EVENT STUDY METHODS AND ESTIMATION OF PATENT VALUE

by

TORUS LU

(Under the Direction of John L. Turner)

ABSTRACT

This paper applies an event study methodology to estimate the impact of patenting on firm value over the most comprehensive data set on patents to date. An event study model that accurately accounts for the ordinary pattern of stock market returns among patenting firms estimates nearly zero effect of patent grant events on firm value. This sample mean does not mask significant heterogeneity related to patent citation. These results may be due to the substantial noise involved in stock market returns, but methods recently devised for separating patent value signal from this noise do not produce credible patent value estimates.

INDEX WORDS: Patents, Event study, CRSP, Patent citations, Innovation

EVENT STUDY METHODS AND ESTIMATION OF PATENT VALUE

by

TORUS LU

A Thesis Submitted to the Graduate Faculty of The University of Georgia in Partial Fulfillment of the Requirements for the Degree

MASTER OF ARTS

ATHENS, GEORGIA

2019

EVENT STUDY METHODS AND ESTIMATION OF PATENT VALUE

by

TORUS LU

Major Professor:	John L. Turner
Committee:	Martin Gervais
	Thomas Quan

Electronic Version Approved:

Ron Walcott
Interim Dean of the Graduate School
The University of Georgia
December 2019

ACKNOWLEDGEMENTS

I would like to thank my committee chair, Dr. John L. Turner, for his formative guidance over the years. As an instructor, he instilled in me a great interest in economics. His willingness to take me on in support of his research developed my interest in the topics that eventually led to this thesis. As my committee chair, he helped me from the start, providing a source of knowledge and motivation. I would also like to thank my other committee members, Dr. Thomas Quan and Dr. Martin Gervais, for their guidance and advice as I pursued my research.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

Patent systems exist to help innovators internalize the socially-beneficial externalities of innovation spillovers. By granting exclusive rights over the use of innovations, patents can allow innovators to appropriate the benefits of their innovations, generating private value. The extent to which patents benefit innovators is therefore relevant to research into how patent systems can be optimally-designed. However, the extent to which patents benefit innovators is also an open empirical question, and a variety of techniques have been applied toward answering this question. I focus on methods applied by researchers to estimate the private value of patents, and in particular, event study methods.

Event studies describe the use of stock market returns near the occurence of an event relevant to the value of a publicly-traded firm to estimate the event's impact on the firm's value. In an efficient market, the impact of the event should be reflected as a reasonably quick change in prices, insofar as the event represents new information about the prospects of the corresponding firm. Event studies originate with Fama et al. (1969), who applied the method toward studying the efficiency of the stock market in response to stock splits. Event studies have since been applied to a wide variety of types of events, including regulatory changes (Lamdin 1999), macroeconomic variable announcements (Andritzky, Bannister, and

Tamirisa 2007), earnings releases (Louhichi 2008), security issues (Asquith and Mullins 1986), and merger announcements (McAfee and Williams 1988).

In the realm of patents, event studies have been previously applied toward studying events related to the patenting process, with major emphasis on patent grants. Patenting in a variety of specific industries has been studied using event studies. Austin (1993) uses event studies to estimate the abnormal returns associated with grant events of patents owned by large biotechnology firms in 1991. Sellers-Rubio, Nicolau-Gonzálbez, and Mas-Ruiz (2007) estimate abnormal returns of 19 patent grant events involving Spanish electrical sector firms. Patel and Ward (2011) estimate abnormal returns of grant events for 23 pharmaceutical firms over 1980-1995. Developments in patents data over time have been essential to expanding the scope of patents research. This is particularly important in the case of event studies. The NBER data files compile millions of patent grants and have been used to study the relationship between citations and abnormal returns (Dornelles and Ali 2014; Patel and Ward 2011) as well as expand the scope of events studied (Dornelles and Ali 2014).[1] Most recently, work by Kogan et al. (2017) has developed the broadest data set of patent grants to date, in particular expanding time coverage relative to the NBER files. These studies have all emphasized used patent grant data derived in some way from US Patent and Trademark Office (USPTO) records, implicitly assuming that when a USPTO action occurred is equivalent or close to equivalent to when that action became public knowledge. Other research has sought alternative ways of determining informative patent grant events. Chen et al. (2018), for example, constructed a data set on patent grant events from announcements in the Lexis/Nexis Newswire database.

[1] The NBER data files are described in detail in Hall, Jaffe, and Trajtenberg (2001).

I analyze the Kogan et al. (2017) data set using three event study models to determine the appropriate method of estimating abnormal returns attributable to patent grants. I find that a market-adjusted return model underestimates the normal returns of publicly-traded patenting firms, thereby overestimating the abnormal returns attributed to patent grants. Reasoning from a market model, this may occur because patenting firms generally outperform the market or because patenting firms have higher return volatility than the market, with higher average returns representing compensation for that volatility. I additionally find that models of normal returns that do not account for day-of-week effects may produce biased estimates of the abnormal returns associated with patent grants. The direction of this bias is dependent on the specified event window, since event window decision may affect the representation of particular days of the week in the event window. A market model specification, with controls for day-of-week effects, resolves both of these biases. When a market model with day-of-week effects is used, the average abnormal return associated with patent grant events is not significant. Using a market-adjusted return model, one would be wrongly led to believe that the average abnormal return in a small window around patent grants is positive and significant at standard levels.

Additionally, I analyze cross-sectional variation in abnormal returns to determine whether the null finding overall masks substantial heterogeneity, finding that there is little meaningful variation related to patenting frequency, citations, or firm size. As measured by event studies, the impact of an event on a firm's value is attenuated by the degree to which it was anticipated by investors. To determine whether variation in the degree to which patent grants are anticipated drives cross-sectional variation in abnormal returns, I separate patenting firms into subgroups based on their patenting frequency, which might proxy for the surprise factor of patent grant events. I find that patenting frequency is not associated with significant differences in abnormal returns. The abnormal returns at the extreme are higher in mean, but not significantly different from zero.

Whether the size of the patentee influences the effect of a patent grant is an open question. If patents granted to small firms and large firms have approximately similar values, then abnormal returns of small firms should be larger, ceteris paribus. On the other hand, patents granted to firms of different sizes might not have similar values. Large firms may produce more valuable patents, or may be better positioned to exploit their patents. Small firms, as measured by market capitalization, have substantially higher abnormal returns. In the size distribution of patenting firms, the bottom decile of firms within each year experience a mean abnormal return of 0.35% in the four days following a patent grant event. Large firms experience significant negative abnormal returns. However, when firm size is instead defined by revenue, there is no significant difference in abnormal returns across the firm size distribution.

Prior literature has suggested a link between forward citations, which may measure the scientific value of a patent, and the private value of the patent. The mean abnormal return of highly-cited patents is significant, which provides some evidence of a relationship between citations and abnormal returns around patent grant dates. However, when firm and year effects are controlled for, patent grant events involving highly-cited patents are not associated with significantly larger abnormal returns.

These results suggest that patents grants are not generally associated with significant positive abnormal returns, and there is little significant variation in abnormal returns by citation counts, patenting frequency, and firm size. Given the substantial body of research linking patenting to value, a reasonable explanation for these findings is that individual patent grant events have relatively minor impacts overwhelmed by noise in stock market returns. Moreover, this noise might also obscure meaningful variation in patent value based on the characteristics discussed.

In the interest of solving this noise issue, Kogan et al. (2017) develop methodology for estimating patent value signals based on abnormal returns. Kogan et al. demonstrate applications of their estimates in studying the relationship between scientific value and private value, the impact of patenting on rivalrous competition between firms, and the impact of innovation on aggregate growth. In the first application, Kogan et al. find that an additional forward citation is associated with a 0.1% to 3.2% increase in the value of a patent. This is a notably different result from the lack of a relationship between citations and abnormal returns.

To examine the assumptions behind the Kogan et al. estimation and why it may produce results that differ substantially from those found with abnormal returns, I review the Kogan et al. method. In particular, the assumptions on the signal-to-noise ratio and the model of normal returns may affect cross-sectional variation between patent value estimates. I also perform a placebo test of the method on randomized grant dates. This placebo test produces signal-to-noise and patent value estimates comparable to the real patent value estimates. Moreover, when placebo value estimates are related to citation counts, I find highly significant coefficient estimates comparable to that of Kogan et al. Separation of patent value from noise in the real abnormal returns therefore cannot explain why significant citations effects are found only after using the method. Endogeneity related to the number of patents granted presents a plausible alternative explanation.

My findings of a null overall effect of patenting and little heterogeneity with respect to citations contrast to some extent with prior event studies of patent grants. Patel and Ward (2011) finds significant positive abnormal returns in their data set of pharmaceuticals. Conditional on some characterstics that may proxy for publicity or high value of the underlying technology, Austin (1993) finds significant positive abnormal returns in his data set of biotechnology firms. Sellers-Rubio, Nicolau-Gonzálbez, and Mas-Ruiz (2007) finds significant positive abnormal returns in their data set of electrical sector firms. A direct comparison is difficult to make, since these papers study specific industries, and I do not. It may be that these industries produce particularly valuable patents, and it is easier to discern the value of these patents from the noise in stock returns. Dornelles and Ali (2014) does not restrict their data to a specific industry or sector and finds more ambiguous results, suggesting that patent grants were not associated with significant abnormal returns from 2000 to 2006, but were associated with significant positive abnormal returns from 1995 to 2000. Alternatively, methodological differences may explain the contrast. The referenced event studies do not include day of week effects in the model of normal returns, although this is not the only difference. Some use the Fama/French factors in their model of normal returns (Dornelles and Ali 2014; Patel and Ward 2011).[2]

[2] Fama/French factors are described in detail in Fama and French (1993).

I note day-of-week effects in the stock returns of firms in the Kogan et al. data set. These firms are, by design, patenting firms, but day-of-week effects are by no means unique to patenting firms. French (1980) finds significant negative average returns on Mondays and particularly large returns on Wednesdays and Fridays for the S&P over 1953 to 1977. Gibbons and Hess (1981) finds significant negative average returns on Mondays for a broader set of portfolios, including the CRSP value-weighted portfolio. Return volatility may be correlated with the day of week, as Fama (1965) finds that Mondays are associated with an approximately 20% increase in return volatility. These particular patterns are consistent with the patterns I observe in patenting firms' returns. Following these findings, research applying event studies frequently make note of whether events cluster around particular days of the week (Andritzky, Bannister, and Tamirisa 2007; Allen and Sirmans 1987; Peress 2014) and/or include day-of-week effects in a specification (Andritzky, Bannister, and Tamirisa 2007; Peress 2014). In the specification Kogan et al. (2017) use to estimate the increase in return volatility associated with a patent grant, day-of-week fixed effects are present.

The literature on estimating patent value is broad, and event study methods are just one area. Prior research has used surveys of patent-holders' for their valuations of their patents (Harhoff, Scherer, and Vopel 2003). Given that USPTO periodically assesses costly renewal fees to extend patent rights through their full potential term, patent values can also be inferred from renewal decisions (Pakes and Schankerman 1984). Regression of a firm's market value, or some variant such as Tobin's Q, on its patent stock may also be used to estimate the contribution of patents (Bessen 2009; Bessen et al. 2018; Hall, Jaffe, and

Trajtenberg 2005; Nicholas 2008). The last method is similar to an event study method insofar as both relate the market value of a firm to patenting.

I discuss the relationship between forward citations and measures of the private value of patents. The theory behind such a relationship is that forward citations may measure the scientific value of a patent, and the relationship between scientific value and private value is of interest to researchers. Alternatively, citations may simply be a useful measure of patent quality, and prior research has emphasized that, if not all patents are made equal, a measure of patent quality is an informative complement to data on the quantity of patents held by innovators. As such, the existing literature studying the relationship between citations and value is substantial. Hall, Jaffe, and Trajtenberg (2005) relate Tobin's Q to patent citations, finding that one additional citation for each patent increase the market value of a firm by 3%. Nicholas (2008) relates firms' returns and Tobin's Q between 1910 and 1939 to patent and citation stocks, finding significant positive effects on both returns and Tobin's Q during the 1920s. Harhoff, Scherer, and Vopel (2003) relate their surveyed patent value estimates with citations, finding significant positive effects of citations. Although citations have frequently been found to have a significant positive relationship to measures of private patent value, this point is not uncontested. Studying non-practicing entities, i.e. firms that primarily license patents rather than directly entering markets in which their patents may be relevent, Abrams, Akcigit, and Grennan (2018) find that citations have an inverted-U relationship with patent value, with the most-valuable patents having relatively few citations. Among event studies, Patel and Ward (2011) finds no relationship between citations and abnormal returns at the patent grant event, but finds significant effects when the citations actually occur. My finding of a lack of significant heterogeneity related to citations is consistent with the former result, and I do not observe these citation events.

CHAPTER 2

DATA

As stated in the introduction section, this paper is focused on analyzing the data compiled by Kogan et al. (2017). Kogan et al. compile their patents data from the source files available via Google Patents and make it publicly-available. They estimate the value of 1,801,879 patents granted over 1926-2010. Theirs is the most comprehensive data set on patents to date. Another data set used by researchers interested in broad study of patents is the NBER Patent Data Project. Comparatively, Kogan et al. include more than 500 thousand patents not included in the NBER data, which also has a more restricted time coverage, spanning 1976-2006. The comprehensiveness of this data set makes it a natural choice for use in patents research.

