact that an application has been made, unless the applicants elect to publicize the event or the contents of the document. Thus, one might expect that application events do not cause a significant stock market effect.

Prior to June 9 of 1995, the date on which the United States began a transition to conform to World Trade Organization standards for patent law, the information in a patent application would be kept secret until the date of grant: if a patent was never granted, the information was never public. Since that date patent applications have been published 18 months after their application date, regardless of when and even whether they are granted, a reduction in secrecy which led many inventors to file just in advance of the change in order to grandfather their patent applications under the old rules. This date of public revelation we call the publication date, representing the first moment at which technical details about the patentable information are known to be public knowledge.

After consideration by patent examiners, and an average of three to five years of communication to clarify and justify the legal claims in the application, a status acceptance letter is sent privately to the applicant, indicating that the legal claim to a patent will be granted. While the public announcement usually follows within two to three weeks, this stage of the process is still private, and may have little to no stock market impact.

Finally, in a very predictable manner after the close of business each Monday, the United States Patent and Trademark Office publishes the details of newly granted patents, presenting the first public announcement of the legally enforceable intellectual asset. In the wake of the creation of the Court of Appeals for the Federal Circuit in 1981, patent defendants have won roughly ninety percent of infringement cases brought to court, so this granting date represents a serious financial claim of ownership or at least of the legal right to exclude others from use and production (Kortum and Lerner, 1998).

When granted, each patent includes a set of required information, which we have found useful to analyze here. Every document includes the names of all applicants (inventors, not firms) so we easily identify patents which show Steve Jobs as a contributor. Every document shows the assignee, which may be the individual but frequently lists the company employing the inventors. In our case, we only consider patents listing Apple as the assignee. All dates related to the document are public knowledge at this point as well.

Since patents vary in value, we search for other methods of quantifying those variations as well. In the process of examination, each patent is assigned a primary technology code, or United States Patent Classification (USPC) code, so that patent examiners, lawyers and researchers may compare it to similar claims by other patent-holders. There are more than 450 such classes, divided into 150,000 subclasses. Further, each patent lists the technological claims made to ownership over intellectual space, in general listing more claims on a more broad or wide-reaching patent document. To justify those claims, patents must include a bibliography, other documents and patents of which the inventors were aware and which are clearly outside the scope of the current claims. We not only identify all citations listed in each Apple patent, but identify all subsequent patent documents which cite them in return. Using the USPC of each forward and backward citation, we reproduce and extend the work of Hall et al. (2001), calculating originality and generality measures for each Apple patent. Originality of each patent is calculated as

$$originality = 1 - \sum_{i=1}^n s_i^2 \quad (2)$$

or a measure of how broadly the references span across technologies, where s_i is the share of all references within a given patent that hail from a particular USPC class. Notice that the summation is a version of the Herfindahl Index of concentration, exponentially emphasizing classes with high shares of the total. Generality is measured symmetrically as

$$generality = 1 - \sum_{i=1}^n s_i^2 \quad (3)$$

where s_i here is the share of all subsequent citations within each USPC, or a measure of how broadly the given patent is cited, across technologies.

Since multiple patents may be granted on the same date, we average their characteristics in order to determine their aggregate effects (if any) on Apple stock prices. In many cases, different types of events occur on a given day, with some patent applications being filed as others are granted, so we summarize the characteristics of patent documents in each event-group (application, publication, status, issue) separately.

Even before a formal multivariate regression test, there is mixed evidence that patent events are correlated with abnormal returns. Table 2 compares the subsamples of 2-week event windows that comprise the lowest abnormal returns and highest abnormal returns over the 27-year period. Notice that with only one exception (design patent events), the highest quintile subsample quite dramatically outperforms the lowest quintile subsample on patent-related events. While none of these differences are statistically significant, they are suggestive. The same basic pattern holds for other event windows, but somewhat less strongly.

Table 2. Variable means by quintile

Variable	Average in lowest quintile (198 obs)	Average in highest quintile (198 obs)
Abnormal return (AR_t)	-7.89	7.63
Steve Jobs patent events	0.04	0.08
Design patent events	0.15	0.11
Patent events	1.18	1.48
Patent claims	10.89	15.60
Patent citations	13.81	19.85
Patent generality	0.23	0.29
Patent originality	0.20	0.28
Number of inventors	1.53	1.94

Now we test the ability of patent-related events to explain AR_t in the standard fashion, permitting each estimated coefficient β to vary by event type (application, publication, status or issue), and so grouping all patents which share a similar event on a given date:

$$AR_t = \beta_0 + \beta_1 Jobs_t + \beta_2 designs_t + \beta_3 patentvalue_t + \delta NASDAQ_t + u \quad (4)$$

where $Jobs$ is an indicator of how many patents events show Steve Jobs name on that date;

$designs$ is the number of design patent events on that date;

$patentvalue$ is a measure of the patent events on that date, either a count of patents, or the average number of subsequent citations, or the average number of claims, or the average originality or generality score for patents, or the average number of inventors listed on patents;

and $NASDAQ$ is the rate of return of the NASDAQ composite index over the same period.