Observations in the patents data set are unique patent grants and contain information on the patent number, dates of application filing and patent grant, patent technology class and subclass, number of forward citations, and Kogan et al. estimated patent value. The data set also contains the patenting firm's CRSP permno, which enables merging the patents data set with CRSP's stock market and firm data. From CRSP, I draw data on share volume, shares outstanding, closing prices, industry codes (SIC), and market returns. Additionally, I draw data on the firms' revenue from Compustat.

Some of the event studies that I discuss were performed through an application provided by CRSP through Wharton Research Data Services. However, one event study model requires controls that are not present in this application and therefore cannot be estimated using the CRSP application. A close replication of the CRSP event study application requires data on the risk-free rate of return as well as the market return, which I retrieve from Kenneth French's website.

The patents data set contains errors in recorded filing and grant dates. I correct dates for observations that clearly indicate a problem.[3] I use an automation script to find the correct dates on Google Patents. Dates are corrected for 5,790 observations. Since this is a small fraction of the entire data set, it is unlikely that the corrections affect results substantially.

[3] Indicators of a problematic observation were as follows: (1) a grant date that was not on a Tuesday, since the USPTO typically issues patents and publishes patent issues in its *Official Gazette* on Tuesday, and (2) a grant date on or before the corresponding filing date, since patents are only granted after an application is filed. Regarding (1), it should be noted that USPTO does occasionally deviate from this policy, e.g. patent 3,588,741 granted on June 28, 1971, a Monday. Still, patents are granted on a Tuesday for approximately 99.9% of the corrected data set.

CHAPTER 3

EVENT STUDY

3.1 MODELLING ABNORMAL RETURNS

Since the purpose of an event study is to identify the impact of a specific event taking place on a specific date on the value of a firms, the ordinary pattern of returns for a firm should not be attributed to the effect of the event. It is therefore necessary to estimate the normal returns of firms studied and difference out these estimates. What constitutes a reasonable model of normal returns is contingent on the particular firms and events studied, so a variety of abnormal returns models have been proposed and applied in various contexts. This section will discuss three models in the context of patent grant events: a market-adjusted returns model, a market model, and a market model controlling for day-of-week effects.

A market-adjusted returns takes the market return as the normal return, and differences it from the firm's return within the relevant event window,

$$AR_i = R_i - E(R_i)$$
$$E(R_i) = R_m \tag{1}$$

where AR_i is the abnormal return under market adjustment, R_i is the firm's return, and R_m is the market return. The basic alternative model to be considered is the market model. In the market model, abnormal returns are the difference between the firm's return and an expected return. The expected return is estimated using the following equation,

$$AR_i = R_i - E(R_i)$$

$$E(R_i) = \beta_0 + \beta_1 R_m \tag{2}$$

where R_m is the market return and β_0 and β_1 are parameters to be estimated by regressing the firm's return against the market return in an estimation window prior to the event window. It can be seen that, if β_0 is set to 0 and β_1 is set to 1, expected returns are simply equal to the market return, and the market model is equivalent to the market-adjusted model. Using market-adjusted returns is therefore equivalent to using the market model and fixing the parameters to these values (MacKinlay 1997). With market-adjusted returns, parameters do not need to be estimated and an estimation window is unnecessary. This could be advantageous if, for example, a patent grant event occurs early in the firm's public trading history, and the sample of returns prior to the event is too small to estimate the parameters. Or if there are patent grant events taking place within the estimation window, the market model estimation might reflect the effect of those events rather than truly normal performance, while the market-adjusted model would be unaffected. Fixing parameters creates a different set of problems, however. If the market model is accurate and $\beta_0 \neq 0$ or $\beta_1 \neq 1$, then fixing the parameters to those values could bias abnormal return estimates.

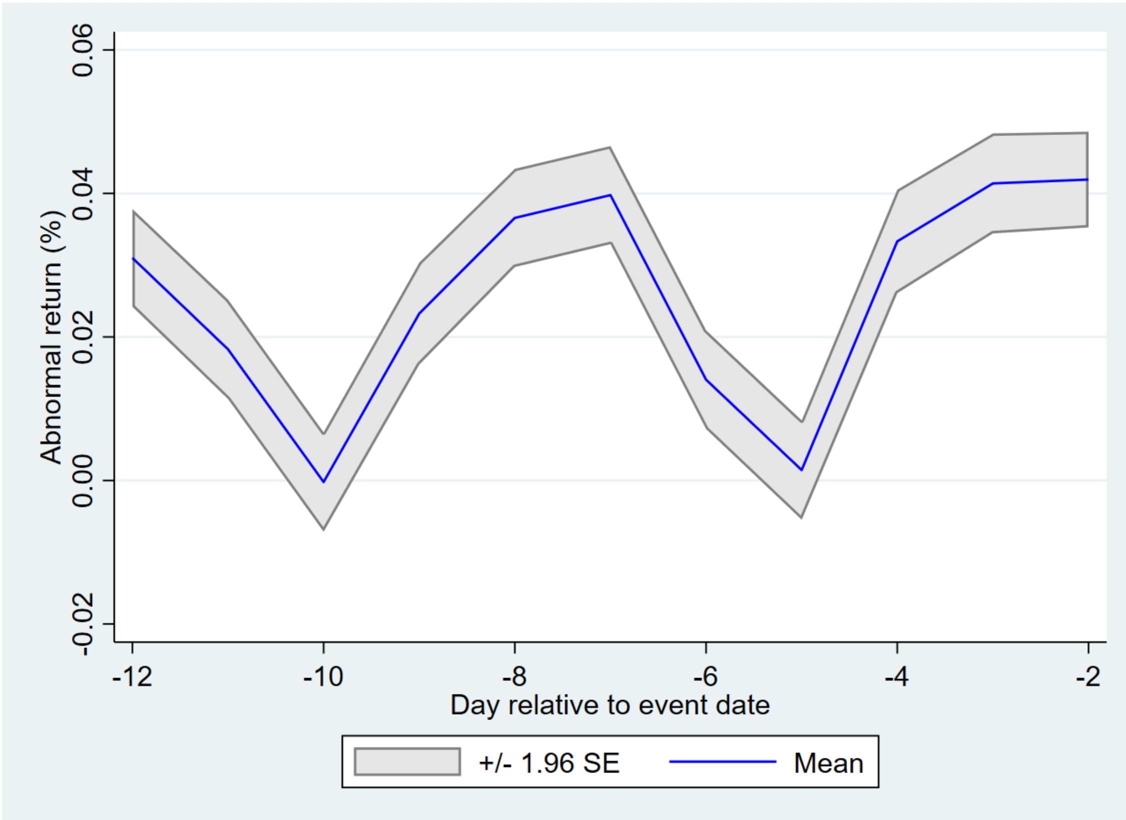

Figure 1: Mean market-adjusted returns, pre-event

Note: Average abnormal returns on days prior to the event date are given with 95% confidence intervals. Abnormal returns are market-adjusted and daily.

To check whether the market-adjusted model of returns is appropriate for an event study of patent grants, I use the CRSP event study application to estimate market-adjusted returns for each event in a $[t-12, t-2]$ window, where t is the day of patent grant. Since patents are typically granted on Tuesdays, this window will typically represent the two weeks prior to the week in which the patent is granted. Figure 1 shows the daily abnormal returns for events within this window. Daily abnormal returns are mostly positive, although some are approximately zero. The cumulative abnormal return, given by summing each daily mean, is 0.28%, and highly significant. Given that the event has not actually occurred within this window, these returns should be normal returns and not reflect the effect of the event.[4] If the cumulative abnormal return over such a window is not approximately zero, then there is reason to believe that the market return does not properly estimate the normal return of an event. A significant positive abnormal return over the depicted window therefore suggests that the market-adjusted model underestimates the firms' normal returns.

The market return could underestimate normal returns for a variety of reasons. Since a market-adjusted model is simply a market model with parameters assumed to be fixed at particular values, failure of these assumptions may explain why the normal return is underestimated. If the market model is true and β_0 is positive, i.e. the firms ordinarily outperform the market, then normal returns would be greater than market returns. It may be that patenting firms are innovative and outperform the market. If β_1 is greater than one, i.e. the firms have higher beta risk than the market, and market returns tend to be positive, then normal returns would be greater than market returns on average. In this case, the average overperformance relative to the market is compensation for greater beta risk among patenting firms. This explanations are not mutually exclusive, and a combination of these might also explain why market-adjusted returns underestimate normal returns.

The firms in the data set also appear to experience day-of-week effects in their returns. Specifically, returns are on average lower on Mondays and Tuesdays and higher on Wednesdays, Thursdays, and Fridays. In the presence of day-of-week effects, normal returns models

[4] It could be that patent grant events predict other patent grant events in this window, which would explain why daily market-adjusted returns are, on average, positive immediately after $t-5$ and $t-10$, which correspond to Tuesdays when the event itself takes place on Tuesdays and there are five trading days to a week. In this case, the estimation window returns are not normal. However, market-adjusted returns continue to be large and positive even when restricting to events where it is not the case that there were other patent grant events taking place in the $[t-12, t-2]$ window. The weekly pattern is better-explained by day-of-week effects.

are often specified with fixed effects for day-of-week (Andritzky, Bannister, and Tamirisa 2007; Peress 2014). These effects are present in non-patent grant weeks, which suggests that they are unrelated to the effect of the patent grant. Day-of-week effects are apparent in Figure 1, with market-adjusted returns following a cyclical pattern. Although only two weeks are depicted, this pattern holds in larger windows. Because patents are almost always granted on Tuesdays, the abnormal returns reflect the day-of-week effects, even when the market model is used. These effects may cause both the market-adjusted model and the market model to produce misleading estimates of abnormal returns, if the window is selected in such a way that overrepresents later days in the week. The third model is therefore an adjustment of the market model to account for day-of-week effects by allowing parameters to vary by day-of-week, so that,

$$\begin{aligned} AR_i &= R_i - E(R_i) \\ E(R_i) &= \beta_0 + \beta_1 R_m + \sum_{i=1}^{6} \alpha_i DAY_i + \sum_{i=1}^{6} \gamma_i DAY_i R_m \end{aligned} \tag{3}$$

where DAY_i are dummy variables equal to one if the return occurs on a particular day, e.g. $DAY_1 = 1$ if the return occurs on a Monday, $DAY_2 = 1$ if the return occurs on a Tuesday, etc. This model additionally allows the beta risk paremeter to vary by day, which may be useful in predicting normal variation in firms' returns if day-of-week heterogenously affects the firm's volatility and correlation with market returns depending on the day of week.

3.2 RESULTS

Average cumulative abnormal return estimates, using the three discussed models of normal returns and multiple window lengths, are displayed in Table 1. Since cumulative abnormal returns are simply the sum of daily abnormal returns within the relevant window, differences between the rows represent daily abnormal returns. The first column provides cumulative market-adjusted returns, i.e. abnormal returns estimated using the expected return provided by equation (1). Market-adjusted returns are larger than the abnormal returns produced by other two models. Market-adjusted returns also display a cyclical pattern, with abnormal returns being close to zero on the day of grant, four days after, and five days after. These days correspond to Mondays and Tuesdays if the patent grant event occurred on a Tuesday after 1952.[5] All other days see positive market-adjusted returns. The mean cumulative market-adjusted return for the $[t, t+2]$ window is 0.09%, while the market-adjusted return for the $[t, t+3]$ window is 0.12%. Both of these means are highly signficant statistically, although not informative about the average effect of the patent grant event, if they are driven primarily by underestimated normal returns and day-of-week effects. In the longer windows, market-adjusted returns consistently grow, with the cumulative abnormal return of the $[t, t+8]$ window tripling that of the $[t, t+2]$ window. Although not depicted, market-adjusted returns consistently grow beyond this window. This suggests that the market return is underestimating the normal return for the firms studied. Due to this underestimation, cumulative abnormal returns are biased upwards regardless of window choice.

[5] Prior to 1952, the New York Stock Exchange opened on Saturdays, so the typical week contained six trading days. Otherwise, weeks typically contain five trading days.

Table 1: Cumulative abnormal returns

Window	Market-adjusted	Market	Market with day effects
t	0.002	-0.019^{***}	-0.001
	(0.003)	(0.003)	(0.004)
t,t+1	0.048^{***}	-0.005	0.000
	(0.005)	(0.005)	(0.005)
t,t+2	0.091^{***}	0.014^{*}	0.008
	(0.006)	(0.006)	(0.006)
t,t+3	0.122^{***}	0.024^{***}	0.005
	(0.007)	(0.007)	(0.007)
t,t+4	0.122^{***}	0.006	0.011
	(0.007)	(0.007)	(0.008)
t,t+5	0.124^{***}	-0.014	0.004
	(0.008)	(0.008)	(0.008)
t,t+6	0.175^{***}	0.006	0.012
	(0.009)	(0.009)	(0.009)
t,t+7	0.213^{***}	0.021^{*}	0.013
	(0.009)	(0.009)	(0.009)
t,t+8	0.239^{***}	0.028^{**}	0.014
	(0.010)	(0.010)	(0.010)

Note: Abnormal returns are in %. Column 1 estimates abnormal returns where the normal return is given by $E(R_i) = R_m$. Column 2 estimates abnormal returns where the normal return is given by $E(R_i) = \beta_0 + \beta_1 R_m$. Column 3 estimates abnormal returns where the normal return is given by $E(R_i) = \beta_0 + \beta_1 R_m + \sum_{i=1}^{6} \alpha_i DAY_i + \sum_{i=1}^{6} \gamma_i DAY_i R_m$. Columns 2 and 3 require an estimation window prior to the event to over which parameters are estimated. The estimation window is 200 trading days long, ending six days prior to the event window. If an estimation window has less than 150 trading days with returns, then the event is excluded. Standard errors are in parentheses. Asterisks denote significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The second column of Table 1 provides the cumulative abnormal returns estimated with the market model. For these, the expected return is given by equation (2). The abnormal returns produced by the market model are smaller than the market-adjusted returns, which is what would be expected if the restrictive assumptions of the market-adjusted model were causing expected returns to be underestimated. Market model cumulative abnormal returns are smaller than the market-adjusted cumulative abnormal return in all windows, and negative in some windows. As with the market-adjusted returns, the market model abnormal returns have a cyclical pattern, with negative abnormal returns the day of grant, four days after, and five days after, and positive abnormal returns on other days. Due to the weekly pattern, the sign of the cumulative abnormal returns is sensitive to the choice of window. The mean cumulative abnormal return for the $[t, t+2]$ window is 0.01% and for the $[t, t+3]$ window is 0.02%. Unlike with market-adjusted returns, the cumulative abnormal returns do not increase substantially in the longer windows.

The third column of Table 1 provides the cumulative abnormal returns estimated with the market model, with day-of-week effects accounted for. The expected return is given by equation (3). The abnormal returns produced by this model do not have the same weekly pattern as those produced by the other two models, which is to be expected if the weekly pattern was driven by day-of-week effects that were not accounted for in the previous two models. The mean cumulative abnormal return is 0.008% for the $[t, t+2]$ window and 0.005% for the $[t, t+3]$ window. Neither of these are significant, which may suggest that patent grants do not generally represent significant positive information for the firm. Longer windows suggest that abnormal returns on later days also fluctuate around zero.

Examination of abnormal returns produced in longer event windows and under alternative expected return models reveals that the market-adjusted returns model biases estimates of abnormal returns upwards and is sensitive to window specifications that over or under represent particular days of the week. These issues are substantial. Upwards bias due to underestimating normal returns causes the market-adjusted returns model to misleadingly produce highly-significant mean abnormal returns in all windows starting from $[t, t+1]$. Day-of-week effects can cause a market-adjusted returns model or a market model to misleadingly suggest significant abnormal returns in windows representing a short period of time after a patent grant event. The overrepresentation of later days of the week is an acute problem because patents are typically granted on Tuesdays. Given these issues, there is reason to believe that a market model with day-of-week effects may better estimate the impact of patent grant events on the value of a firm.

Costs of using a market model were discussed briefly, but warrant more discussion. Using a market model or a market model with day-of-week effects requires a lengthy estimation window prior to the event during which the patenting firm was publicly-traded. Market-adjusted returns simply require a market return on the same days as the event window. As such, patent grant events that occurred early in a firm's trading history are excluded, which is a potential source of selection bias. If these firms react more to patent grants than more older firms, then the mean abnormal returns may understate the impact of patent grants on firms generally. In my event study estimates, an event was excluded if there were not at least 150 trading days with non-missing returns prior to the event. This resulted in the exclusion of approximately 1.5% of events. Since a far larger number of patents were granted to firms that are privately-held or unable to be matched to trading data, this is a relatively small selection issue. Firms that have been publicly-traded for at least a short period of time do not experience significant abnormal returns associated with patent grant events, and the vast majority of patents granted to publicly-traded firms are granted to these firms.

Additionally, estimates of equation (3) may be biased due to events occurring during the estimation windows. If estimation window returns are impacted by patent grant events, and patent grants typically have a positive effect, this would likely enter as an overestimate of β_0 that suggests normal overperformance of the firm relative to the market when this is not the case in truth. This would bias downward abnormal return estimates based on the market model. Excluding events for which the relevant firm had other events take place during the estimation window would be highly selective and involve dropping 96% of events. Restricting the data set as such does suggest higher average abnormal returns, but these are still not significantly different from zero.

3.3 CROSS-SECTIONAL VARIATION

When returns are modelled using a market model with day-of-week effects, patents are not associated with significant positive abnormal returns. However, this result might mask substantial cross-sectional variation. Since the abnormal return of an event functions as a proxy measure of the value of a patent grant event, albeit with substantial measurement error, abnormal returns may be associated with descriptive characteristics that have been shown to contribute to patent value in previous work, such as patent citations. For this analysis, the relevant model of normal returns is the market model with day-of-week fixed effects, given by equation (3). I use a $[t, t+3]$ window, which in the typical case will include Tuesday through Friday the week of the patent grant. The unit of observation is a patent grant event, which may comprise of multiple patents if the patents were granted on the same day to the same firm. Summary statistics are given in Table 2.

Table 2: Summary statistics

Variable	Mean	Median
Abnormal return	0.005	-0.153
Total citations	34.268	12
Patents granted	3.050	1
Market cap	7,717	779
Revenue	8,138	1,522

Note: Summary statistics describe observations at the event-level. Abnormal returns are in %. Market capitalization is the market capitalization of the firm at the end of trading the day prior to the event. Revenue is the prior fiscal year's revenue of the firm. Market capitalization and revenue are in nominal millions of dollars.

Prior research has emphasized a link between patent value and forward citations, which measure the number of times a patent is cited by later patents. Although forward citations are by definition post-determined, they can be seen as a proxy for the scientific value of a patent, with highly-cited patents regarded as more scientifically valuable due to relevance to later innovation. In most cases, this link is positive, with patent citations being associated with increased value (Harhoff, Scherer, and Vopel 2003; Nicholas 2008; Hall, Jaffe, and Trajtenberg 2005). However, this point has not been uncontested. Abrams, Akcigit, and Grennan (2018) find that, among the patent portfolios of non-practicing entities, citations have an inverted-U relationship with revenues attributable to the patent, with the most-valuable patents having relatively few citations.

Figure 2 plots the mean abnormal return and the corresponding 95% confidence interval for patent grant events across average citation quintiles. Average citations are measured as the total number of forward citations of the patents that were granted in the event, divided by the number of patents granted. This measure is used to avoid conflating grant events involving high-citation count patents and grant events involving many patents, which might

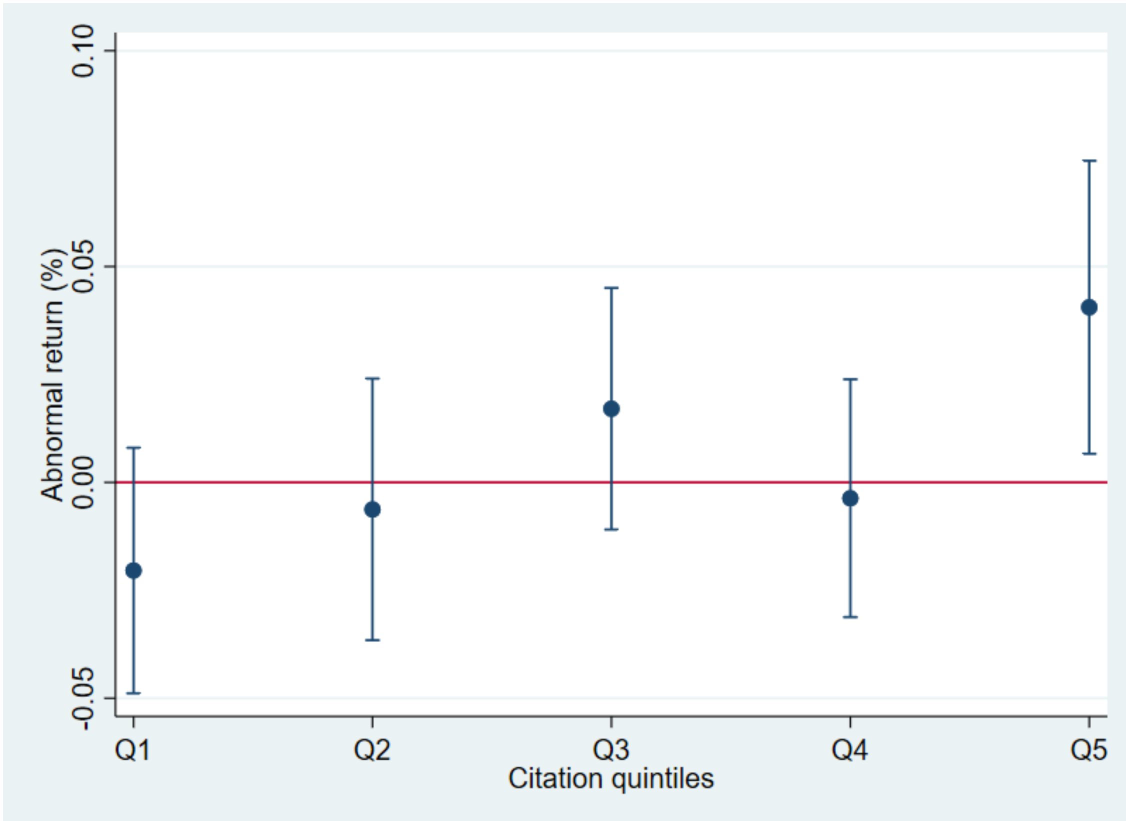

Figure 2: Mean abnormal returns by average citation count

Note: Events are binned into five approximately equally-sized groups based on the average citation count of patents granted during this event. Citation bins are defined within year, such that the bins contain approximately one-fifth of the observations in each year. 95% confidence intervals for the mean abnormal return over the $[t, t+3]$ window are given.

have the same number of total citations. As with the abnormal returns generally, the means of the four lower quintiles do not differ from zero significantly. The fifth quintile has a mean abnormal return of 0.041%, and differs from zero at a 5% significance level. Events in this category had patents receiving on average more than 16 citations. This could indicate that highly-cited patents are more likely to be significant and valuable to the patenting firm. This result should be interpreted with caution, however, since the mean is still close to zero.

The extent to which patent grants are actually surprising to patenting firms is an open question. One reason why patent grants may be predictable to investors is that some firms

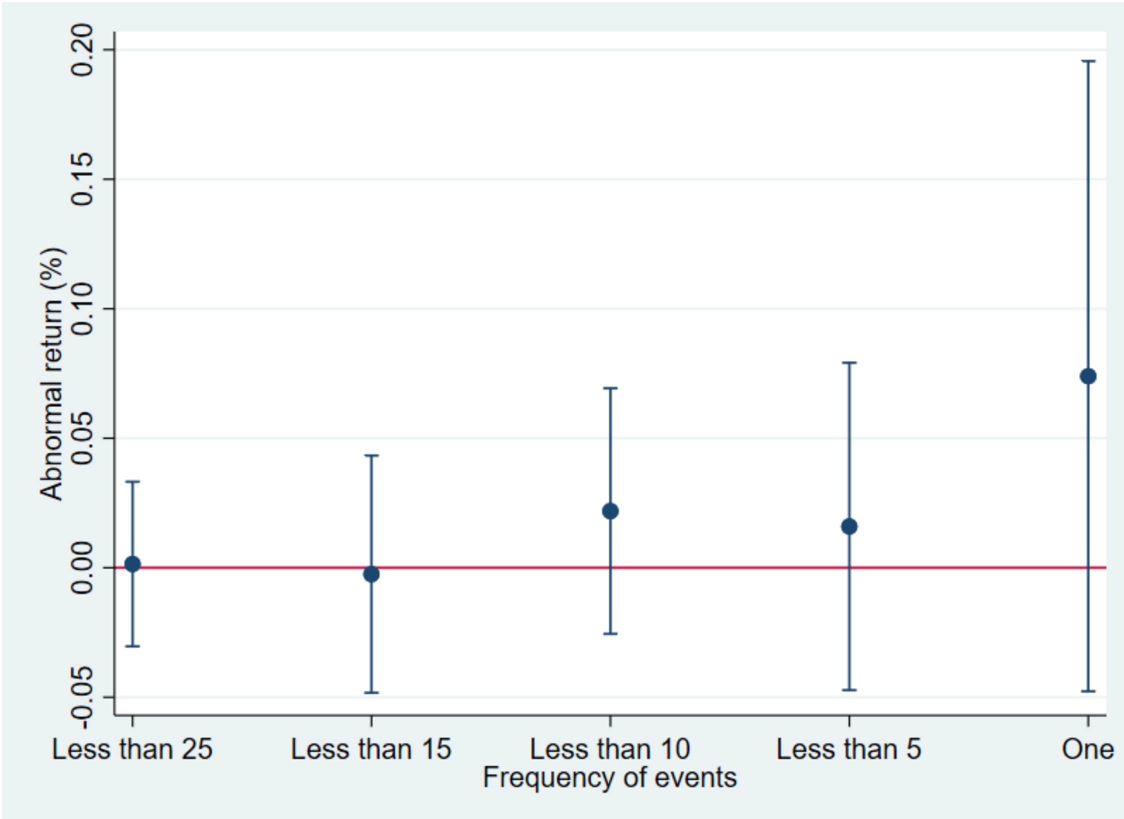

Figure 3: Mean abnormal returns by event frequency

Note: Mean abnormal returns are given under various restrictions on the data set. Restrictions are based on the number of events the patenting firm experienced in the year of grant. For example, the "Less than 10" restriction implies that only grant events to firms experiencing less than 10 patent grant events in the same year were included. 95% confidence intervals for the mean abnormal return over the $[t, t+3]$ window are given.

patent regularly. Even if information about a specific patent is not public, the market could anticipate a patent grant if there is a past pattern of receiving patents on a weekly basis. Some firms do actually patent with sufficient frequency for this to be plausible. For example, IBM received $5,052$ patents in 2010. These patent grants occurred on 44 weeks, and represent most of the days on which the USPTO ordinarily would have granted patents in 2010. In order to proxy for this patenting frequency and examine whether less frequent patentees see larger abnormal returns, I construct a variable recording the number of patent grant events the patenting firm experiences in the same year as the observation. Figure 3 plots the mean abnormal returns in five groups defined by restrictions of the number of events variable, and the corresponding 95% confidence intervals. There does not appear to be significant differences across these groups. The group of events where observations represent the only event that the patenting firm experienced during the year has a mean abnormal return of 0.074%, and higher than the sample means of the other groups, but this estimate still does not significantly differ from zero.