While design patents are only represented in one way in the analysis, utility-model patents are represented in alternative methods (which are strongly positively correlated, so are not included simultaneously). Clearly, if patent events signal an asset's acquisition, then the estimated coefficients of patents and designs should be positive (at least for public event types). We include the NASDAQ composite index to capture other market events which might confound our identified events but which have not already been netted out by the Fama-French method.

Obviously, it would be optimal to identify and code all other significant events that occurred in the lifetime of Apple, Inc. over the observation period in order to avoid misattribution of effects. However, that task is a daunting one, and may be attempted in a subsequent study. We trust that our sample size will reduce any occasional events to a spurious correlation, insignificantly changing the pattern (if any) seen with patent events.

4. Results

Table 3 presents our unweighted results using counts of patent events, with all observations included in the first column, only the lowest quintile AR_t results in the second column, and only the highest quintile AR_t results shown in the third column. All results are presented here corrected for both heteroskedasticity and first-order autocorrelation.

First, notice that on average, patent-related events do not correlate strongly with abnormal returns. Instead, it is simply industry-wide returns that correlate with Apple price movements. This result accords with the previous literature which shows limited explanatory power from patents (Hall, 2000; Hall et al., 2007). However, that average relationship masks an underlying difference between quintiles, at a level where the NASDAQ has no explanatory power for Apple's stock price.

Second, we note that application events appear to have no statistically significant effect on stock values, presumably because they are both private and uncertain in outcome, but also because we are simply counting patent events here without regard to their value or importance.

In contrast, stock prices appear to change in a pattern correlated with publication events, responding strongly to published news that Steve Jobs's name appears on a patent application under consideration. That reputation effect seems to add volatility to prices, making negative days much more negative and making positive days more positive. In contrast, the publication of design patent applications is countercyclical, dampening good days.

On status dates, the effects parallel publication dates, which is a little surprising given the private but finalized nature of the status event. There are no Steve Jobs-related status events in the bottom quintile, offering further evidence that they are associated with higher price change outcomes.

Table 3. Regression results of two-week AR_t on unweighted counts of patent events

Variable	All observations		Bottom quintile		Top quintile	
	Coefficient	t-statistic	Coefficient	t-statistic	Coefficient	t-statistic
Application						
Jobs	-0.692	(0.26)	-54.152	(0.12)	-18.104	(0.51)
Design	-0.156	(0.16)	94.625	(1.07)	10.416	(0.31)
Patent	0.262	(0.47)	-19.692	(0.89)	-5.012	(0.54)
Publication						
Jobs	-3.013	(1.03)	-458.281	(3.08) ^{***}	63.947	(1.68) [*]
Design	0.687	(0.33)	48.166	(0.32)	-39.484	(2.19) ^{**}
Patent	1.197	(1.53)	-15.149	(0.42)	-14.274	(1.01)
Status						
Jobs	3.046	(0.84)	n/a	n/a	59.170	(2.62) ^{***}
Design	-1.494	(0.57)	-64.553	(2.04) ^{**}	-140.152	(4.33) ^{***}
Patent	0.251	(0.42)	41.978	(2.01) ^{**}	20.635	(1.02)
Issue						
Jobs	1.885	(0.55)	94.625	(1.07)	-130.870	(2.98) ^{***}
Design	-1.231	(0.51)	68.625	(1.76) [*]	57.132	(1.47)
Patent	-0.477	(0.70)	34.955	(1.52)	13.837	(0.79)
Time	0.001	(0.38)	-0.012	(1.00)	0.003	(0.57)
NASDAQ	14.636	(18.19) ^{***}	6.326	(1.22)	1.976	(0.84)
Constant	-10.819	(0.54)	-641.828	(11.0) ^{***}	552.235	(17.5) ^{***}
Summary Statistics						
Obs	9502		198		198	
F-statistic	23.08 ^{***}		73.97 ^{***}		109.76 ^{***}	

All coefficients are multiplied by 100 for readability. Significance is indicated in the table as ^{***} for 99%, ^{**} for 95%, ^{*} for 90%.

On issue dates, Steve Jobs acts as an ameliorating force, exerting negatively correlation with already highly positive outcomes. Perhaps this suggests that his positive effect is earlier in the process, bumping or hyping up prices at the publication date and status date but showing an offsetting decline at the time of issuance. This interpretation would lead to the conclusion that Jobs patents are no more valuable than other patents, but simply exert an earlier impact on the market due to a name recognition effect.