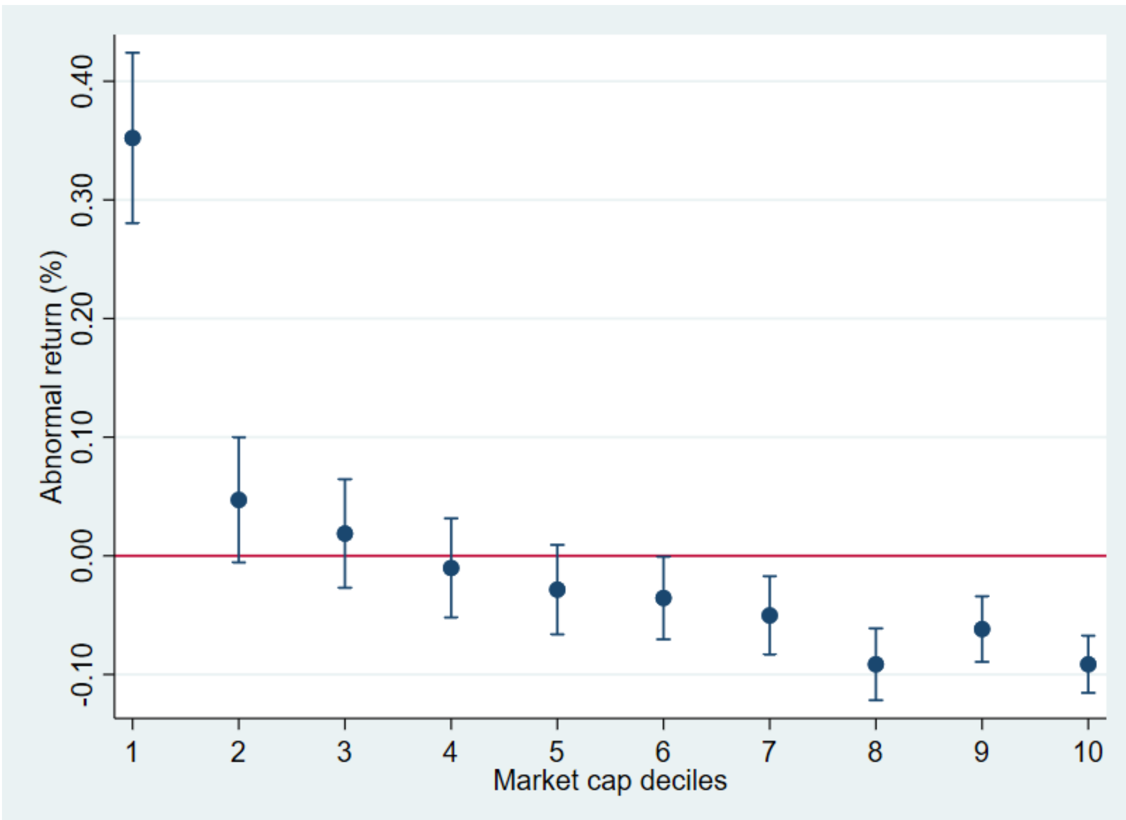

Figure 4: Mean abnormal returns by patenting firm market cap

Note: Events are binned into ten approximately equally-sized groups based on the patenting firm's market capitalization one day prior to the event. Market capitalization bins are defined within year, such that the bins contain approximately one-tenth of the observations in each year. 95% confidence intervals for the mean abnormal return over the $[t, t+3]$ window are given.

There are a number of reasons why firm size would be related to abnormal returns of patent grant events. Firm size may be related to the abnormal returns of patent grant events, mechanically, since a patent should create a smaller abnormal return at a large firm than a similarly-valued patent, in dollar terms, would create at a small firm, all else equal. Firm size may also be related to the abnormal returns of patent grant events if larger firms produce more valuable patents or are better positioned to exploit the patent grant for private gain. Previous work with event studies has suggested a negative association between abnormal returns and the size of the firm. Austin (1993) finds that larger biotechnology firms tend to have more valuable patent grants in dollar terms than smaller firms, but this association does not scale with the size of the firm. The abnormal return is therefore smaller for larger firms than smaller firms, but larger than would be expected if patents were equally-valued across firms of different size. To examine whether firm size is related to the abnormal returns of patent grant events, I construct firm size deciles, based on the relative position of the market capitalization of the patenting firm among all firms receiving patent grants within the year of grant.[6] Relative size within year is used to avoid conflating variation in firm size with variation over time, since firms have on average become larger over time. Figure 4 plots the mean abnormal returns of grant events in each firm size decile. Abnormal returns are larger for smaller firms, with the smallest decile of firm size experiencing highly significant mean abnormal returns of 0.352%. Firms in the next four deciles experience smaller abnormal returns that do not differ significantly from zero. Firms larger than the median firm experience abnormal returns that are negative in mean and statistically significant.

[6] Provided market capitalization is the measure of firm size, the following result that small firms experience larger returns is not sensitive to defining relative size within year or across all years. When firms are sorted into size deciles by CPI-deflated market capitalization across all years, mean abnormal returns are approximately the same for firms in the smallest decile and decrease in larger deciles in similar fashion to Figure 4. The only notable difference is that firms in the second decile experience mean abnormal returns that are somewhat larger and significant at a 5% level.

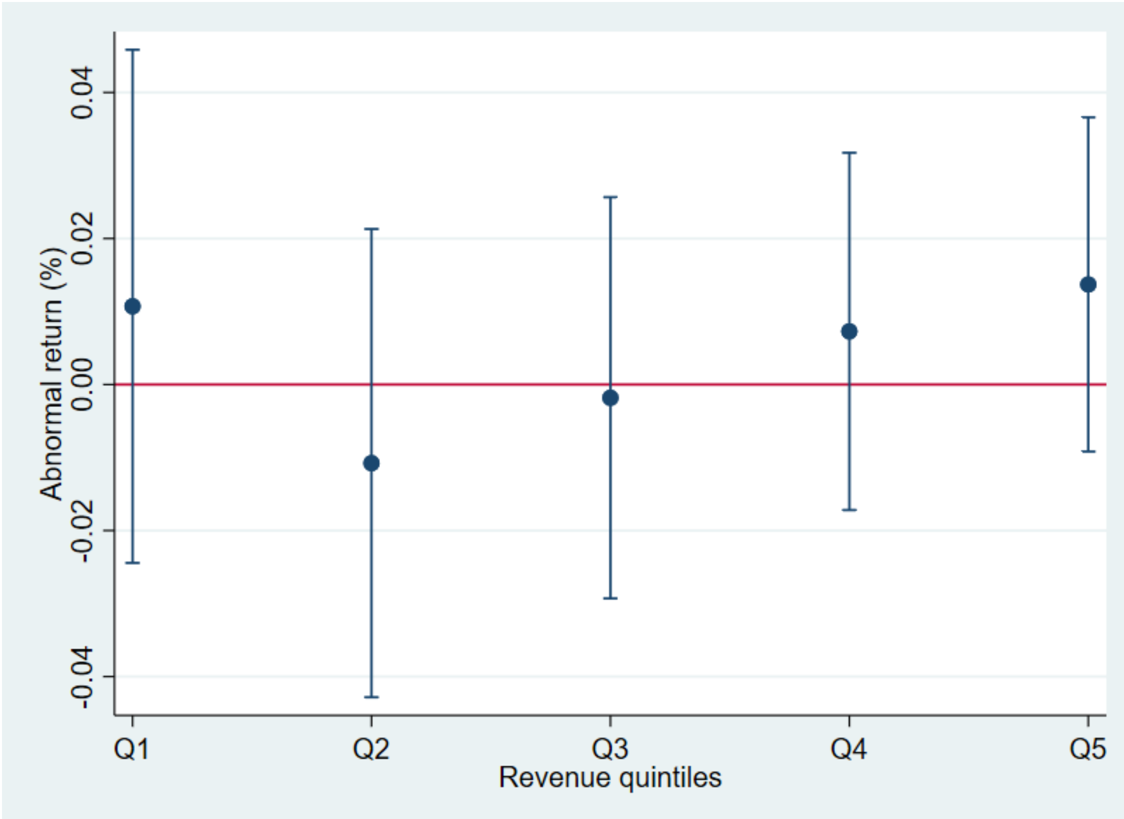

Figure 5: Mean abnormal returns by patenting firm revenue

Note: Events are binned into five approximately equally-sized groups based on the patenting firm's revenue in the fiscal year prior to the event. Revenue bins are defined within year, such that the bins contain approximately one-fifth of the observations in each year. 95% confidence intervals for the mean abnormal return over the $[t, t+3]$ window are given.

Firm size might alternatively be measured by the revenue of the firm. This measure of firm size has been used previously in work explaining the cross-sectional variation of abnormal returns in response to patent litigation events (Bessen et al. 2018). As far as this data set is concerned, defining firm size by revenue imposes substantial limitations, since revenue data typically comes from Compustat, which does not have data prior to 1950. Furthermore, a substantial number of publicly-traded firms do not have revenue data. From 1951 onwards, this missing revenue data affects approximately 6% of observations. To examine whether revenue is related to the abnormal returns of patent grant events, I construct revenue quintiles, defined relative to other firms receiving a patent in the same year as the observation. Patent grant events experienced by firms that did not have revenue data in the prior fiscal year are omitted. Figure 5 plots mean abnormal returns of firms across the revenue quintiles. Mean abnormal returns for all quintiles are not significantly different from zero. This holds true for the lowest quintile, which contrasts with the result for small firms when size is defined by market capitalization.

3.4 CITATIONS

Figure 2 suggests some evidence of a relationship between citations and abnormal returns of patent grant events. In order to examine this relationship more closely, I estimate the following equation:

$$AR_j = \alpha + \alpha_1 \log((C_j + 1)/N_j) + bZ_j + e_j \tag{4}$$

where AR_j is the four-day abnormal return of event j over the window $[t, t+3]$, t being the day of patent grant, C_j is the total number of forward citations of patents granted on that day, and N_j is the number of patents granted to the firm on that day. The total number of citations is incremented by one to avoid undefined logarithms in the case of grant events associated with patents that were not cited by other patents. Total citations are divided by the number of patents granted in order to avoid conflating events involving large citations counts among few patents granted with events involving small citation counts among many patents granted. Z_j is a vector of controls, which vary among specifications, but can include size quintiles defined by revenue or alternatively by market capitalization, which control for variation in how patents are valued against the total value of the firm and variation in ability to exploit patents by firm size; industry effects defined by three-digit SIC codes, which control for variation in the importance of patents across industries; firm effects, which control for time-invariant firm effects, and year effect, which control for variation in patent values over time.

Table 3 presents estimates of equation (4). The first column presents results without any controls, and suggests that larger citation counts are associated with modest, although statistically significant, increases in the abnormal returns of patent grant events. A 10% increase in the average number of forward citations of the patents granted in the event is associated with a 0.002 percentage point average increase in the cumulative abnormal return associated with the event.

The second column presents estimates of equation (4) with size quintile controls, defined as market capitalization relative to the market capitalization of other firms patenting in the same year. The citations estimate suggests a similar small, positive association between citation counts and abnormal returns. A 10% increase in the average number of forward citations of the patents granted in the event is associated with a 0.003 percentage point average increase in the cumulative abnormal return associated with the event. Coefficients on the size quintiles are in line with basic cross-sectional variation of abnormal returns with market capitalization, as documented in Figure 4. The smallest quintile of firms, which are the omitted reference category, have signficantly larger abnormal returns than firms belonging to the other quintiles. Although the largest difference is between the smallest quintile of firms and the second-smallest quintile of firms, the relationship between size and

Table 3: Cross-sectional analysis of abnormal returns

	(1)	(2)	(3)	(4)	(5)	(6)
$\log((C_j + 1)/N_j)$	0.020**	0.023**	0.019*	0.019*	0.010	-0.001
	(0.007)	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)
Size quintile (market cap)						
2		-0.195***				
		(0.028)				
3		-0.232***				
		(0.026)				
4		-0.272***				
		(0.025)				
5		-0.280***				
		(0.025)				
Size quintile (revenue)						
2			-0.022	-0.021	-0.022	0.042
			(0.029)	(0.030)	(0.030)	(0.035)
3			-0.012	-0.004	-0.006	0.078
			(0.028)	(0.029)	(0.030)	(0.040)
4			-0.004	0.012	0.010	0.082
			(0.027)	(0.029)	(0.030)	(0.045)
5			-0.014	0.004	0.003	0.091
			(0.027)	(0.030)	(0.031)	(0.049)
Missing revenue			0.054	0.071	0.083	0.153*
			(0.040)	(0.043)	(0.044)	(0.068)
Industry effects	N	N	N	Y	Y	N
Firm effects	N	N	N	N	N	Y
Year effects	N	N	N	N	Y	Y
N	585,510	585,510	538,776	538,776	538,776	538,776

Note: All models estimate the relationship between abnormal returns produced by the market model with day fixed effects $AR_i = R_i - (\beta_0 + \beta_1 R_m + \sum_{i=1}^{6} \alpha_i DAY_i + \sum_{i=1}^{6} \gamma_i DAY_i R_m)$, and log average citations of patents granted $\log((C_j + 1)/N_j)$. Coefficient estimates are scaled to reflect percentage point effects on abnormal returns. Market capitalization is the market capitalization of the firm at the end of trading the day prior to the event. Revenue is the prior fiscal year's revenue of the firm. Quintiles are defined within each year, so that a given quintile contains approximately 20% of the observations in each year. For columns 3 through 6, observations prior to 1950 are dropped, since there is no revenue data. Observations for which revenue is missing after 1950 are identified with the missing revenue dummy. Robust standard error are in parentheses. Asterisks denote significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

abnormal return appears to be monotonic, with average abnormal returns decreasing as size increases.

The third column presents estimates of equation (4) with size quintiles controls, defined instead as revenue relative to the revenue of other firms patenting in the same year. For columns involving size quintiles defined by revenue, observations prior to 1951 are dropped, as no firm has revenue data in Compustat in this time period. This results in the loss of $32,851$ observations. A variable identifying observations that are missing revenue data from 1951 onwards is also included, to proxy for any differences between firms that are and are not missing revenue data that may be relevant to patenting. The citations estimate suggests a similar association between citation counts and abnormal returns as the first two specifications, although standard errors are larger and the significance is more questionable. A 10% increase in the average number of forward citations of the patents granted in the event is associated with a 0.002 percentage point average increase in the cumulative abnormal return associated with the event. Relative to the market cap-size quntiles, revenue-size quintiles do not have a clear relationship with abnormal return. Although the coefficient estimates are again larger for the smallest quintile than any of the other quintiles, this difference is minor and not statistically significant.

The fourth column presents estimates of equation (4) with revenue-size quintile controls and industry effects. Including industry effects does not affect the citations estimate. Size coefficients show no clear relationship between size and abnormal return in this specification.

The fifth column presents estimates of equation (4) with year effects added. Including year effects reduces the citations estimate, which remains postive but is no longer statistically significant. As in the fourth specification, the size coefficients do not show any clear relationship between size and abnormal return.

The sixth specification uses firm effects instead of industry effects. The estimate of the citations coefficient is insignificant and negative. This indicates that increased citation counts do not have a substantial impact on the abnormal return of the patent grant event. The size coefficients are increasing in size quintile, but these differences are not significant. The coefficient on the missing revenue dummy is positive and significantly different from zero. Since firm effects are present in the regression and since the revenue variable uses the prior fiscal year's revenue, the missing revenue dummy might select for early periods in a firm's life when it experiences particularly large abnormal returns around patent grant events.

CHAPTER 4

PATENT VALUE ESTIMATION

The abnormal returns studied in the previous section are quite noisy, and a plausible explanation for the findings of a null overall effect and little meaningful variation is that patent values and their cross-sectional variation is obscured by this noise. In order to produce more informative estimates of patent values, Kogan et al. (2017) develop an estimation methodology for separating the patent value signal from the noise in abnormal returns estimated by event study. Analyzing the same data set discussed in the previous section, Kogan et al. find a highly-significant relationship between forward citations and their estimated patent values, estimating that a 10% increase in citations increases the value of a patent by 0.04% to 1.7%. This is clearly a substantially different result from the lack of a relationship between abnormal returns and citations. To better understand this discrepancy, I review the Kogan et al. methodology and conduct a placebo test in this section.

4.1 OVERVIEW

The method is designed to estimate the value of individual patents, based on the stock market returns of the patenting firm around the time of patent grant. Patent grants are assumed to affect the value of the patenting firm over a three day window $[t, t+2]$, where t is the day of patent grant. Abnormal returns are modelled as market-adjusted returns, and the observed market-adjusted three-day return R_j (as a % of firm value) is assumed to have a patent grant component v_j and an uncorrelated noise component ϵ_j.

$$R_j = v_j + \epsilon_j \tag{5}$$

v_j is relevant to estimating patent value, but it is unobserved. Instead, R_j is observed. The value of patent j is therefore estimated as:

$$\xi_j = (1-\bar{\pi})^{-1}\frac{1}{N_j}E[v_j|R_j]M_j \tag{6}$$

where $\bar{\pi}$ is the ex-ante probability of a successful patent grant. It is included to accomodate the anticipation of a successful patent grant. If it is known that a patent application is pending, and that a pending application will be successful with some non-zero probability, then the effect of a patent grant on the patenting firm's stock will not be the full value of the patent, but moderated by that probability. Kogan et al. set $\bar{\pi}$ equal to 0.56, following research by Carley, Hegde, and Marco (2015) showing that 56% of applications filed between 1996 and 2005 resulted in patent grants. N_j is the number of patents granted to patent j's assignee on the same grant day. This method cannot distinguish between the value of multiple patents granted to the same firm on the same day, but attributing the patent-related value accruing to the firm to each of multiple patents will result in double-counting. M_j is the patenting firm's market capitalization on the day prior to grant. $E[v_j|R_j]$ is the conditional expectation of the patent grant component given the noisy observation of market-adjusted return R_j.

Distributional assumptions are required to calculate $E[v_j|R_j]$. For their main estimation, Kogan et al. assume that the patent value component is distributed truncated normal, such that it is supported over the positive half of a normal distribution centered at zero. The noise component is assumed to be uncorrelated and normally-distributed.

$$\epsilon_j \sim N(0, \sigma^2_{\epsilon ft})$$

$$v_j \sim N^+(0, \sigma^2_{vft})$$

f and t subscripts allow for variation in these parameters by firm and year, although practically, Kogan et al. allow $\sigma^2_{\epsilon ft}$ to vary by firm-year, but require σ^2_{vft} to satisfy a ratio elaborated upon further below. The desired conditional expectation is then given by the following:

$$E[v_j|R_j] = \delta_{ft} R_j + \sigma_{\epsilon ft}\sqrt{\delta_{ft}} \frac{\phi(-\sqrt{\delta_{ft}}\frac{R_j}{\sigma_{\epsilon ft}})}{1 - \Phi(-\sqrt{\delta_{ft}}\frac{R_j}{\sigma_{\epsilon ft}})} \tag{7}$$

$$\delta_{ft} = \frac{\sigma^2_{vft}}{\sigma^2_{\epsilon ft} + \sigma^2_{vft}}$$

δ_{ft} may be understood as a signal-to-noise ratio, the proportion of return variance around patent grant event dates that is attributed to variance in patent value. To make the estimation problem feasible, Kogan et al. fix δ_{ft} to a benchmark signal-to-noise ratio δ, and estimate $\sigma^2_{\epsilon ft}$ for each firm f and year t using average return volatility. To identify δ, Kogan et al. regress return variance against a patent grant event dummy,

$$\log(R_{fd}^2) = \gamma D_{fd} + cZ_{fd} + u_{fd} \tag{8}$$

where R_{fd} is the three-day market-adjusted return of firm f starting on day d, D_{fd} is a dummy variable equal to one if firm f received a patent grant on day d, and Z_{fd} is a vector of controls, including firm-year fixed effects and day-of-week fixed effects. From $\hat{\gamma}$ the signal-to-noise ratio can be recovered as $\hat{\delta} = 1 - e^{-\hat{\gamma}}$. The noise variance parameter ϵ_{ft} is estimated as

$$\sigma_{\epsilon ft}^2 = 3\sigma_{ft}^2(1 + 3d_{ft}(e^{\hat{\gamma}} - 1))^{-1}$$

where σ_{ft}^2 is the mean of squared market-adjusted returns of firm f in year t and d_{ft} is the proportion of trading days that are patent grant days for firm f in year t. That adjustment is made to recover noise variance from observed return variance, which will by assumption also include patent value variance within patent grant event windows.

4.2 SIGNAL-TO-NOISE RATIO

This section will focus on the estimation step where signal-to-noise ratio δ is fixed to a single value. Fixing the signal-to-noise ratio simplifies the estimation, but may lead to misleading value estimates. A fixed signal-to-noise ratio will cause the implied patent value variance parameter to vary with the noise variance estimates. The noise variance estimates measure return volatility not attributed to the patent grant, so variation in estimated patent values may be driven by variation in return volatility. There is no direct evidence that patent values covary in such a way with the volatility of the firm's stock value.

Kogan et al. do suggest that allowing the signal-to-noise ratio to vary between firm size quintiles does not produce significantly different regression estimates except in the smallest quintile. A different estimate for the lowest firm size quintile, however, may be relevant to research into the relationship between firm size and patent value. Furthermore, if the regression estimates of equation (8) are noisy, it is difficult to conclude that there are no economically meaningful differences in the signal-to-noise ratio between firms in the various size quintiles, just because of statistical insignificance.

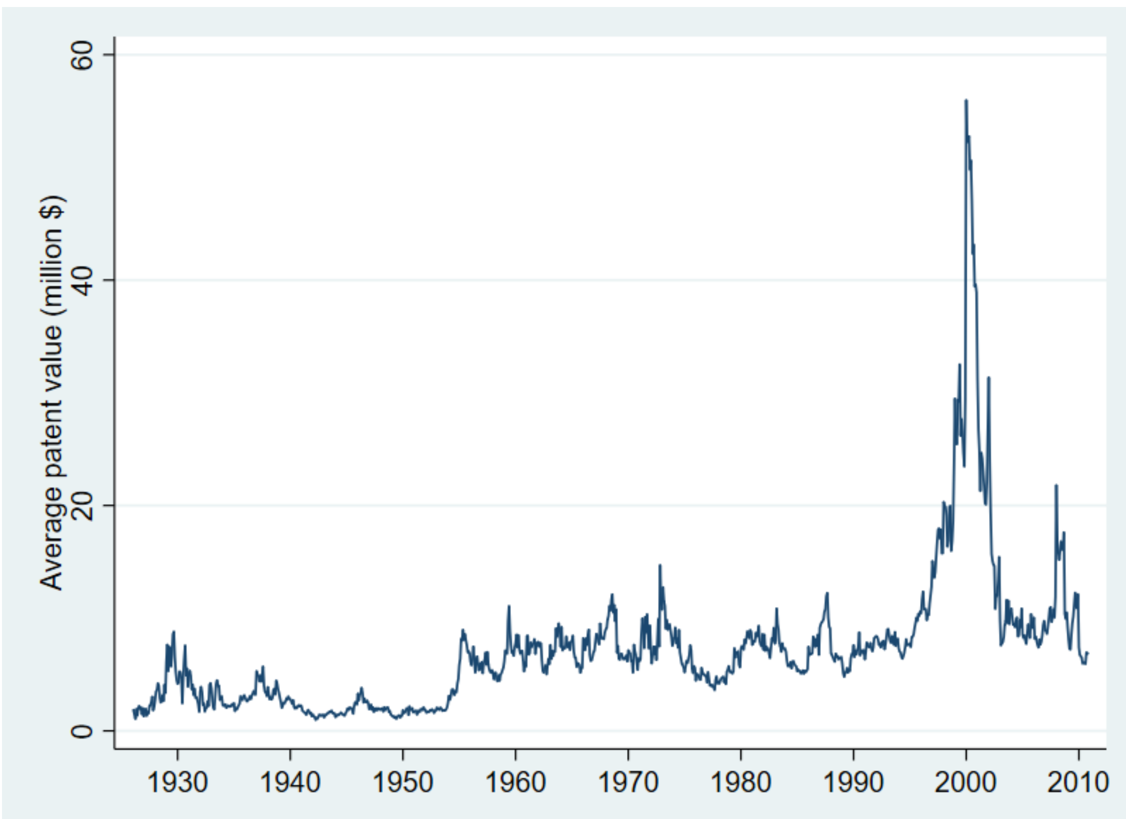

Figure 6: Average patent value time series

Note: Monthly averages of Kogan et al. estimates of patent value are provided.

I focus on a different source of variation in the signal-to-noise ratio, temporal variation. Firms may experience periods of especially high return volatility due to market-wide conditions, which is unlikely to be driven by patent grants to particular firms. Since the signal-to-noise ratio is fixed across firms and time, the same proportion of return variance will be attributed to patent grants in these volatile periods as in more stable periods. This may cause high patent value estimates to simply reflect the fact that they were granted to firms in these periods of volatility, rather than an actual higher patent value.

Figure 6 shows the monthly averages of the Kogan et al. patent value estimates over the timeframe of the Kogan et al. data set. There are local peaks in 1929, 2000, and 2008 which coincide with periods of high volatility in the stock market. This suggests that temporal variation in return volatility may contribute substantially to the variation in patent value estimates.