Overall, the only occasion on which patent counts act as a statistically significant correlate with abnormal returns is on their status date, and even then not in all quintiles. This clearly calls for more investigation.

Robustness tests were performed in two broad methods: using alternatives to patent counts, which should better reflect the variation in quality of the patent-related events (results using originality measures are in Table 4, with results using other measures available from the authors), and using alternative event windows, which might represent a more immediate response to news (a 2-day event window is shown in Table 5).

While similar to the primary results in most ways, there are some different results to note. Using originality-weighted counts of patents, the publication event becomes statistically insignificant. More importantly, when considering a shorter event window, the application event becomes statistically significant, at least for the bottom quintile.

Table 4. Regression results of two-week AR_t on originality-weighted counts of patent events

Variable	All observations		Bottom quintile		Top quintile	
	Coefficient	t-statistic	Coefficient	t-statistic	Coefficient	t-statistic
Application						
Jobs	-0.664	(0.26)	-223.543	(0.43)	-29.508	(0.77)
Design	-0.124	(0.13)	29.237	(0.09)	19.734	(0.53)
Originality	1.638	(0.45)	56.233	(0.61)	69.497	(0.66)
Publication						
Jobs	-2.624	(0.92)	-391.581	(2.69) ^{***}	79.667	(0.91)
Design	1.175	(0.59)	-8.408	(0.06)	-43.804	(0.92)
Originality	-6.130	(0.67)	-431.720	(1.42)	-54.154	(0.27)
Status						
Jobs	2.852	(0.78)	n/a	n/a	68.159	(2.11) ^{**}
Design	-1.115	(0.46)	-68.671	(2.59) ^{***}	-111.045	(6.11) ^{***}
Originality	0.634	(0.12)	264.928	(2.25) ^{**}	0.947	(0.02)
Issue						
Jobs	0.892	(0.26)	207.601	(1.83) [*]	-134.330	(1.92) [*]
Design	-0.412	(0.17)	82.888	(2.83) ^{***}	34.174	(0.52)
Originality	10.042	(1.06)	294.101	(1.19)	16.745	(0.10)
Time	0.001	(0.39)	-0.013	(1.02)	0.003	(0.59)
NASDAQ	14.616	(18.19) ^{***}	6.710	(1.32)	1.816	(0.73)
Constant	-10.915	(0.54)	-649.433	(11.1) ^{***}	5.567	(17.8) ^{***}
Summary Statistics						
Obs	9502		198		198	
F-statistic	22.86 ^{***}		75.02 ^{***}		109.18 ^{***}	

All coefficients are multiplied by 100 for readability. Significance is indicated in the table as ^{***} for 99%, ^{**} for 95%, ^{*} for 90%.

5. Conclusion

The goal of this paper was to examine the role that patent-related events have on share prices, given the historical difficulties in identifying their effects. We proposed that two factors account for that quantitative difficulty: the variability in patent value and the long and uncertain timeframe during which information leaks to the market. This paper has modeled both of those factors for Apple Inc., while incorporating two other elements which we believe to be new to the patent-finance literature: a quantified personality effect and the analysis of design patents in contrast to utility-model patents.

Our findings are admittedly limited in usefulness. We confirm that (for the most part), application events are unimportant to stock prices, presumably because they are events private to the applicant and represent only an uncertain amount of information about the final asset. In contrast, publication events appear to matter for certain quintiles of the data, as a true market event. Since publication has only been separated from application since 1995, it is worth remarking on this effect of policy which (to our knowledge) has never been evaluated previously. Status events appear to matter, a fact which is surprising because such events are only privately communicated to the applicants and should represent no public market information. As status events have not been evaluated in previous literature, we have no backdrop against which to evaluate our results. Finally, the issuance of a patent shows no residual stock market effect, confirming the literature's findings of negligible impacts, which in our model is unsurprising given the amount of information leaked to the market in advance of the final issuance.

Table 5. Regression results of two-day AR_t on unweighted counts of patent events