Firms are only represented in the data set if they receive at least one patent, so their returns may behave differently from that of the market as a whole. Figure 7 shows annual averages of return volatility, as measured by annual mean squared market-adjusted firm returns σ^2_{ft}, for the firms represented in the data set. There are local peaks in 1932, 2000, and 2008, which coincide with the high average patent value estimates, except in the case of the late 1920s/early 1930s, where the peak in average return volatility occurs about three years after the peak in average patent value.

In order to study how event-induced return volatility varies over time, I replicate the squared return regression for observations within each year, producing separate estimates of $\hat{\gamma}$ from equation (8) for each year. Figure 8 displays implied δ_t estimates for each year

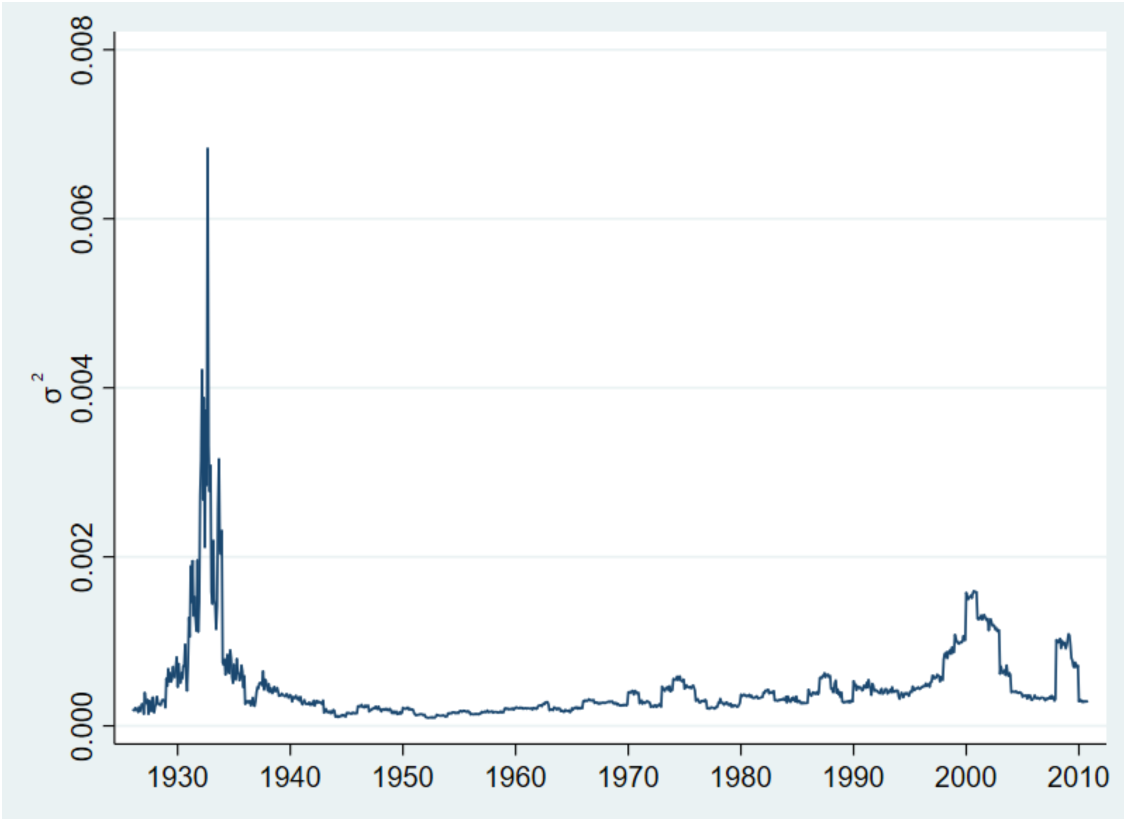

Figure 7: Average volatility time series

Note: Monthly averages of Kogan et al. volatility estimates for patenting firms are provided. Volatility is estimated for firm-years, but a firm's estimates is only included if it patented during that month, hence the variation between months.

from 1926 to 2010. The benchmark value of δ_t used for the entire data set, from Kogan et al., is marked in red. Estimates vary substantially from year to year, with early years' estimates appearing to vary more. The latter may simply reflect the fact that these are noisy estimates, and estimates of $\hat{\gamma}$ are made using fewer observations in earlier years, since there are far fewer patent grants, and by extension, fewer trading days of patenting firms.

For some of the years, $\hat{\gamma}$ is negative. Since the signal-to-noise ratio δ is taken as $1 - e^{-\hat{\gamma}}$, a negative $\hat{\gamma}$ implies a negative signal-to-noise ratio. A negative signal-to-noise ratio is incompatible with the assumptions of the estimation methodology, as it represents the ratio between two variances, and cannot be used to estimate equation (7). $\hat{\gamma}$ is negative in 38 of

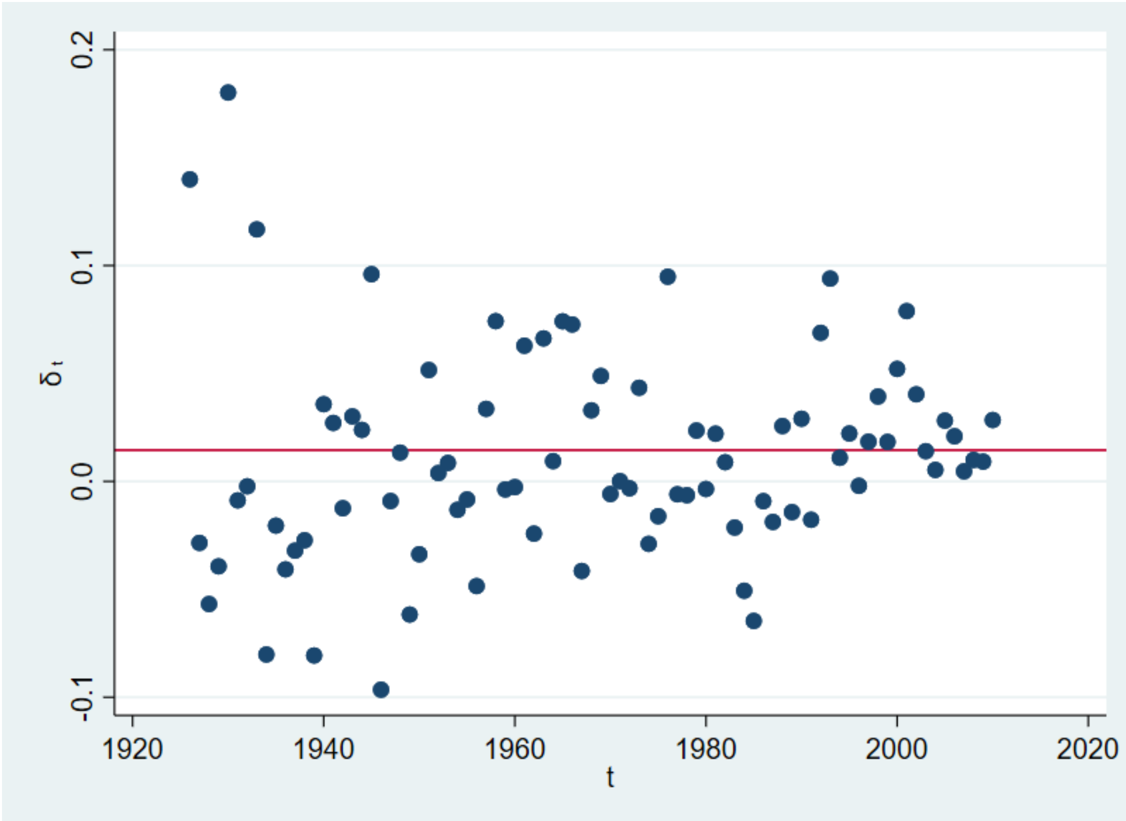

Figure 8: Time-varying signal-to-noise ratio

Note: Estimates of $\delta_t = 1 - e^{-\hat{\gamma}_t}$, where $\log(R^2_{fd}) = \gamma_t D_{fd} + cZ_{fd} + u_{fd}$ is estimated separately for each year, are provided. The aggregate estimate $\delta = 0.0145$ is marked with a red line.

the 85 years of data, so these estimates do not permit further calculation of patent value estimates.

Fixing signal-to-noise ratio to a single value is a restrictive assumption that may lead to patent value estimates for which cross-sectional variation is driven by the firm's volatility around the time of grant. However, steps to relax this assumption, such as allowing signal-to-noise ratio to vary by year results in noisy estimates of the return variance regression, a substantial fraction of which are negative, which prevents estimation of patent values.

4.3 PLACEBO TEST

In order to test the external validity of the Kogan et al. estimation methodology, I perform a placebo test. I reassign patent grant dates to other days in the same year, with patents granted on the same day reassigned together, and perform the estimation as described in section 4.1.[7] This is repeated 100 times. Distributions of key variables in the Kogan et al. estimation across iterations are plotted below.

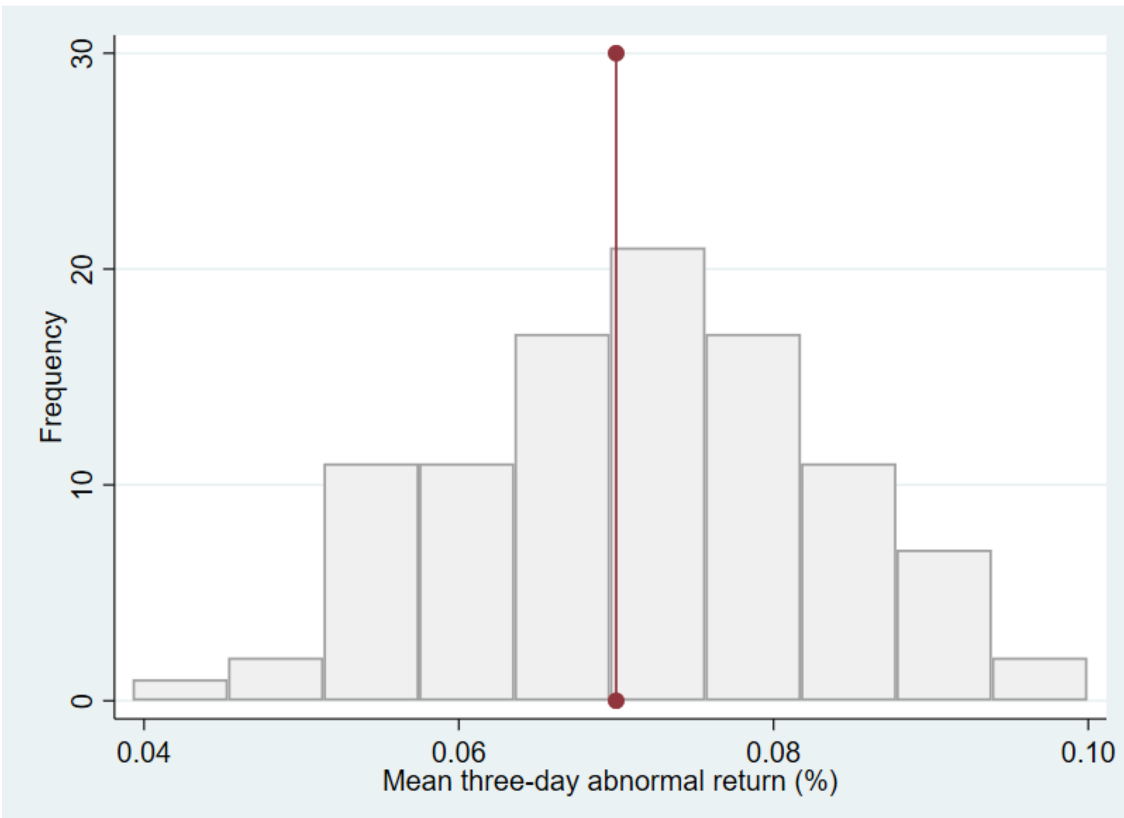

Figure 9: Mean market-adjusted returns of placebo events

Note: The distribution of the mean market-adjusted return, across iterations of the placebo test, is provided. The actual mean market-adjusted return 0.07% is marked with a red line.

[7] The goal of reassigning patents together is to prevent random reassignment from increasing the total number of patent grant days, which may easily happen if a firm received many patents in the same year. In the case of particularly prolific patentees, completely random reassignment may cause all days to become patent grant days. This would effectively remove a very selected group of firms from the signal-to-noise ratio estimation.

Figure 9 plots the mean market-adjusted returns over the three-day window $[t, t+2]$ for placebo event dates. The actual mean market-adjusted return, 0.07%, is marked in red. The mean placebo returns are similar to the actual market-adjusted return, falling between 0.03% and 0.10%, indicating that, in about half of placebo iterations, market-adjusted returns on placebo dates are as large or larger than market-adjusted returns on the actual patent grant dates. Since placebo dates are often not patent grant days, this indicates that returns on patent grant days are not especially large relative to non-patent grant days. Additionally, the positive means of the placebo market-adjusted returns provides additional evidence that the market-adjusted returns model underestimates the normal returns of firms in the data set.

As part of the placebo test, the signal-to-noise ratio δ is estimated for the placebo events in each iteration. Figure 10 plots the signal-to-noise ratio for the placebo events. The actual $\delta = 0.0145$ used by Kogan et al. is marked in red. The placebo events tend to produce smaller δ, which are approximately centered on zero, and negative in 41 interations. This is expected since placebo events have random dates. In six of the hundred iterations, however, the placebo δ is larger than the actual δ. Since the placebo tests using random event dates produce δ as large or larger than the actual δ with this frequency, it is unclear whether the actual patent grant events produce a significant increase in the volatility of firm returns.

Negative δ in the placebo tests presents a practical problem to further estimation. As can be seen in equation (7), $E[v_j|R_j]$ cannot be estimated if δ is negative. In the placebo test iterations producing negative δ, patent values therefore cannot be estimated. For the 59 iterations in which δ is positive, patent values are estimated and plotted in Figure 11. The

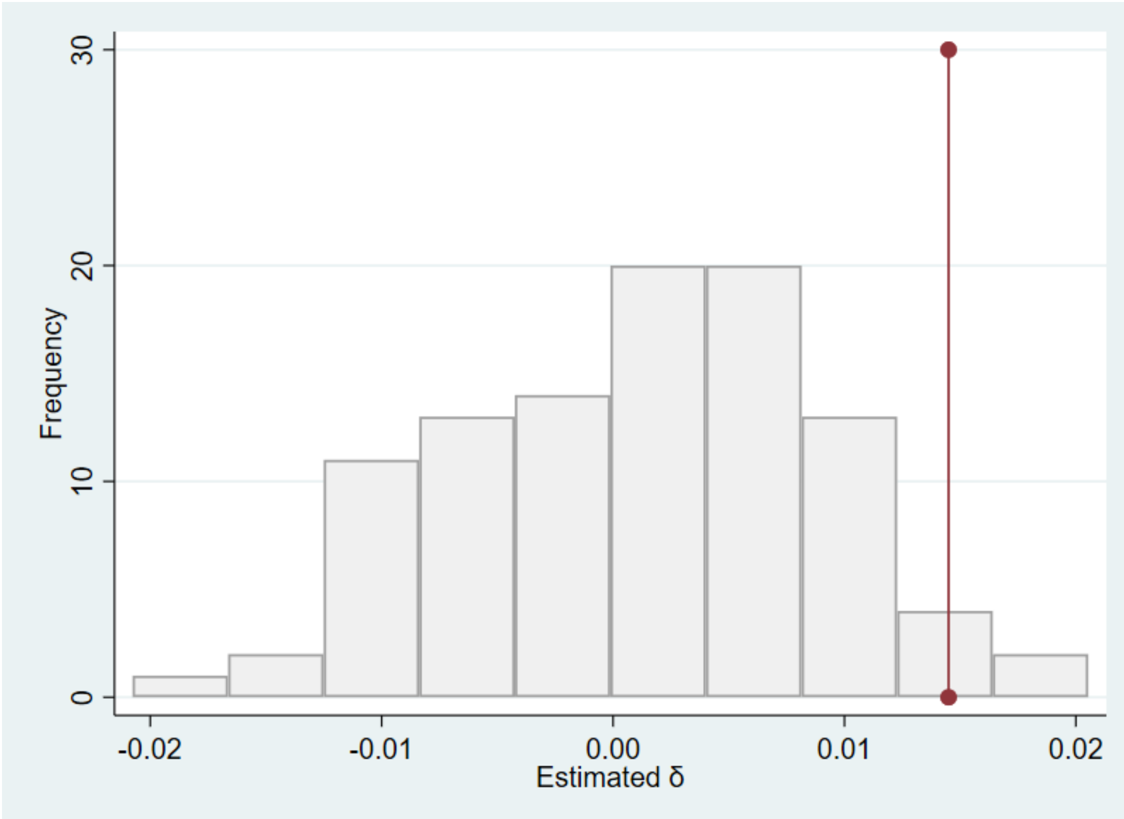

Figure 10: Signal-to-noise ratio of placebo events

Note: Estimates of $\delta = 1 - e^{-\hat{\gamma}}$, where $\log(R^2_{fd}) = \gamma D_{fd} + cZ_{fd} + u_{fd}$ is estimated by regression, are provided. The distribution represents the frequency of δ estimates across iterations of the placebo test. The actual signal-to-noise ratio $\delta = 0.0145$ is marked with a red line.

mean of the actual patent values, \$10.36 million, is marked in red. Placebo patent values tend to be smaller, but in six iterations, the placebo mean is as large or larger than the actual mean patent value. These iterations are the same iterations in which placebo δ is as large or larger than the actual δ. In fact, larger δ is closely associated with larger mean value estimate among the 59 iterations.

The distribution of placebo market-adjusted returns confirms that the market return underestimates the normal return of patenting firms. Market-adjustment produces postive abnormal returns, even on days that are not necessarily associated with a patent grant event. Placebo tests also suggest that the estimation of signal-to-noise ratio δ is noisy, and

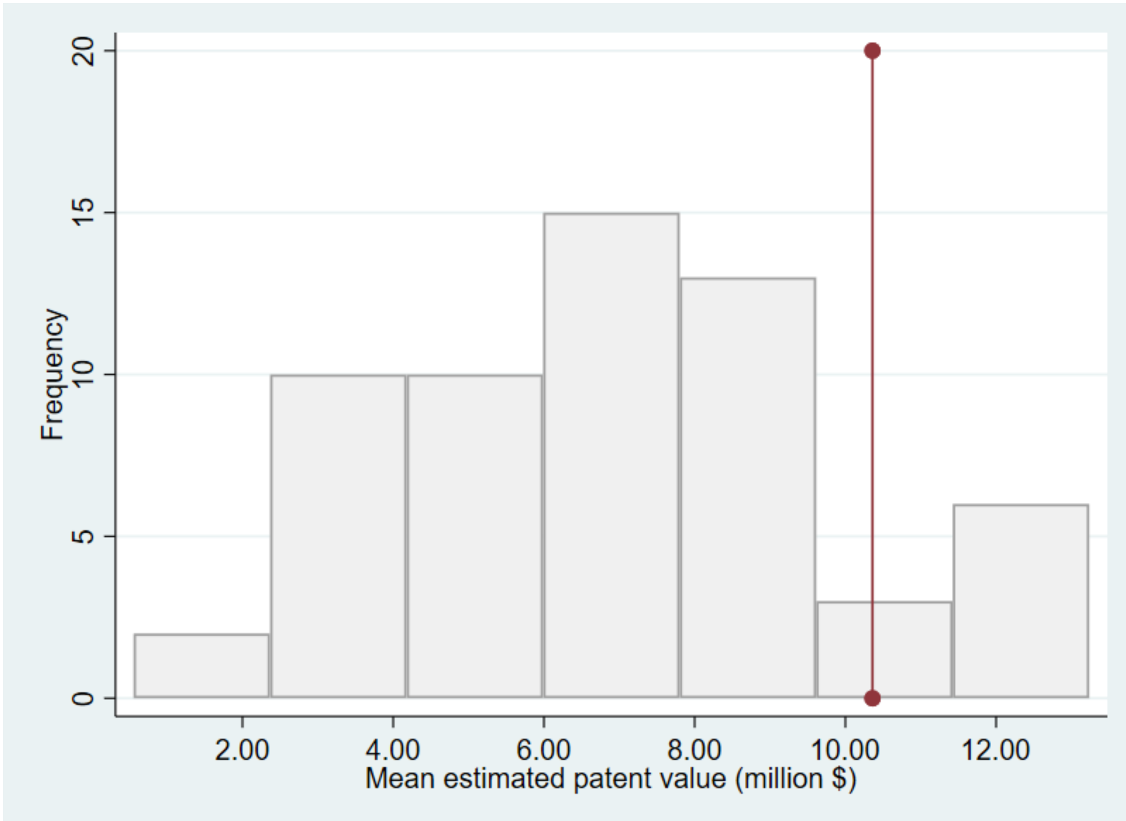

Figure 11: Mean value estimates based on placebo dates

Note: The mean Kogan et al. patent value estimate across iterations of the placebo test, where the full procedure is repeated for each iteration of the placebo test, are provided. The actual average Kogan et al. patent value estimate, $10.36 million, is marked with a red line. Total frequency is less than a hundred despite a hundred iterations being run because patent values cannot be estimated if δ estimates, reported in Figure 10 are negative.

the actual signal-to-noise ratio might not represent an actual significant increase in return volatility due to patent grants. The placebo events may produce a signal-to-noise ratio of the same size and sign as the Kogan et al. δ spuriously. As δ is a crucial input to the patent value estimates, the placebo events can produce patent value estimates of a similar size.

While these statistics of the Kogan et al. estimation are replicated approximately in the placebo tests, this does not speak to whether variation between the patent value estimates is informative. Kogan et al. make such a point, demonstrating that cross-sectional variation in their patent values has a close positive association with patent citations. Specifically, they estimate the following equation,

$$\log(\xi_j) = a + b \log(1 + C_j) + cZ_j + u_j \tag{9}$$

where ξ_j is their value estimate for patent j, C_j is the number of forward citations of patent j and Z_j is a vector of controls, which in their fifth specification[8], includes controls for firm size measured as the log of the market capitalization of the patenting firm one day prior to the patent grant date, fixed effects for patent technology class interacted with year, and fixed effects for firm interacted with year.

Since the outcome variable of equation (9) is the patent value estimates, $\hat{b}$ can only be estimated in the 59 iterations of the placebo test in which patent values were estimated. Coefficient estimates are plotted in Figure 12. The actual coefficient estimate, 0.004, is marked in red. Placebo estimates of the citations coefficients are quite close to the actual estimate. Although the magnitude of the actual estimate is larger than all but two of the placebo estimates, all placebo estimates are within rounding error of the actual estimate. This suggests that even when the abnormal returns attributed to the patent grant event are simply drawn from some random date in the year, the Kogan et al. method produces similar estimates of the citations coefficient. It is unlikely, then, that the relationship found between

[8]The choice of this specification is somewhat arbitary, although it is the specification Kogan et al. select for comparison in their own placebo test.

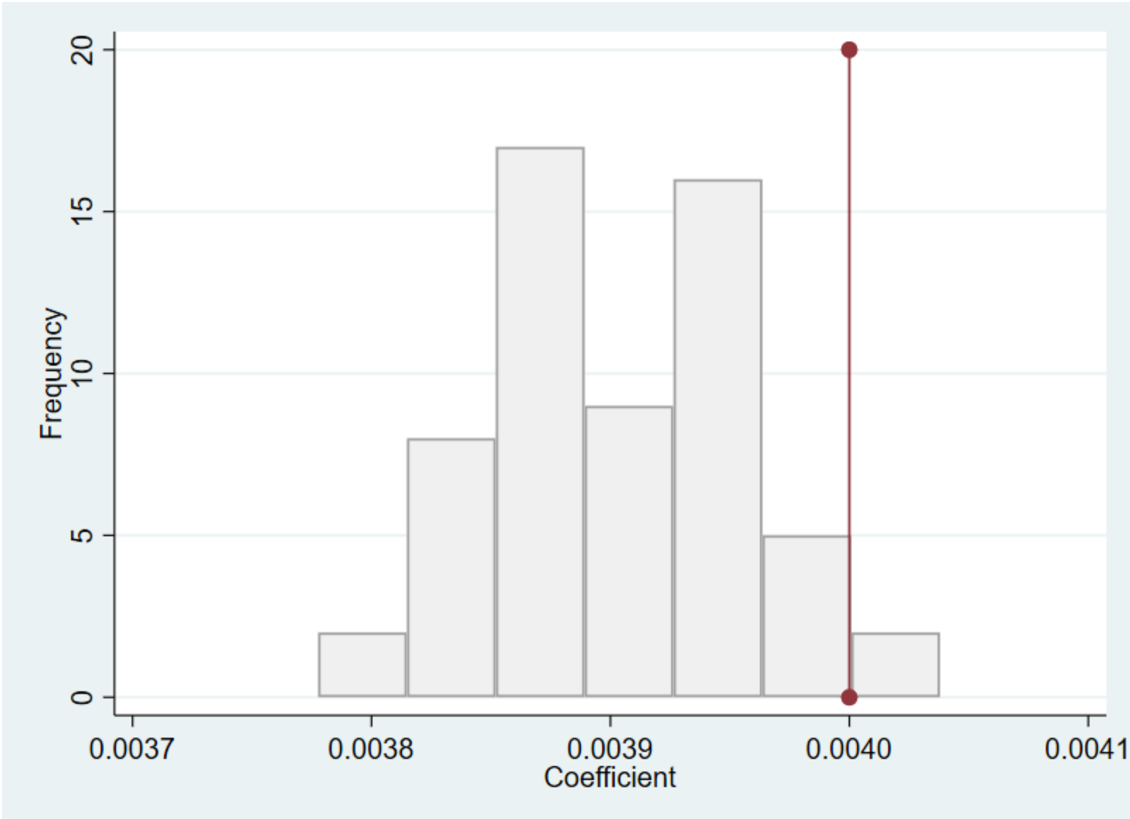

Figure 12: Citations coefficient estimates based on placebo dates

Note: Plotted are coefficient estimates $\hat{b}$ from the equation $\log(\xi_j) = a + b\log(1 + C_j) + cZ_j + u_j$ across placebo iterations. The actual coefficient estimated by Kogan et al., 0.004, is marked with a red line. Total frequency is less than a hundred despite a hundred iterations being run because patent values cannot be estimated if δ estimates, reported in Figure 10 are negative.

citations and the real estimated patent values is due to effective separation of signal from noise in the abnormal returns; a similar relationship is found when events are purely noise.