Variable	All observations		Bottom quintile		Top quintile	
	Coefficient	t-statistic	Coefficient	t-statistic	Coefficient	t-statistic
Application						
Jobs	1.719	(0.22)	-54.250	(2.12)**	200.371	(0.91)
Design	-1.643	(0.58)	170.972	(2.78)***	-10.811	(1.13)
Patent	-1.217	(0.54)	-197.047	(2.33)**	-4.035	(0.65)
Publication						
Jobs	-10.379	(1.28)	-192.994	(0.93)	-542.251	(2.83)***
Design	4.504	(0.78)	281.907	(1.67)*	-33.984	(2.79)***
Patent	-2.265	(1.00)	-72.854	(0.85)	-12.488	(0.67)
Status						
Jobs	-6.911	(0.69)	54.589	(0.26)	-34.604	(1.34)
Design	5.365	(0.78)	238.435	(2.74)***	119.033	(3.02)***
Patent	1.339	(0.74)	23.293	(0.58)	11.835	(0.80)
Issue						
Jobs	-1.313	(0.15)	-585.475	(1.46)	-61.574	(0.75)
Design	2.817	(0.50)	157.111	(1.13)	-56.585	(1.97)*
Patent	1.613	(0.91)	94.557	(1.71)*	-14.023	(0.95)
Time	0.001	(0.07)	-0.001	(0.68)	0.001	(0.10)
NASDAQ	45.531	(23.68)***	-4.695	(0.61)	7.053	(0.95)
Constant	-4.273	(0.44)	-716.735	(7.80)***	757.669	(16.1)***
Obs	9502		198		198	
F-statistic	39.47***		23.22***		97.86***	

All coefficients are multiplied by 100 for readability. Significance is indicated in the table as *** for 99%, ** for 95%, * for 90%.

Personality or reputation appears to matter as well. Steve Jobs's name alone on a patent appears to correlate with stock price variations, frontloading the value onto the publication and status events at the expense of the issuance event. The net effect is statistically indistinguishable from zero, making the entire story coincide nicely with a "marketing" interpretation where name-brand recognition nets early investment which may later trail off.

Design patents matter to the market, and they matter differently than utility-model patents matter. This is presumably a result specific to our sample, which evaluates the stock price of a company known for both types of patents. However, it suggests that future analysis should separate out the effects; in our case, designs had much greater impact, particularly negative impact, at various quintiles of the sample.

Naturally, there are many weaknesses in our analysis which we hope to address in future work. For example, it is remarkable that we see any correlations at all, given that we have not coded the institutional history of the company under consideration. We would anticipate that subsequent work will identify important events and therefore quantify patent-related event effects more cleanly, not only for this company example but for others in the industry and more broadly.

References

Adams, R., Almeida, H., & Ferreira, D. (2005). Powerful CEOs and their impact on corporate performance. *Review of Financial Studies*, 18, 1403-1432. <http://dx.doi.org/10.1093/rfs/hhi030>

- Fama, E., & French, K. (1996). Multifactor explanations of asset pricing anomalies. *Journal of Finance*, 51, 55-84. <http://dx.doi.org/10.1111/j.1540-6261.1996.tb05202.x>
- French, K. (2013). Data Retrieved February 14, 2013, from http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html
- Hall, B. H. (2000). Innovation and Market Value. In R. Barrell, G. Mason, and M. O'Mahoney (Eds.), *Productivity, Innovation and Economic Performance* (2000). Cambridge, England: Cambridge University Press.
- Hall, B. H., Jaffe, A. B., & Trajtenberg, M. (2001). The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. *NBER Working Paper 8498*.
- Hall, B. H., Thoma, G., & Torrisi, S. (2007). The market value of patents and R&D: Evidence from European firms. *NBER Working Paper 13426*.
- Johnson, D. K. N., & Popp, D. (2003). Forced Out of the Closet: The Impact of the American Inventors Protection Act on the Timing of Patent Disclosure. *Rand Journal of Economics*, 34(1), Spring, 96-112. <http://dx.doi.org/10.2307/3087445>
- Kortum, S., & Lerner, J. (1998). Stronger protection or technological revolution: what is behind the recent surge in patenting?. *Carnegie-Mellon Conference Series on Public Policy*, 48, 247-304. [http://dx.doi.org/10.1016/S0167-2231\(98\)00023-2](http://dx.doi.org/10.1016/S0167-2231(98)00023-2)
- Lanjouw, J. O., & Schankerman, M. (1999). The Quality of Ideas: Measuring Innovation with Multiple Indicators. *NBER Working Paper #7345*.
- Lanjouw, J. O., Pakes, A., & Putnam, J. (1996). How to Count Patents and Value Intellectual Property: Uses of Patent Renewal and Application Data. *NBER Working Paper #5741*.
- Pakes, A., & Simpson, M. (1989). Patent Renewal Data. *Brookings Papers: Microeconomics*, 331-410. <http://dx.doi.org/10.2307/2534724>
- Sneed, K. A., & Johnson, D. K. N. (2009). Selling Ideas: The Determinants of Patent Value in an Auction Environment. *R&D Management*, 39(1), 87-94. <http://dx.doi.org/10.1111/j.1467-9310.2008.00540.x>
- Trajtenberg, M. (1990). A Penny for Your Quotes: Patent Citations and the Value of Innovations. *Rand Journal of Economics*, 21, 172-187. <http://dx.doi.org/10.2307/2555502>