A plausible explanation for why the placebo estimates are similar to the actual estimate is endogeneity in the number of patents granted. The citation count of a patent has a strong negative relationship with the number of patents granted, and this relationship is not controlled for by any of the covariates in the specification. The number of patents granted to the same firm on the same day as patent j, N_j, is a divisor in equation (6), so larger N_j mechanically reduces the estimated patent value. Since the number of patents granted is negatively correlated with citation count, and the number of patents granted is negatively correlated with the estimated patent value, the estimated coefficient on citations is biased upwards. Adding a control for the log of the number of patents granted to the specification reduces the estimate of the citations coefficient to approximately zero. Recalling the randomization procedure used in this placebo test, grant dates are randomized, but patents granted on the same day receive the same randomized date. As such, N_j holds constant across iterations of the placebo test, even though the grant date of patent j is randomized, introducing a similar relationship between citations and placebo estimates..

CHAPTER 5

CONCLUSION

Using an event study framework, I analyze the effect of patent grant events on the value of patenting firms. For patent grant events, the question of how to model abnormal returns is critical. Because patenting firms are substantially different from the market as a whole, a market-adjusted returns model underestimates the normal returns of patenting firms, introducing upward bias in the estimated abnormal returns. Because patent grant events typically occur on Tuesdays, estimates of abnormal returns attributable to patent grant events may be significantly confounded by day-of-week effects present in the returns of patenting firms. These biases are significant. Using two commonly-applied abnormal returns models, the market-adjusted returns model and the market model, results in misleading significant, positive estimates of the impact of patent grant events on firm value.

Applying a market model with day-of-week effects, I find little evidence of significant average effects across the broadest data set compiled to date. Analysis of heterogeneity in abnormal returns by firm size, patenting frequency, and forward citations yields little evidence of meaningful cross-sectional variation in abnormal returns. Patenting frequency, which might proxy for the degree to which patent grant events are unanticipated news for the firm, is not associated with significant variation in abnormal returns. Forward citations are not significantly related to abnormal returns when year effects are controlled for. Analysis of heterogeneity in abnormal returns by firm size does suggest that small firms appear to experience larger abnormal returns. In particular, the smallest decile of firms, defined by market capitalization, experiences significant mean abnormal returns of 0.35%. However, since this relationship does not hold when size is instead defined by the firm's revenue, it seems likely that this merely a mechanical relationship. Individual patent grants or even patent grant events have only relatively minor impacts on the value of firms, and this impact is only discernible is among small firms. In both the overall analysis and the analysis of heterogeneity, a principle challenge is the noise in abnormal returns.

Turning to methodology that has recently been developed by Kogan et al. (2017) for the purpose of addressing separating patent value from the noise in abnormal returns, I examine the underlying assumptions and conduct a placebo test. One strong assumption is that of a fixed ratio between the variance in patent values and the variance in unrelated return noise. Not allowing this ratio to vary over time may lead one to estimate larger patent values in a particular time period merely because the market is more volatile in the same time period. Empirically, the Kogan et al. patent value estimates spike during periods in which patenting firms experienced large amounts of return volatility.

In conducting my placebo test, I find that similar estimates of key statistics in the estimation method may be produced even when patent grant events are assigned randomized dates. Placebo estimates of the signal-to-noise ratio, which measures the importance of variance in patent values to the overall variance in returns around patent grant dates, are as large or larger than the real estimate with sufficient frequency as to call into question whether patent value variance actually contributes significantly to return variance. One particularly interesting difference between abnormal returns and Kogan et al. patent value estimates is that citations are not meaningfully related to abnormal returns but are meaningfully related to the patent value estimates. Replicating the estimation of that relationship, I find that citations have a similar relationship to placebo patent value estimates. A plausible explanation for why this occurs is endogeneity. The estimation method mechanically introduces a relationship between the number of patents granted and the patent value estimate, and the number of patents granted is negatively correlated with citation count.

REFERENCES

Abrams, David S., Ufuk Akcigit, and Jillian Grennan. 2018. “Patent Value and Citations: Creative Destruction or Strategic Disruption?” *University of Pennsylvania Institute for Law & Economics Research Paper.*

Allen, Paul R., and C.F. Sirmans. 1987. “An Anlaysis of Gains to Acquiring Firm’s Shareholders”. *Journal of Financial Economics* 18:175–184.

Andritzky, Jochen R., Geoffrey J. Bannister, and Natalia T. Tamirisa. 2007. “The Impact of Macroeconomic Announcements on Emerging Market Bonds”. *Emerging Markets Review* 8:20–37.

Asquith, Paul, and David W. Mullins. 1986. “Signalling with Dividends, Stock Repurchases and Equity Issues”. *Financial Managment* 15:27–44.

Austin, David H. 1993. “An Event Study Approach to Measuring Innovative Output: The Case of Biotechnology”. *American Economic Review* 83:253–258.

Bessen, James. 2009. “Estimates of Patent Rents from Firm Market Value”. *Research Policy* 38:1604–1616.

Bessen, James, et al. 2018. “Trends in Private Patent Costs and Rents for Publicly-Traded United States Firms”. *International Review of Law and Economics* 56:53–69.

Carley, Michael, Deepak Hegde, and Alan Marco. 2015. “What is the Probability of Receiving a U.S. Patent”. *Yale Journal of Law and Technology* 17.

Chen, Li-Yu, et al. 2018. “Organisational Form and the Economic Impacts of Innovation: The Case of Patent Grant Announcements”. *Technology Analysis & Strategic Management* 30:1073–1083.

Dornelles, Juliana, and Ayfer Ali. 2014. “Patent Value: Exclusivity or Signal of Research Productivity”. Paper presented at DRUID Society Conference.

Fama, Eugene F. 1965. “The Behavior of Stock-market Prices”. *Journal of Business* 38:34–105.

Fama, Eugene F., and Kenneth R. French. 1993. “Common Risk Factors in the Returns on Stocks and Bonds”. *Journal of Financial Economics* 33:3–56.

Fama, Eugene F., et al. 1969. “The Adjustment of Stock Prices to New Information”. *International Economic Review* 10:1–21.

French, Kenneth R. 1980. “Stock Returns and the Weekend Effect”. *Journal of Financial Economics* 8:55–69.

Gibbons, Michael R., and Patrick Hess. 1981. “Day of the Week Effects and Asset Returns”. *The Journal of Business* 54:579–596.

Hall, Brownwyn H., Adam Jaffe, and Manuel Trajtenberg. 2005. “Market Value and Patent Citations”. *RAND Journal of Economics* 36:16–38.

Hall, Brownyn H., Adam B. Jaffe, and Manuel Trajtenberg. 2001. “The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools”. *NBER Working Paper*.

Harhoff, Dietmar, Frederic M. Scherer, and Katrin Vopel. 2003. “Citations, Family Size, Opposition and the Value of Patent Rights”. *Research Policy* 32:1343–1363.

Kogan, Leonid, et al. 2017. “Technological Innovation, Resource Allocation, and Growth”. *Quarterly Journal of Economics* 132:665–712.

Lamdin, Douglas J. 1999. “Event Studies of Regulation and New Results on the Effect of the Cigarette Advertising Ban”. *Journal of Regulatory Economics* 16:187–201.

Louhichi, Wael. 2008. “Adjustment of Stock Prices to Earnings Announcements: Evidence from Euronext Paris”. *Review of Accounting and Finance* 7:102–115.

MacKinlay, A. Craig. 1997. “Event Studies in Economics and Finance”. *Journal of Economic Literature* 35:13–39.

McAfee, R. Preston, and Michael A. Williams. 1988. “Can Event Studies Detect Anticompetitive Mergers”. *Economics Letters* 28:199–203.

Nicholas, Tom. 2008. “Does Innovation Cause Stock Market Runups? Evidence from the Great Crash”. *American Economic Review* 98:1370–1396.

Pakes, Ariel, and Mark Schankerman. 1984. “The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources”. *R&D, Patents, and Productivity.*

Patel, Darshak, and Michael R. Ward. 2011. “Using Patent Citation Patterns to Infer Innovation Market Competition”. *Research Policy* 40:886–894.

Peress, Joel. 2014. “The Media and the Diffusion of Information in Finanical Markets: Evidence from Newspaper Strikes”. *Journal of Finance* 69:2007–2043.

Sellers-Rubio, Ricardo, Juan L. Nicolau-Gonzálbez, and Francisco Mas-Ruiz. 2007. “The Economic Value of Patent Protection and Rivalry in the Spanish Electrical Sector”. *European Journal of Innovation Management* 10:434–452.

APPENDIX A

DERIVATION OF EQUATION (7)

Equation (7) follows from the distributional assumptions. If x_1, x_2 are uncorrelated normal variables with $x_1 \sim N(\mu_1, \sigma_1^2)$ and $x_2 \sim N(\mu_2, \sigma_2^2)$, their sum $x_1 + x_2$ is distributed $N(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$. x_1 is clearly correlated with $x_1 + x_2$, and the conditional distribution $x_1 | x_1 + x_2$ is given by:

$$N(\mu_1 + \rho_{x_1, x_1+x_2} \frac{\sigma_1}{\sqrt{\sigma_1^2 + \sigma_2^2}} (x_1 + x_2 - \mu_2), (1 - \rho_{x_1, x_1+x_2}^2) \sigma_1^2)$$

The correlation between x_1 and the sum, ρ_{x_1, x_1+x_2} simplifies:

$$\begin{aligned} \rho_{x_1, x_1+x_2} &= \frac{\sigma_1^2}{\sigma_1 \sqrt{\sigma_1^2 + \sigma_2^2}} \\ &= \frac{\sigma_1}{\sqrt{\sigma_1^2 + \sigma_2^2}} \end{aligned}$$

Substituting this into the conditional distribution:

$$\begin{aligned} x_1 | x_1 + x_2 \sim& N(\mu_1 + \frac{\sigma_1}{\sqrt{\sigma_1^2 + \sigma_2^2}} \frac{\sigma_1}{\sqrt{\sigma_1^2 + \sigma_2^2}} (x_1 + x_2 - \mu_2), (1 - (\frac{\sigma_1}{\sqrt{\sigma_1^2 + \sigma_2^2}})^2) \sigma_1^2) \\ x_1 | x_1 + x_2 \sim& N(\mu_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2} (x_1 + x_2 - \mu_2), (1 - \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}) \sigma_1^2) \\ x_1 | x_1 + x_2 \sim& N(\mu_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2} (x_1 + x_2 - \mu_2), \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2} \sigma_1^2) \end{aligned}$$

Given that $\epsilon_j \sim N(0, \sigma_\epsilon^2)$, $v_j \sim N^+(0, \sigma_v^2)$ are uncorrelated, and R_j is the observed sum of these two variables, $v_j|R_j$ has the distribution similar to the general conditional form, but truncated on the left at zero. Let $\psi(\mu, \sigma^2, l, u)$ denote a truncated normal distribution, with μ, σ^2 the parameters of the corresponding normal distribution, and l, u the lower and upper truncation points, respectively. Then:

$$
\begin{aligned}
v_j|R_j &\sim \psi(0 + \frac{\sigma_v^2}{\sigma_v^2 + \sigma_\epsilon^2}(R - 0), \frac{\sigma_\epsilon^2}{\sigma_v^2 + \sigma_\epsilon^2}\sigma_v^2, 0, \infty) \\
v_j|R_j &\sim \psi(\frac{\sigma_v^2}{\sigma_v^2 + \sigma_\epsilon^2}R, \frac{\sigma_\epsilon^2\sigma_v^2}{\sigma_v^2 + \sigma_\epsilon^2}, 0, \infty)
\end{aligned}
$$

Truncation adds an Inverse Mill's Ratio term,

$$
E[v_j|R_j] = \frac{\sigma_v^2}{\sigma_v^2 + \sigma_\epsilon^2}R + \frac{\sigma_\epsilon\sigma_v}{\sqrt{\sigma_v^2 + \sigma_\epsilon^2}}\frac{\phi(\alpha)}{1 - \Phi(\alpha)}
$$

where α is l, standardized using the first two parameters of the distribution.

$$
\begin{aligned}
\alpha &= (0 - \frac{\sigma_v^2}{\sigma_v^2 + \sigma_\epsilon^2}R)/(\frac{\sigma_\epsilon\sigma_v}{\sqrt{\sigma_v^2 + \sigma_\epsilon^2}}) \\
&= \frac{-\sigma_v R}{\sigma_\epsilon\sqrt{\sigma_v^2 + \sigma_\epsilon^2}}
\end{aligned}
$$

Substituting α into the expectation and substituting $\delta = \frac{\sigma_v^2}{\sigma_v^2+\sigma_\epsilon^2}$ where appropriate leads to:

$$E[v_j|R_j] = \frac{\sigma_v^2}{\sigma_v^2 + \sigma_\epsilon^2}R + \frac{\sigma_\epsilon \sigma_v}{\sqrt{\sigma_v^2 + \sigma_\epsilon^2}} \frac{\phi(\frac{-\sigma_v R}{\sigma_\epsilon \sqrt{\sigma_v^2+\sigma_\epsilon^2}})}{1 - \Phi(\frac{-\sigma_v R}{\sigma_\epsilon \sqrt{\sigma_v^2+\sigma_\epsilon^2}})}$$

$$E[v_j|R_j] = \delta R + \sigma_\epsilon \sqrt{\delta} \frac{\phi(-\sqrt{\delta}\frac{R}{\sigma_\epsilon})}{1 - \Phi(-\sqrt{\delta}\frac{R}{\sigma_\epsilon})}$$